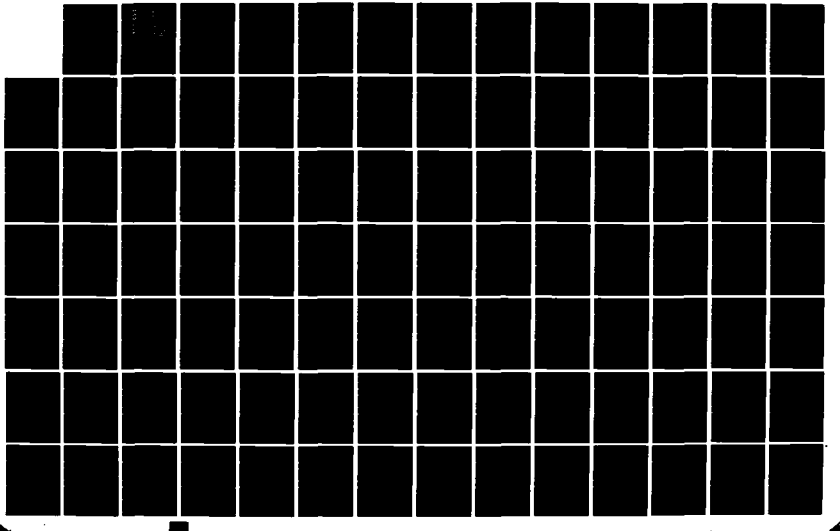
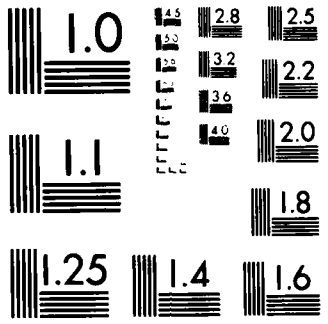
AD-A131 954 GENERALIZED MONITORING FACILITY USERS MANUAL CHANGE 6 1/8
(U) COMMAND AND CONTROL TECHNICAL CENTER WASHINGTON DC
B M WALLACK ET AL. 01 JUL 83 CCTC-CSM-UM-246-82-CHG-6
UNCLASSIFIED DCA/DF-83/001B F/G 9/2 NL





MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS 1963-A

AD A131954

WWMCCS

WORLDWIDE MILITARY COMMAND AND CONTROL SYSTEM

STANDARD SYSTEM SOFTWARE

DTIC FILE COPY

DEFENSE COMMUNICATIONS AGENCY

APPROVED FOR PUBLIC RELEASE:
DISTRIBUTION UNLIMITED

COMPUTER SYSTEM MANUAL

CSM UM 246-82



COMMAND & CONTROL TECHNICAL CENTER

GENERALIZED MONITORING FACILITY

DTIC
1983

USERS MANUAL

REPRODUCED BY
NATIONAL TECHNICAL
INFORMATION SERVICE
U.S. DEPARTMENT OF COMMERCE
SPRINGFIELD, VA 22161

83 08 10 041

COMMAND AND CONTROL TECHNICAL CENTER

Computer System Manual CSM UM 246-82

GENERALIZED MONITORING FACILITY

Users Manual, Change 6

1 July 1983

SUBMITTED BY:

BARRY M. WALLACK
GEORGE H. GERO, JR.
Project Officers

APPROVED BY:

for

ALAN GROSS
Colonel, USAF
Deputy Director for
Computer Services

50272-101

REPORT DOCUMENTATION PAGE	1. REPORT NO. DCA/DF-83/001b	2.	3. Recipient's Accession No.

4. Title and Subtitle

GENERALIZED MONITORING FACILITY, Users Manual, Change 6

5. Report Date July 1, 1983

6.

7. Author(s)

Barry M. Wallack, George H. Gero, Jr.

8. Performing Organization Rept. No. CSM UM 246-82

9. Performing Organization Name and Address

Defense Communications Agency/CCTC
The Pentagon, C751
Room BE 685
Washington, D.C. 20301

10. Project/Task/Work Unit No.

11. Contract(C) or Grant(G) No.

(C)

(G)

12. Sponsoring Organization Name and Address

(same)

13. Type of Report & Period Covered

14.

15. Supplementary Notes

For magnetic tape, see AD-A122 553

This document supersedes AD-A122 554

16. Abstract (Limit: 200 words)

This Users Manual provides the information necessary for an individual to run all of the programs that are comprised in the Honeywell H6000 Generalized Monitoring Facility (GMF). Program interrelationships are presented, as well as a general overview of the processing, input, and output procedures for each program. Formats and examples of user-controlled input data and sample program outputs are shown and explained. Additionally, the Job Control Language (JCL) deck setups necessary to run the program are provided.

17. Document Analysis a. Descriptors

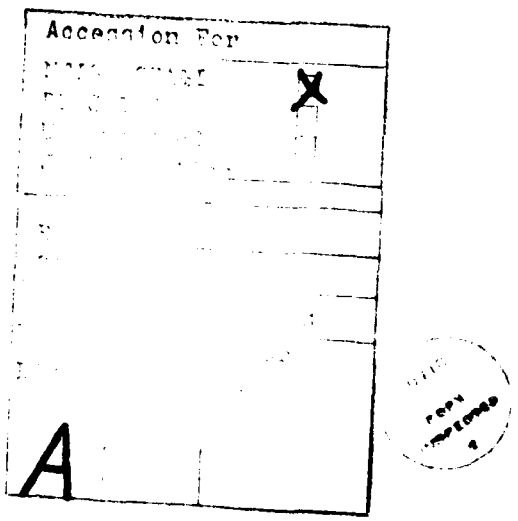
b. Identifiers/Open-Ended Terms

c. COSATI Field/Group

18. Availability Statement:	19. Security Class (This Report) UNCLASSIFIED	21. No. of Pages
	20. Security Class (This Page) UNCLASSIFIED	22. Price

ACKNOWLEDGMENT

This manual was prepared for the Computer Services Directorate of the Command and Control Technical Center (CCTC) under the direction of the Computer Performance Evaluation Office (C751).



CONTENTS

Section		Page
	ACKNOWLEDGMENT	ii
	ABSTRACT	xxviv
1.	GENERAL	1-1
1.1	Purpose of the Users Manual	1-1
1.2	Project References	1-1
1.3	Terms and Abbreviations	1-1
1.4	Security and Privacy	1-2
1.5	Manual Format	1-2
2.	SYSTEM SUMMARY	2-1
2.1	System Application	2-1
2.2	System Operation	2-1
2.2.1	The Resource Monitor Collection (RMC) Subsystem	2-1
2.2.2	The Generalized Monitor (GMC) Subsystem	2-2
2.3	System Configuration	2-2
2.4	System Organization	2-7
2.5	Performance	2-7
2.6	GMF Installation	2-7
2.6.1	Creation of GMF Files	2-7
2.6.2	GMF Release Dependent Parameters	2-10
3.	RESOURCE MONITOR COLLECTOR	3-1
3.1	Introduction	3-1
3.2	RMC Input Options and JCL	3-1
3.3	Processing	3-1
3.3.1	Collection Routine	3-1
3.3.1.1	Program Data - Subtype 1	3-3
3.3.1.2	Peripheral Data - Subtype 2	3-4
3.3.1.3	Scheduler Data - Subtype 3	3-4
3.3.1.4	Configuration Data - Subtype 4	3-5
3.4	Outputs	3-5
3.5	Catalog Description	3-5
3.5.1	Files	3-5
3.5.1.1	FILEDIT Source File	3-5
3.5.1.2	FILEDIT Object File	3-6
3.5.1.3	SYSEDIT Loadable File	3-6
3.5.1.4	JCL	3-6
3.5.1.4.1	DUPJCL	3-6
3.5.1.4.2	DALT	3-6
3.5.1.4.3	COLL	3-6
3.5.1.4.4	PSUMR	3-6
3.5.1.4.5	RPTSUM	3-6
3.5.1.4.6	CUPJCL	3-6
3.5.2	Update Procedures	3-6

Section		Page
4.	RESOURCE MONITOR DATA REDUCTION	4-1
4.1	PSUMR	4-1
4.1.1	PSUMR Inputs	4-1
4.1.1.1	SCF Records (File IN)	4-1
4.1.1.2	Intermediate Record - Old Master (File OM)	4-1
4.1.1.3	Parameter File (File PF)	4-1
4.1.2	PSUMR Output	4-1
4.1.2.1	Parameter File Listing (File RP)	4-1
4.1.2.2	Intermediate Record - New Master File (File NM)	4-1
4.1.2.3	Intermediate Record - Current File (File DR)	4-3
4.1.3	PSUMR Deck Setup	4-3
4.2	DEMON	4-3
4.2.1	DEMON Inputs	4-3
4.2.1.1	History File (File HF)	4-4
4.2.1.2	Parameter File (File PF)	4-4
4.2.1.3	Intermediate Record (File IN)	4-4
4.2.2	DEMON Outputs	4-4
4.2.2.1	Updated History File (File HF)	4-4
4.2.2.2	Listing File (File RP)	4-4
4.2.3	DEMON Deck Setup	4-4
4.3	RPTSUM	4-4
4.3.1	RPTSUM Inputs	4-5
4.3.1.1	Intermediate Master Record (File IF)	4-5
4.3.1.2	Parameter File (File PF)	4-5
4.3.2	RPTSUM Outputs	4-5
4.3.2.1	Parameter File Listing (File RP)	4-5
4.3.2.2	Report Listing (File GP)	4-5

Section		Page
4.3.3	RPTSUM Deck Setup	4-5
4.4	Parameter Language	4-5
4.4.1	Statement Identifiers	4-6
4.4.2	Report Statement Keywords and Parameters	4-6
4.4.2.1	NAME	4-6
4.4.2.2	SYSTEM	4-6
4.4.2.3	PERIOD	4-7
4.4.2.4	TIME	4-7
4.4.2.5	INCLUDE	4-7
4.4.2.6	EXCLUDE	4-8
4.4.2.7	Sample Parameter File	4-8
4.4.3	CONFIG Statement Keywords and Parameters	4-9
4.4.3.1	PROCESSOR	4-9
4.4.3.2	MEMORY	4-10
4.4.3.3	TAPES	4-10
4.4.3.4	DISKS	4-10
4.4.3.5	PRINTERS	4-10
4.4.3.6	Sample Parameter File	4-11
4.4.4	SUMMARY Statement Keywords and Parameters	4-11
4.4.4.1	NOSUM	4-11
4.4.4.2	SUM	4-12
4.4.5	SNUMB Statement Keywords and Parameters	4-12
4.4.5.1	SYSTEM	4-12
4.4.5.2	SPECIAL	4-12
4.4.5.3	Sample SNUMB Statement	4-12
4.5	Description of Reports	4-14
4.5.1	Report Overview Page	4-14
4.5.2	Graphs	4-14
4.5.2.1	Usage Codes	4-14
4.5.2.2	Report Graphs	4-16
4.6	Usage Concepts	4-26
4.7	Abort Codes	4-27
5.	THE GENERAL MONITOR COLLECTOR - DATA COLLECTION PROGRAM	5-1
5.1	Introduction	5-1
5.2	GMC Monitor Subroutines	5-3
5.2.1	Memory Utilization Monitor	5-4
5.2.2	Mass Storage Monitor	5-14
5.2.3	CPU Monitor	5-16
5.2.4	Tape Monitor	5-17
5.2.5	Channel Monitor	5-17
5.2.6	Communications Analysis Monitor	5-18
5.2.7	GRTS Monitor	5-19
5.2.7.1	Configuration Requirements	5-19
5.2.7.2	H6000 Configuration Requirements	5-19
5.2.7.3	Altering of Phase II-A Software	5-20
5.2.7.4	FEP Configuration Requirements	5-20

Section		Page
5.2.7.5	Interface Requirements	5-20
5.2.7.6	Abort Codes	5-21
5.2.7.7	DATANET Monitor Software Description	5-21
5.2.7.7.1	DATANET-Host Interface	5-22
5.2.7.7.2	Monitoring of CPU	5-22
5.2.7.7.3	Host/DATANET Response Time	5-23
5.2.7.7.4	Terminal Monitoring	5-24
5.2.8	Idle Monitor	5-25
5.2.9	Transaction Processing System Monitor	5-25
5.2.9.1	TPS Trace Collection	5-25
5.2.9.2	Modifying the Transaction Processing System	5-26
5.2.10	Timesharing Subsystem Monitor	5-26
5.2.10.1	Locations of TSS Trace Points	5-27.2
5.2.10.2	Formats of TSS Traces	5-27.15
5.2.10.3	Installation Procedures	5-27.16
5.2.10.3.1	Description of Monitor Software	5-27.16
5.2.10.3.2	Software Installation	5-27.16
5.2.10.3.3	Software Activation	5-27.17
5.2.10.3.3.1	Overview of TSS INIT File Changes	5-27.17
5.2.10.3.3.2	Definition of the Master Subsystem Name	5-27.17
5.2.10.3.3.3	Definition of the UST Origin	5-27.17
5.2.10.3.3.4	Installation from a Permanent File	5-27.18
5.2.10.4	Production Use of the Monitor	5-27.19
5.2.10.5	Monitor Limitations	5-27.19
5.3	Processing	5-27.20
5.3.1	Executive Routine	5-28
5.3.2	Output	5-31
5.4	GMC Data Records	5-32
5.4.1	GMC Executive	5-32
5.4.2	MUM	5-37
5.4.2.1	Trace Type 10	5-37
5.4.2.2	Trace Type 46	5-40
5.4.3	MSM	5-40
5.4.3.1	Trace Type 7	5-40
5.4.3.2	MSM Special Record	5-41
5.4.3.3	Device Name Record	5-42
5.4.3.4	FILSYS Catalog Structure Record	5-42
5.4.3.5	FMS CACHE Record	5-43
5.4.4	CPUM	5-44
5.4.4.1	Trace Type 70 - Standard	5-44
5.4.4.2	Trace Type 63 - Initial	5-46
5.4.4.3	Trace Type 63 - Activity Termination	5-46
5.4.4.4	Trace Type 63 - Termination Record	5-46
5.4.5	TM	5-46.1
5.4.5.1	Trace Type 50	5-46.1
5.4.5.2	Special Trace Type 50	5-46.1
5.4.5.3	Trace Type 51	5-47
5.4.5.4	Trace Type 52	5-47
5.4.6	CM	5-47

Section		Page
5.4.6.1	Trace Type 4	5-47
5.4.6.2	Trace Type 7	5-48
5.4.6.3	Trace Type 22.	5-48
5.4.7	CAM	5-48
5.4.7.1	Trace Type 14	5-48
5.4.7.2	Special Trace Type 14	5-49.1
5.4.7.3	Special TSS Trace Type 14	5-49.1
5.4.7.4	Special GEROUT Trace Type 14	5-49.2
5.4.8	GRTM	5-49.2
5.4.8.1	Trace Type 62.	5-49.2
5.4.9	Idlem	5-52
5.4.9.1	Trace Type 21	5-52
5.4.10	TPS	5-53
5.4.10.1	Trace Types 0, 1, 2, 4, 5, 6 and 65	5-53
5.4.10.2	Trace Types 13, 42 and 51	5-53
5.4.10.3	Trace Type 74	5-53
5.4.11	TSS Monitor	5-53
5.4.11.1	TSSM Trace Type 74	5-53
5.4.12	Special Records	5-53.22
5.4.12.1	Lost Data Record	5-53.22
5.4.12.2	Termination Record	5-54
5.4.12.3	End-of-Reel Flag	5-54
5.4.12.4	MUM Lost Data	5-54
5.4.12.5	Reconfiguration Record	5-54
5.5	GMF User Input Parameter Options	5-54
5.5.1	ON/OFF Option	5-56
5.5.2	Tape Selection Option	5-57
5.5.3	Terminal Specification Option	5-57
5.5.4	Move Option	5-58
5.5.5	CPU SNUMB Option	5-58
5.5.6	Connect Option	5-58
5.5.7	Time Option	5-58
5.5.8	Specifying High Density Tape	5-59
5.5.9	Limited Mass Store Monitor/Channel Monitor Collection	5-59.1
5.5.10	Request Next Data Card	5-60
5.5.11	Disable the Dispatcher Queuing Option	5-60
5.5.12	Specifying Monitoring Requirements for the GRTM	5-60
5.5.13	General Rules of the GMC Data Parameter Card	5-61
5.6	JCL for Creation of an Object File	5-61
5.6.1	Introduction to JCL	5-61
5.6.2	Creation of an Object File	5-62
5.7	JCL for Executing the GMC	5-68
6.	MEMORY UTILIZATION DATA REDUCTION PROGRAM	6-1
6.1	Inputs	6-1

Section		Page
6.1.1	Report Options	6-1
6.1.2	Default Options	6-1
6.1.3	Histogram Options	6-1
6.1.4	Plot Options	6-6
6.1.5	Default Option Alteration	6-6
6.1.6	Histogram Alterations (Action Code HISTG)	6-12
6.1.7	Plot Alterations (Action Code PLOT)	6-12
6.1.8	Turn a Report On (Action Code ON)	6-12.3
6.1.9	Turn a Report Off (Action Code OFF)	6-12.3
6.1.10	Set a Time Span of Measurement (Action Code TIME)	6-12.3
6.1.11	Turn All Reports Off Except Those Specified (Action Code ALLOFF)	6-15
6.1.12	Turn All Reports On Except Those Specified (Action Code ALLON)	6-15
6.1.13	Continue Data Reduction After an Input Option Error (Action Code ERROR)	6-15
6.1.14	Debug for a Given Program Number (Action Code DEBUG)	6-15
6.1.15	Stop After a Specified Number of Tape Records Processed (Action Code NREC)	6-15
6.1.16	Suppress USERID (Action Code NOUSER)	6-18
6.1.17	Turn Idle Reports Off (Action Code IDLE)	6-18
6.1.18	Change Excessive Resource Limits Used in Excessive Resource Report (Action Codes WASTED, CORE, IO, CPU, and RATIO)	6-18
6.1.19	Eliminate SNUMBs From Abort Report (Action Code ABORT)	6-18
6.1.20	Change the Plot Interval (Action Code PLTINT)	6-18
6.1.21	Change the Program Number for the First Slave Job (Action Code FSTSLV)	6-19
6.1.22	Request That Certain Jobs be Considered System Jobs (Action Code MASTER)	6-19
6.1.23	PALC Report Print Control (Action Code PALC)	6-19
6.1.24	Request the Special Job Memory Reports (Action Code SPECL)	6-19
6.1.25	Process Data on a WW6.4 System (Action Code RN)	6-19.1
6.1.26	Produce a Memory Map Only Under Certain Memory Demand Conditions (Action Code MAPART)	6-19.1
6.1.27	Include Zero Urgency Jobs in all Memory Calculations (Action Code ZERO)	6-19.1
6.2	Processing	6-21
6.2.1	General	6-21
6.2.2	JCL	6-21
6.3	Outputs	6-23
6.3.1	MUM Title Page	6-23

Section		Page
6.3.2	System Program Usage	6-28
6.3.3	MUM Reports	6-31
6.3.3.1	Report 1 - Memory Demand Sizes of New Activities in 1K Word Blocks	6-31
6.3.3.2	Report 2 - The Memory Demand Size of All Demand Types	6-34
6.3.3.3	Report 3 - The Total Memory Demand Outstanding	6-34
6.3.3.4	Report 4 - The Demand That was Outstanding When a Processor Went Idle	6-34
6.3.3.5	Report 5 - The Total Amount of Available Memory	6-34
6.3.3.6	Report 6 - The Memory Available When a Processor Went Idle	6-35
6.3.3.7	Report 7 - The Time Corrected Total Demand Outstanding	6-35
6.3.3.8	Report 8 - The Time Corrected Memory Available	6-35
6.3.3.9	Report 9 - The Number of Activities Waiting for Memory in Allocator Queue	6-35
6.3.3.10	Report 10 - The Number of User Activit Waiting Memory in Allocator Queue	6-35
6.3.3.11	Report 11 - The Time Corrected Number of Activities Waiting Memory	6-36
6.3.3.12	Report 12 - The Time Corrected Number of User Activities Waiting Memory	6-36
6.3.3.13	Report 13 - The Number of Activities Waiting Memory When a Processor Went Idle	6-36
6.3.3.14	Report 14 - The Number of Activities Residing in Memory	6-36
6.3.3.15	Report 15 - The Number of User Activities in Memory	6-36
6.3.3.16	Report 16 - The Time Corrected Number of Activities in Memory	6-36
6.3.3.17	Report 17 - The Time Corrected Number of User Activities in Memory	6-37
6.3.3.18	Report 18 - The Number of Activities in Memory When a Processor Went Idle	6-37
6.3.3.19	Report 19 - The Ratio of User Activity Duration Versus Its Memory Use Time	6-37
6.3.3.20	Report 20 - The Elapsed Duration of User Activity in 10ths of a Second	6-37
6.3.3.21	Report 21 - The Total Elapsed Time a User Activity was in Memory	6-37
6.3.3.22	Report 22 - The GEMORE Service or Denial Time - 1/10 Second, Elapsed	6-38
6.3.3.23	Report 23 - The Request Size of GEMOREs	6-38

Section		Page
6.3.3.24	Report 24 - Delay Time in the System Scheduler	6-38
6.3.3.25	Report 25 - Delay Time Until Core Allocator	6-38
6.3.3.26	Report 26 through 31 - The Elapsed Time of a Busy State of the Processors	6-38
6.3.3.27	Report 32 - The Elapsed Time of a Busy State of Processors	6-38
6.3.3.28	Report 33 - Elapsed Time Between Allocator Calls in 1/100 of a Second	6-39
6.3.3.29	Report 34 - The I/O Time Charged per User Activity in Seconds	6-39
6.3.3.30	Report 35 - The CP Time Charged per User Activity in Seconds	6-39
6.3.3.31	Report 36 - The Number of Times a User Activity was Swapped	6-39
6.3.3.32	Report 37 - The Total Elapsed Time a User Activity was Swapped	6-39
6.3.3.33	Report 38 - The Number of Times a System Activity was Swapped	6-39
6.3.3.34	Report 39 - The Total Elapsed Time a System Activity was Swapped	6-39
6.3.3.35	Report 40 - Number of Extra Activities That Might Fit in Memory Using Compaction	6-39
6.3.3.36	Report 41 - Number of Extra Activities That Might Fit Memory Without Compaction	6-40
6.3.3.37	Report 42 - The Percent of Size-Time Product Used by a User Activity	6-40
6.3.3.38	Report 43 through 49 - The Length of Idle State in the Processors	6-40
6.3.3.39	Report 50 - The Original Allocation Time for User Memory in 1/10 Second	6-41
6.3.3.40	Report 51 - The Time Corrected Percent of Assigned Memory Used	6-41
6.3.4	Activity Resource Usage Report	6-41
6.3.5	SNUMB-IDENT Report	6-43
6.3.6	Memory Map Report	6-43
6.3.7	Demand List Report	6-48
6.3.8	Activity Abort Report	6-48
6.3.9	Jobs Out of Core Report	6-50
6.3.10	Excessive Resource Use Report	6-50
6.3.11	Allocation Status Report	6-53
6.3.12	Plots Reports	6-53.1
6.3.12.1	Plot 1 - Available Memory vs. Outstanding Demand in Core Allocator Queue vs. Outstanding Demand in Core Allocator + Peripheral Allocator Queue	6-56

Section		Page
6.3.12.2	Plot 2 - Memory Shortfall in Core Allocator vs. Memory Shortfall in Core Allocator + Peripheral Allocator	6-56
6.3.12.3	Plot 3 - Number of Activities in Core Queue vs. Number of Activities in Peripheral Allocator Queue	6-56
6.3.12.4	Plot 4 - Average Size of TSS, FTS and NCP	6-56
6.3.13	Memory Statistics Report	6-58
6.3.14	Special Job Memory Reports	6-58
6.3.15	Distribution of Urgency Over Time Report	6-58
6.3.16	Zero Urgency Job Report	6-58.3
6.4	Error Messages	6-58.4
6.5	Multireel Processing	6-62
6.6	TAPE Error Aborts	6-63
7.	MASS STORE DATA REDUCTION PROGRAM (MSDRP)	7-1
7.1	Introduction	7-1
7.2	Data Collection Methodology	7-3
7.3	Analytical Methodology	7-3
7.4	Data Reduction Methodology	7-4
7.5	MSMDRP Output	7-5
7.5.1	System Configuration and Channel Usage Report (File 42)	7-5
7.5.2	System Summary Report (File 42)	7-9
7.5.3	System Traces Captured by Monitor Report (File 42)	7-10
7.5.4	Channel Status Changes Report (File 29)	7-10
7.5.5	Physical Device, Device ID Correlation Table (File 42)	7-10
7.5.6	Device Space Utilization Report (File 42)	7-10
7.5.7	Device Seek Movement Report (File 42)	7-15
7.5.8	Head Movement Efficiency Report (File 42)	7-17
7.5.9	System File Use Summary Report (File 21)	7-19
7.5.10	Individual Module Activity Report (File 21)	7-21
7.5.11	SSA Module Usage Report by Job (File 21)	7-23
7.5.12	File Code Summary Report (File 23) (NAME=FILECODE)	7-23
7.5.13	Cat/File String Report (File 23)	7-26
7.5.14	Connect Summary Report By Userid/SNUMB (File 23)	7-29
7.5.15	Activity Summary Report (File 24)	7-29
7.5.16	Device Area File Code Reference Report (File 22)	7-32
7.5.17	Device File Use Summary Report (File 21)	7-32
7.5.18	Chronological Device Utilization Report (File 26)	7-32
7.5.19	FMS Cache Report (File 21)	7-36

Section		Page
7.5.20	Proportionate Device Utilization Report (File 42)	7-36.4
7.5.21	Elapsed Time Between Seeks Report (File 42)	7-39
7.5.22	Data Transfer Size Report (File 42)	7-39
7.5.23	Data Transfer Sizes For TSS Swap Files (File 42)	7-42
7.5.24	Connects Per Minute Report (File 20)	7-42
7.5.25	Special FTS File Access Time Report (File 42)	7-42
7.5.26	TSS Swap File Usage Over Time Report (File 42)	7-42.1
7.5.27	Device Seek Movement Summary Report (File 29)	7-42.1
7.5.28	Special Processing Messages	7-42.2
7.6	Default Option Alteration	7-47
7.6.1	Monitor a Specific Device Area (Action Code AREA)	7-48
7.6.2	System Debug (Action Code DEBUG)	7-48
7.6.3	Continue Data Reduction After an Input Option Error (Action Code ERROR)	7-48
7.6.4	Specify System File Names (Action Code FILDEF)	7-48.1
7.6.5	End Card (Action Code END)	7-50
7.6.6	Produce the SSA Module Usage Report by Job (Action Code MODULE)	7-50
7.6.7	Record Limitation by Connects (Action Code NCONN)	7-50
7.6.8	Record Limitation by Records (Action Code NREC)	7-50
7.6.9	Turn a Report Off (Action Code OFF)	7-50
7.6.10	Turn a Report On (Action Code ON)	7-51
7.6.11	Produce Connect Summary Report by Userid/SNUMB (Action Code PROJ)	7-51
7.6.12	Reduce WW6.4 Data or Process MSMDRP on a WW6.4 System (Action Code RN)	7-51
7.6.13	Set a Timespan of Measurement (Action Code TIME)	7-51
7.6.14	Change the Time Quantum Value for the Chronological Device Utilization Report (Action Code TIMEQ)	7-53
7.6.15	Suppress the USERIDs (Action Code USERID)	7-53
7.6.16	Change the Time Quantum Value for the Connect Per 10 Minute Report (Action Code RATECH)	7-55
7.6.17	Turn on the Cat/File String Report (Action Code CAT)	7-55
7.6.18	Request the Connect Per 10 Minute Report for Specific User Job (Action Code RATE)	7-55
7.6.19	Limit the Processing and Output (Action Code LIMITS)	7-55
7.6.20	Zero Activity Processing (Action Code ZERO)	7-55
7.7	JCL	7-56
7.8	Multireel Processing	7-56

Section		Page
7.9	Tape Error Aborts	7-58
8.	CHANNEL MONITOR DATA REDUCTION PROGRAM (CMDRP)	8-1
8.1	Introduction	8-1
8.2	Data Collection Methodology	8-1
8.3	Analytical Methodology	8-3
8.4	Data Reduction Methodology	8-7
8.5	CMDRP Output	8-8
8.5.1	System Configuration and Channel Usage Report (File 57)	8-8
8.5.2	System Summary Report (File 57)	8-10
8.5.3	System Traces Captured by Monitor Report (File 57)	8-13
8.5.4	Channel Status Changes Report (File 57)	8-13
8.5.5	Physical Device, Device ID Correlation Table (File 57)	8-13
8.5.6	Channel Statistics Report (File 57)	8-13
8.5.6.1	Channel Busy and Device Busy Report	8-13
8.5.6.2	Channel Busy and Device Free Report	8-13
8.5.6.3	Channel Free and Device Busy Report	8-13
8.5.6.4	Channel Free and Device Free Report	8-20
8.5.6.5	GEPR Connect Report	8-20
8.5.6.6	Lost Interrupt Report	8-20
8.5.6.7	Device ID STIOS Not Connected Report	8-20
8.5.6.8	Entries Still in Queue Report	8-20
8.5.7	Idle Monitor Report (File 57)	8-20
8.5.8	Proportionate Device Utilization Report (File 57)	8-27
8.5.9	Queue Length and Queue Time Histograms (File 57)	8-29
8.5.10	Service Time Histograms (File 57)	8-34
8.5.11	Activity Statistic Report (Files 23 and 24)	8-37
8.5.12	Job Conflict Report (Files 31, 32, 33, 34)	8-37
8.5.13	Special Job Processing Report by Device (File 32)	8-41
8.5.14	Special Job Processing Report Per 10 Minutes (File 32)	8-41
8.5.15	Special Processing Messages	8-44
8.6	Default Option Alteration	8-45
8.6.1	Job Device Conflict Report (Action Code QDEV)	8-46
8.6.2	Program Debug Options	8-46
8.6.2.1	Program Number Debug (Action Code DPRG)	8-46
8.6.2.2	Device Debug (Action Code DDEV)	8-46
8.6.2.3	Queue Location Debug (Action Code DQUE)	8-46
8.6.2.4	Random Access File Debug (Action Code DEBUG)	8-46

Section		Page
8.6.2.5	Channel Debug (Action Code DCHN)	8-46
8.6.3	Removal of Queue Entries (Action Code DELTA) . . .	8-46
8.6.4	Set a Timespan of Measurement (Action Code TIME)	8-47
8.6.5	Turn a Report On/Off (Action Code ON/OFF)	8-47
8.6.6	Continue Data Reduction After an Input Option Error (Action Code ERROR)	8-47
8.6.7	W6.4/2H Data Reduction (Action Code RN)	8-47
8.6.8	Record Limitation by Connects (Action Code NCONN)	8-47
8.6.9	Record Limitation by Records (Action Code NREC)	8-47
8.6.10	Special Job Processing Report (Action Code JOB)	8-49
8.6.11	Change the Time Quantum Value for the Special Job Processing Report Per 10 Minutes (Action Code RATE)	8-49
8.6.12	END Card (Action Code END)	8-49
8.6.13	Limit the Processing and Output (Action Code LIMITS)	8-49
8.6.14	Zero Activity Processing (Action Code ZERO) . . .	8-49
8.7	JCL	8-50
8.8	Multireel Processing	8-50
8.9	Tape Error Aborts	8-52
9.	COMMUNICATIONS ANALYSIS MONITOR DATA REDUCTION PROGRAM (CAMDRP)	9-1
9.1	Introduction	9-1
9.2	Data Collection Methodology	9-1
9.3	Analytical Methodology	9-1
9.4	Data Reduction Methods	9-2
9.5	CAMDRP Output	9-2
9.5.1	Header Page	9-4
9.5.2	Trace Dumps	9-4
9.5.2.1	355 Mailbox Report-Trace Dump	9-4
9.5.2.2	Terminal Mailbox Dump	9-8
9.5.3	Statistical Summary Reports	9-11
9.5.3.1	DAC Devices Summary Report	9-11
9.5.3.2	DAC Subsystem Summary Report	9-13
9.5.3.3	Remote Batch Device Summary Report	9-13
9.5.3.4	Terminal ID Summary Report	9-16
9.5.4	Delta Time Period Summary Report	9-18
9.5.5	Histogram Reports	9-18
9.5.6	Response Time Limit Report	9-20
9.5.7	User Think Time Limit Report	9-24
9.5.8	Terminal Session and High Terminal Usage Reports	9-24
9.5.9	Opcode Count Report	9-24

Section		Page
9.5.10	Response Time Report	9-29
9.5.11	Error Messages	9-29.1
9.5.12	H6000-DN355 Reject Report	9-29.1
9.5.13	Abort/Disconnect Report	9-29.1
9.5.14	TS1 Initial Parameter Report	9-29.1
9.5.15	Mailbox Busy Report	9-29.2
9.6	Default Option Alteration	9-31
9.6.1	Timeframe Reduction Report (Action Code TIME)	9-31
9.6.2	Delta Timeframe Report (Action Code DELTA)	9-31
9.6.3	Histogram Report (Action Code HISTG)	9-33
9.6.4	Trace Dump Report (Action Code LIST or ALL)	9-33
9.6.5	Record Count Limitation (Action Code NREC)	9-33
9.6.6	Response Time Limit (Action Code RESP)	9-33
9.6.7	Think Time Limit (Action Code THINK)	9-33
9.6.8	Terminal Mailbox Dump (Action Code MAIL)	9-36
9.6.9	Terminal Busy Limit (Action Code BUSY)	9-36
9.6.10	W6.4/2H Data Reduction (Action Code RN)	9-36
9.6.11	Response Time Report Timeframe (Action Code RATE)	9-36
9.6.12	Response Time Report SNUMB (Action Code SNUMB)	9-36
9.6.13	Disconnect Report Line IDs (Action Code DISC)	9-36
9.6.14	Terminate Input Options (Action Code END)	9-36
9.6.15	Default Options	9-36.1
9.7	JCL	9-38
9.8	Multireel Processing	9-38
9.9	Tape Error Aborts	9-40
10.	DATANET-355 DATA REDUCTION PROGRAM (DDRP)	10-1
10.1	Introduction	10-1
10.2	Data Collection Methodology	10-1
10.3	Analytical Methodology	10-1
10.4	Data Reduction Methods	10-3
10.5	DDRP Output	10-4
10.5.1	DATANET-355 Tabular Reports	10-9
10.5.1.1	Average Response Time for all Users by DATANET (Report RESPD)	10-9
10.5.1.2	HSLA/Subchannel Transmission Report by DATANET (Report BUFF)	10-12
10.5.1.3	Average Response Time for Specifically Designated Line IDs (Report RESPL)	10-12
10.5.1.4	HSLA Subchannels Being Monitored (Report HSLA)	10-12
10.5.1.5	Card Images from I* of the GRTS-II Data Collection (Report CARD)	10-16
10.5.1.6	List of Active Line IDs	10-16
10.5.1.7	HSLA Threshold Report (Report THRESH)	10-16

Section		Page
10.5.1.8	Response Interval Unmatched Pairs Verification	10-16
10.5.1.9	Response Interval Matched Pairs Verification (Report MATCH)	10-20
10.5.1.10	Annotated List of Datanet Traces (Report TRACE)	10-20
10.5.2	Plots	10-20
10.5.2.1	Number of Transactions Sent to Host (Plot ID PLOT1)	10-22
10.5.2.2	Number of Transactions Received From Host (Plot ID PLOT2)	10-22
10.5.2.3	Number of 36-Bit Words Sent to Host (Plot ID PLOT3)	10-24
10.5.2.4	Number of 36-Bit Words Received From Host (Plot ID PLOT4)	10-24
10.5.2.5	Percent Idle Time Over a Plot Interval (Plot ID PLOT5)	10-24
10.5.2.6	Percent Buffer Requests Denied (Plot ID PLOT6)	10-24
10.5.2.7	Number of 18-Bit Words Available for Buffers (Plot ID PLOT7)	10-24
10.5.2.8	Number of Users Logged on the System (Plot ID PLOT8)	10-28
10.5.2.9	Number of Host RSVPs Received (Plot ID PLOT9)	10-28
10.5.3	Histograms	10-28
10.6	Default Option Alteration	10-31
10.6.1	Histogram Desired (Action Code HISTG)	10-33
10.6.2	Plot Alteration (Action Code PLOT)	10-33
10.6.3	Turn a Report On (Action Code ON)	10-33
10.6.4	Turn a Report Off (Action Code OFF)	10-35
10.6.5	Set a Timespan of Measurement (Action Code TIME)	10-35
10.6.6	Process Reduction on a WW6.4/2H System (Action Code RN)	10-35
10.6.7	Turn All Reports Off Except Those Specified (Action Code ALLOFF)	10-35
10.6.8	Turn All Reports On Except Those Specified (Action Code ALLON)	10-35
10.6.9	Continue Data Reduction After an Input Option Error (Action Code ERROR)	10-39
10.6.10	Stop Processing (Action Code NREC)	10-39
10.6.11	Plots Only (Action Code REPORT)	10-39
10.6.12	Print Formatted DATANET-355 Record Types (Action Code TRACE)	10-39
10.6.13	DATANETs Not to Plot (Action Code NOPLOT)	10-40
10.6.14	Set a Plot Interval (Action Code INTERV)	10-40

Section		Page
10.6.15	Response Time Records (Type 2) Matched (Action Code MATCH)	10-40
10.6.16	Accept Line IDs for Special Analysis (Action Code SPECL)	10-42
10.6.17	Modify Threshold Values (Action Code THRESH)	10-42
10.6.18	Suppress Nonplot Reports for Selected Nets (Action Code NORPT)	10-42
10.6.19	Debug (Action Code DEBUG)	10-45
10.6.20	Alter Histogram Default Parameters (Action Code CHANGE)	10-45
10.6.21	Unmatched/Duplicate Error Report (Action Code DUPLIC)	10-45
10.6.22	Special Reject Report (Action Code REJECT)	10-45
10.7	JCL	10-45
10.8	Multireel Processing	10-50
10.9	Tape Error Aborts	10-51
11.	CENTRAL PROCESSING UNIT AND TAPE REDUCTION PROGRAM (CPUTRP)	11-1
11.1	Introduction	11-1
11.2	Data Collection Methodology	11-1
11.3	Analytical Methodology	11-1
11.4	Data Reduction Methodology	11-1
11.5	CPUTRP Output	11-2
11.5.1	Execution and Error Reports (Files 6, 7 and 8)	11-2
11.5.2	Central Processing Unit Monitor Reports	11-6
11.5.2.1	CPU Queue Length Distribution (File 14)	11-6.1
11.5.2.2	CPU Burst Distribution for User Activities/ All Activities (File 14)	11-9
11.5.2.3	Percent of Memory Time in Queue	11-9
11.5.2.4	CPU Time Reports (File 10)	11-9
11.5.2.5	CPU Plot of Percent Idle (File 31)	11-17
11.5.2.6	WIN Report	11-17
11.5.2.7	CPU Access by SNUMB Report	11-17
11.5.3	Tape Monitor Reduction Reports	11-17
11.5.3.1	Number of Tape Drives in Use (Time Corrected) (File 14)	11-17.1
11.5.3.2	Tape Activity Report (File 13)	11-17.1
11.5.3.3	Tape Status Report (File 12)	11-23

Section		Page
11.6	Default Option Alteration	11-23
11.6.1	General Format	11-23
11.6.1.1	Valid Commands	11-23
11.6.1.2	Parameter List	11-26
11.6.1.3	Command Syntax	11-29
11.6.1.4	General Control Command Syntax	11-29
11.6.1.5	CPU Reduction Command Syntax	11-32
11.6.1.6	Tape Reduction Command Syntax	11-33
11.6.1.7	Card Input Errors	11-33
11.6.1.8	Examples of Command Use	11-33
11.7	JCL	11-34
11.8	Multireel Processing	11-35
11.9	Tape Error Aborts	11-37
12.	TRANSACTION PROCESSING SYSTEM DATA REDUCTION PROGRAM (TPSDRP)	12-1
12.1	Introduction	12-1
12.2	Transaction Processing System Trace Number	12-1
12.3	Transaction Processing System Trace Collection	12-1
12.3.1	Sample Decks	12-1
12.4	TPE Tuning Guide Reports	12-1
12.4.1	Input	12-1
12.4.2	Output	12-6
12.4.2.1	TPE Summary Report	12-6
12.4.2.2	NONEW Option Report	12-8
12.4.2.3	Profile ID Report	12-8
12.4.2.4	Buffer Space Report	12-8
12.4.2.5	Table Reaching Threshold Size	12-8
12.4.2.6	Master Terminal Report	12-8
12.4.2.7	Impasse Report	12-8
12.4.2.8	Output Transmission Report	12-16
12.4.2.9	TPAP ABORT Report	12-16
12.4.2.10	UNDOT Entry Report	12-16
12.4.2.11	Reissuing Intercom Message Report	12-16
12.4.2.12	Receiving Intercom Message Report	12-16
12.4.2.13	TPAP Modify Report	12-16
12.4.2.14	TPAP Status Report	12-16
12.4.2.15	TPAP Run Times Report	12-16
12.5	TPE Tuning Guide Formatted Dump	12-26
12.5.1	Input	12-26
12.5.2	Formatted Output	12-27
12.5.3	Summary Report	12-27
13.	EXTENDED USE OF GMC	13-1
13.1	Introduction	13-1
13.2	Monitor Structure	13-1
13.2.1	Initialization Section	13-1

Section		Page
13.2.2	Monitor Subroutine	13-1
13.3	User Generated Trace	13-4
13.4	Data Reduction	13-4
13.4.1	Subroutine INITZ	13-4
13.4.2	Subroutine NXTRECRD	13-6
13.4.3	Subroutine RLOVER	13-6
13.4.4	Subroutine TIMEDAY	13-7
13.4.5	Subroutine MAIN	13-7
13.4.6	Other Routines	13-7
13.5	Utility Tape Dump	13-7
14.	CONDUCTING A SITE COMPUTER PERFORMANCE EVALUATION USING THE GMC	14-1
14.1	Introduction	14-1
14.2	General Definitions	14-1
14.2.1	Computer Performance Evaluation (CPE)	14-1
14.2.2	Computer System Performance Variables	14-1
14.2.2.1	System Design	14-1
14.2.2.2	Programming	14-1
14.2.2.3	Hardware Configuration	14-1
14.2.2.4	System Software	14-2
14.2.2.5	Operations	14-2
14.2.2.6	Communications Hardware and Software	14-2
14.2.2.7	Computer System Performance Tuning	14-2
14.2.2.8	Turnaround Time	14-2
14.2.2.9	CPE User Objectives	14-2
14.3	Solutions to Performance Problems	14-2
14.3.1	Scheduling Solutions	14-2
14.3.2	Parameter Solutions	14-2
14.3.3	Programming Solutions	14-3
14.3.4	Sizing Solutions	14-3
14.4	Structure of the Analysis Process	14-3
14.4.1	Starting the Process	14-3
14.4.1.1	Direct Request	14-3
14.4.1.2	Internal Review	14-3
14.4.1.3	User Requests	14-3
14.4.2	Problem Definition Phase	14-3
14.4.2.1	Define and Verify the Problem	14-6
14.4.2.2	Gain Understanding of Facility Environment	14-6
14.4.2.2.1	Hardware Configuration	14-6
14.4.2.2.2	Software Configuration and Development Practices	14-6
14.4.2.2.3	Existing CPE Practices	14-6
14.4.2.2.4	Site Workload Characteristics	14-7
14.4.2.2.5	System Users	14-7
14.4.2.2.6	Operations Practices	14-7
14.4.2.2.7	Batch Job Scheduling	14-7

Section		Page
14.4.2.2.8	Site and Computer Facility Organization . . .	14-7
14.4.2.3	Understand Installation Service Objectives and Priorities	14-7
14.4.2.3.1	Installation Service Objectives	14-7
14.4.2.3.1.1	Production Objective	14-8
14.4.2.3.1.2	Availability Objective	14-8
14.4.2.3.1.3	Mixed Objectives	14-8
14.4.2.3.2	Installation Priorities	14-9
14.4.2.3.2.1	Service Priorities	14-9
14.4.2.3.2.2	Evaluation/Tuning Solution Priorities . . .	14-9
14.4.2.4	Specify Current Evaluation	14-9
14.4.2.4.1	Objective	14-9
14.4.2.4.2	Objective Decision	14-9
14.4.2.4.2.1	Attainable Objectives	14-9
14.4.2.4.2.2	Realistic Objectives	14-9
14.4.2.4.2.3	Cost-Effective Objectives	14-10
14.4.2.4.3	Determine if Worth Continuing	14-10
14.4.2.4.4	Begin Problem Analysis	14-10
14.4.3	Problem Analysis Phase	14-10
14.4.3.1	Run Appropriate Analysis Tool	14-10
14.4.3.2	Evaluate System Output	14-10
14.4.3.3	Follow Tuning Procedures	14-10
14.4.3.4	Evaluate Need to Continue the Tuning	14-10
14.5	Composition of a Performance Evaluation Team	14-11
14.6	System Evaluation	14-11
14.6.1	Introduction	14-11
14.6.2	Selecting a Representative Value From GMC Histogram Reports	14-11
14.6.2.1	Symmetric Distribution Closely Clustered Around a Single Point	14-11
14.6.2.2	Skewed Distribution	14-13
14.6.2.3	Distribution With Outliers	14-13
14.6.3	Memory Evaluation	14-13
14.6.3.1	Obtaining the Data	14-15
14.6.3.2	Evaluating the Data	14-15
14.6.4	CPU Evaluation	14-18.1
14.6.4.1	Data Recording	14-18.1
14.6.4.2	Evaluating the Data	14-19
14.6.5	I/O Evaluation	14-22
14.6.5.1	Data Recording	14-22
14.6.5.2	Evaluating the Data	14-22
14.6.5.3	Mass Storage Operation	14-25
14.6.6	Communication Evaluation	14-31
14.6.6.1	Data Recording	14-31
14.6.6.2	Evaluating the Data	14-31

Section		Page
15.	TIMESHARING EMULATION AND RESPONSE SYSTEM (TEARS)	15-1
15.1	Introduction	15-1
15.2	Data Collection Methodology	15-1
15.3	Analytical Methodology	15-2
15.4	Data Reduction Methodology	15-3
15.4.1	Response Pass Reduction Methodology	15-4
15.4.2	Emulation Reduction Methodology	15-5
15.5	Output	15-7
15.5.1	TEARS Response Pass Reports	15-38
15.5.1.1	TSS Reduction Event Log	15-38
15.5.1.2	Response Time Logplots	15-41
15.5.1.3	Subdispatch Time Logplots	15-41
15.5.1.4	CPU Times Allocated to TS1-Exec and TS1-Subdispatch Report	15-42
15.5.1.5	TSS Subtraces Encountered, and Derails Executed by Users	15-43
15.5.2	TEARS Emulation Reports	15-43
15.5.2.1	Exception Message Report	15-44
15.5.2.2	Users' Map	15-44
15.5.2.3	TSS Delay - Users Delay Report	15-45
15.5.2.4	TSS Core Map	15-47
15.5.2.5	Error Message Report	15-48
15.5.2.6	Intertrace Gap - Duration	15-49
15.5.2.7	Memory Activity Report	15-49
15.5.2.8	Eligibles for Subdispatch - Duration	15-50
15.5.2.9	In Subdispatch - Duration	15-50
15.5.2.10	Eligibles for and In Subdispatch - Duration	15-50
15.5.2.11	User Swaps - Swap Rate	15-50
15.5.2.12	User Swaps - Swap Amount	15-50
15.5.2.13	User Swaps - Duration	15-51
15.5.2.14	User I/Os - Duration	15-51
15.5.2.15	In System Master Catalog Wait - Duration	15-51
15.5.2.16	Intertrace Time (Adjusted for TSS GEWAKES)	15-51
15.5.3	Formatted Dump Reports	15-51
15.5.3.1	Formatted Dump	15-51
15.5.3.2	Intertrace Time	15-52
15.6	Reduction Run-Time Options	15-52
15.6.1	Parameter List	15-53
15.6.2	Command Syntax	15-56
15.6.3	Examples	15-61
15.6.4	Card Input Errors	15-61
15.7	JCL for Reduction Execution	15-62
15.8	Multireel Processing	15-62
15.9	Tape Error Aborts	15-64

DISTRIBUTION 16-1
DD FORM 1473 17-1

ILLUSTRATIONS

Figure		Page
2-1	RMC Flowchart	2-3
2-2	Programs in the RMC Subsystem	2-4
2-3	Programs in the GMC Subsystem	2-5
2-4	Subroutines and Traces in GMC Data Collector Programs	2-6
2-5	B29IDPXO/SOURCE and B29IDPXO/OBJECT Catalog Structure	2-8
2-6	B29IDPXO/JCL Catalog Structure	2-9
3-1	RMC JCL	3-2
4-1	RMDR Overview	4-2
4-2	Default Systems SNUMBs	4-13
4-3	Daily Resource Monitoring Overview	4-15
4-4	Processor Utilization	4-17
4-5	Memory Utilization	4-18
4-6	I/O Channel Time	4-19
4-7	Disk Space Summary	4-20
4-8	Peripheral Summary	4-21
4-9	Scheduled Jobs	4-22
4-10	Execution/Swapped Summary	4-23
4-11	Waiting Memory/Peripherals	4-24
4-12	Remote Status	4-25
4-13	TSS Usage	4-28
5-1	GMC Concept	5-2
5-2	Data Card Examples	5-55
5-3	GMC Catalog File Structure	5-63
5-4	GMC JCL Structure	5-65
5-5	Creation of R* File for Memory, Idle, and CPU Monitor	5-66
5-6	Creation of R* File for Mass Store Monitor, Idle and Channel Monitor	5-67
5-7	JCL for Executing the GMC	5-69
5-8	GRTM JCL	5-70
6-1	Standard Histogram	6-12.1
6-2	HISTG Action Code Format	6-12.2

Figure		Page
6-3	Standard Plot	6-13
6-4	PLOT Action Code Format	6-14
6-5	TIME Action Code Format	6-16
6-6	ALLOFF/ALLON Action Code Format	6-17
6-7	System Bottleneck Chart	6-20
6-8	JCL to RUN MUDRP	6-22
6-9	MUM Title Page Report - Idle Monitor Active	6-25
6-10	MUM Title Page Report - Idle Monitor Off	6-27
6-11	System Program Load	6-30
6-12	Standard Histogram Report	6-32
6-13	Out-of-Range Histogram	6-33
6-14	THIS FIGURE DELETED	
6-15	Activity Resource Usage Report	6-42
6-16	SNUMB IDENT Report	6-44
6-17	Memory Map Report	6-45
6-18	Demand List Report	6-49
6-19	Abort Report	6-51
6-20	Jobs Out of Core Report	6-52
6-21	Excessive Resource Usage Report	6-54
6-22	Allocation Status Report	6-55
6-23	Standard Plot	6-57
6-24	Memory Statistics Report	6-59
6-25	Special Job Memory Demand Report	6-60
6-26	Special Job Memory Size Report	6-61
6-27	Distribution of Urgency Over Time Report	6-61.1
6-28	Zero Urgency Job Report	6-61.2
7-1	System Configuration and Channel Usage Report	7-6
7-2	MSM System Summary Report	7-8
7-3	System Traces Captured by Monitor Report	7-11
7-4	Channel Status Changes Report	7-12
7-5	Physical Device, Device ID Correlation Table	7-13
7-6	Device Space Utilization Report	7-14
7-7	Device Seek Movement Report	7-16
7-8	Head Movement Efficiency Report	7-18
7-9	System File Use Summary Report	7-20
7-10	Individual Module Activity Report	7-22
7-11	SSA Module Usage Report by Job	7-24
7-12	File Code Summary Report	7-25
7-13	Cat/File String Report	7-27
7-14	Connect Summary Report by USERID/SNUMB	7-30
7-15	Activity Summary Report	7-31
7-16	Device Area File Code Reference Report	7-33
7-17	Device File Use Summary Report	7-34
7-18	Chronological Device Utilization Report	7-35
7-19	FMS Cache Report	7-37
7-20	Proportionate Device Utilization Report	7-38

Figure		Page
7-21	Elapsed Time Between Seeks Report	7-40
7-22	Data Transfer Size Report	7-41
7-23	Data Transfer Sizes For TSS Swap Files Report	7-43
7-24	Connect Per 5 Minute Report	7-44
7-24.1	Special FTS File Access Time Report	7-44.1
7-24.2	TSS Swap File Usage Over Time Report	7-44.2
7-24.3	Device Seek Movement Summary Report	7-45
7-25	Specific Device Area Report Card Input	7-49
7-26	Limited File Code Summary Input Card Format	7-52
7-27	Input Option TIME Card Format	7-54
7-28	MSMDRP JCL	7-57
8-1	IOM Configuration	8-2
8-2	System Configuration and Channel Usage Report	8-9
8-3	System Summary Report	8-11
8-4	System Traces Captured by Monitor Report	8-14
8-5	Channel Status Changes Report	8-15
8-6	Physical Device, Device ID Correlation Table	8-16
8-7	Channel Busy and Device Busy Report	8-17
8-8	Channel Busy and Device Free Report	8-18
8-9	Channel Free and Device Busy Report	8-19
8-10	Channel Free and Device Free Report	8-21
8-11	GEPR Connect Report	8-22
8-12	Lost Interrupt Report	8-23
8-13	Device ID STIOS Not Connected	8-24
8-14	Entries Still in Queue	8-25
8-15	Idle Report	8-26
8-16	Proportionate Device Utilization Report	8-28
8-17	Device Queue Length and Time Histograms	8-30
8-18	Channel Queue Length and Queue Time Report	8-32
8-19	I/O Service Time Report	8-35
8-20	Activity Statistic Report	8-38
8-21	Job Conflict Report	8-40
8-22	Special Job Processing Report by Device	8-42
8-23	Special Job Processing Report Per 5 Minute Report	8-43
8-24	Input Option TIME Format	8-48
8-25	CMDRP JCL	8-51
9-1	Communications Analysis Concept	9-3
9-2	CAMDRP Header Page	9-5
9-3	Mailbox Report	9-6
9-4	DAC Device Summary Report	9-12
9-5	DAC Subsystem Summary Report	9-14
9-6	Remote Batch Device Summary Report	9-15
9-7	Terminal ID Summary Report	9-17
9-8	Delta Time Period Summary Report	9-19
9-9	Machine Response Time Report	9-21
9-10	User Think Time Report	9-22
9-11	Session Length Report	9-23

Figure		Page
9-12	Response Time/User Think Time Limit Report	9-25
9-13	Terminal Session Report	9-26
9-14	High Terminal Usage Report	9-27
9-15	Opcode Count Report	9-28
9-16	Response Time Report	9-30
9-16.1	H6000-DN355 Reject Report	9-30.1
9-16.2	Abort Report	9-30.2
9-16.3	TS1 Initial Parameter Report	9-30.3
9-16.4	Mailbox Busy Report	9-30.4
9-17	Card Format - Input Option TIME	9-32
9-18	Histogram Reports, Input Option HISTG	9-35
9-19	Terminal Mailbox Dump, Input Option MAIL	9-37
10-1	Default Values For Plots	10-5
10-2	Average Response Time for All Uses by DATANET (Report RESPD)	10-10
10-3	HSLA/Subchannel Transmission Report (Report BUFF)	10-13
10-4	Average Response Time for Designated Lines (Report RESPL)	10-14
10-5	HSLA Subchannels Being Monitored (Report HSLA)	10-15
10-6	Card Images From I* FILE of the GMF Monitor (Report CARD)	10-17
10-7	List of Active Line IDs	10-18
10-8	HSLA Threshold Report (Report Thresh)	10-19
10-9	Number of Transactions Send to Host (Plot 1)	10-23
10-10	Number of 36 Bit Words Received From Host (Plot 4)	10-25
10-11	Percent Idle Time (Plot 5)	10-26
10-12	Percent Buffer Requests Denied (Plot 6)	10-27
10-13	Histogram of DATANET Response	10-29
10-14	Histogram Special Line ID Response	10-30
10-15	Plot Alteration (Action Code PLOT)	10-34
10-16	Set Timespan(s) For Processing (Action Code TIME)	10-36
10-17	Action Codes For Turning Reports Off	10-37
10-18	Action Codes For Turning Reports On	10-38
10-19	Turn Off Plots (Action Code NOPLOT)	10-41
10-20	Threshold Values (Action Code THRESH)	10-43
10-21	Turn Off Reports For Selected DATANETs (Action Code NORPT)	10-44
10-22	Debug Option (Action Code DEBUG)	10-46
10-23	Alter Histogram Parameters (Action Code CHANGE)	10-47
10-24	Unmatched/Duplicate Error Report (Action Code DUPLIC)	10-48
10-25	Reject Report (Action Code REJECT)	10-49
11-1	Sample Execution Report	11-4
11-2	CPU Title Page	11-7
11-3	CPU Queue Length Distribution	11-8

Figure		Page
11-4	CPU Burst Length Distribution for User Activities	11-10
11-5	CPU Burst Length Distribution for All Activities	11-11
11-6	Percent of Memory Time in Queue	11-12
11-7	CPU Time Report	11-13
11-8	CPU Plot of Percent Idle	11-18
11-9	WIN Report	11-19
11-10	CPU Access by SNUMB	11-20
11-11	Tape Reduction Title Page	11-21
11-12	Number of Tape Drives in Use (Time Corrected)	11-22
11-13	Tape Activity Report	11-24
11-14	Tape Status Report	11-25
11-15	Sample Deck Setup	11-36
12-1	TPS GMF Compile JCL	12-4
12-2	Control Cards for Running Only the TPS GMF	12-5
12-3	JCL for TPE Data Reduction	12-7
12-4	TPE Summary Report	12-9
12-5	NONEW Option Report	12-10
12-6	Profile ID Report	12-11
12-7	Buffer Space Report	12-12
12-8	Table Reaching Threshold Size Report	12-14
12-9	Master Terminal Report	12-15
12-10	Impasse Report	12-17
12-11	Output Transmission Report	12-18
12-12	TPAP ABORT Report	12-19
12-13	UNDOT Entry Report	12-20
12-14	Reissuing Intercom Message Report	12-21
12-15	Receiving Intercom Message Report	12-22
12-16	TPAP Modify Report	12-23
12-17	TPAP Status Report	12-24
12-18	TPAP Run Times Report	12-25
12-19	JCL For TPE Formatted Dump	12-28
12-20	TPE Formatted Dumpt	12-29
12-21	Transaction Processing System ALTER Cards	12-31
13-1	New Initialization Section	13-2
13-2	Monitor Subroutine	13-3
13-3	Creation of User Trace	13-5
14-1	Flowchart of Evaluation/Tuning Process	14-4
14-2	Sample Symmetric Distribution	14-12
14-3	Sample Skewed Distribution	14-14
14-4	Memory Statistics	14-15
14-5	Temporary Storage Test Form	14-24
14-6	Large TSS Subsystem Users	14-32
14-7	Poor TSS Response Log	14-33
14-8	High Usage Terminals	14-34
14-9	Terminal/Line Error Report	14-35
15-1	Execution Summary Report	15-8
15-2	Timesharing Reduction Event Log Report	15-9

Figure		Page
15-3	Response Time For All Users Report	15-11
15-4	CPU Times Allocated to Exec and Subdispatch Report	15-12
15-5	Response Time For Users Not Requesting More Core Plot	15-13
15-6	Response Time For Users With Core Requests During Line Idle Plot	15-14
15-7	Total Time In Subdispatch Queue Plot	15-15
15-8	Total Time In Subdispatch Queue Waiting Service Plot	15-16
15-9	Processor Time In Subdispatch Plot	15-17
15-10	TSS Subtraces Encountered Histogram	15-18
15-11	Derails Executed By User Histogram	15-20
15-12	Execution Message Report	15-22
15-13	Users Map	15-23
15-14	TSS Delay and Users Delay Report	15-24
15-15	TSS Core Map	15-25
15-16	Error Message Report	15-26
15-17	Intertrace Gap Duration Plot	15-27
15-18	Memory Activity Plot	15-28
15-19	Eligibles For Subdispatch Plot	15-29
15-20	In Subdispatch Plot	15-30
15-21	Eligibles For and In Subdispatch Plot	15-31
15-22	User Swap Rate Plot	15-32
15-23	User Swap Amount Plot	15-33
15-24	User Swap Duration Plot	15-34
15-25	User I/OS Duration Plot	15-35
15-26	In System Master Catalog (SMC) Wait Plot	15-36
15-27	Intertrace Time (Adjusted For GEWAKES) Histogram	15-37
15-28	Formatted Dump Report	15-39
15-29	Intertrace Time Histogram	15-40
15-30	TEARS JCL	15-63

TABLES

Number		Page
2-1	GMC Release Dependent Parameters	2-11
4-1	Resource Monitor Abort Codes	4-30
5-1	Required Trace Type for Each Monitor	5-5
5-2	Abort Codes	5-6
5-3	GMC Catalog and File Index	5-10
6-1	Default Reports	6-2
6-2	Default Values for Histograms	6-7
6-3	Default Values for Plots	6-9
6-4	Available Report Actions and Their (Default) Values	6-10
6-5	File Code for MUM Reports	6-24
7-1	MSM/MSMDRP Reports	7-2
9-1	Octal Codes and Device Names for the Mailbox and Statistical Summary Reports	9-9
9-2	Acceptable Device Type Names	9-34
10-1	DATANET-355 Reports	10-6
10-2	Response Time Error Messages	10-21
11-1	Central Processing Unit and Tape Reduction Reports	11-3
11-2	Histogram Default Parameters	11-27
11-3	Plot Default Parameters	11-28
12-1	Transaction Processing System Traces	12-2
15-1	Emulation Phase User States	15-6
15-2	Histogram Default Parameters	15-54
15-3	Plot Default Parameters	15-55

ABSTRACT

This Users Manual provides the information necessary for an individual to run all of the programs that are comprised in the Generalized Monitoring Facility (GMF). Program interrelationships are presented, as well as a general overview of the processing, input, and output procedures for each program. Formats and examples of user-controlled input data and sample program outputs are shown and explained. Additionally, the Job Control Language (JCL) deck setups necessary to run the programs are provided.

This manual supersedes Command and Control Technical Center CSM UM 246-81, 1 May 1981.

SECTION 1. GENERAL

1.1 Purpose of the Users Manual

The Users Manual describes each of the programs in the Generalized Monitoring Facility (GMF), discusses input options required to run each program, and provides sample outputs generated by each program.

The Generalized Monitoring Facility is delivered on a FILSYS SAVE tape. A description of the software is presented in section 2. Installation procedures for the GMF Monitor and guidance on how to run GMF can be found in section 5 of this manual.

1.2 Project References

The Generalized Monitoring Facility was originally developed for the Government by Honeywell Information Systems. Since delivery of the completed software in 1975, the Computer Performance Evaluation Branch (C751) of the Command and Control Technical Center has extensively modified and rewritten the GMF system. Numerous software errors have been corrected and many new features have been added.

1.3 Terms and Abbreviations

The following abbreviations will be used throughout the document.

CAM	-	Communications Analysis Monitor
CM	-	Channel Monitor
CPU	-	Central Processing Unit
CPUM	-	CPU Monitor
FTS	-	File Transfer System (a program within the WWMCCS Network System)
GCOS	-	Generalized Comprehensive Operating System
GMC	-	Generalized Monitoring Collector
GMF	-	Generalized Monitoring Facility
GRTS	-	General Remote Terminal Supervisor
GRTM	-	GRTS Monitor
IDLEM	-	Idle Monitor

MSM - Mass Storage Monitor

MUM - Memory Utilization Monitor

NCP - Network Control Program (a program within the WWMCCS Network System)

RMC - Resource Monitor Collector

RMDRx - Resource Monitor Data Reduction Program 1 through 3

RMON - Resource Monitor

TM - Tape Monitor

TPEM - Transaction Processing Executive Monitor

TSS - Timesharing Subsystem

TSSM - Timesharing Subsystem Monitor

WWMCCS - World Wide Military Command and Control System

1.4 Security and Privacy

There are no classified data collected in the GMF, but there is one exception. If the Communications Analysis Monitor (CAM) is being run with the Specific Terminal Option, classified data may be collected.

1.5 Manual Format

Section 2 of this Manual provides an overview of the GMF system. A brief description of each program is included. Sections 3 through 12 and 15 describe the programs of the GMF system in detail. In the presentation the user will find the appropriate information to successfully operate each program. The format of the sections includes a discussion of each program in the input, processing, and output phases. Section 13 of the document describes the procedures that a user should follow if it is desired to create a new GMF monitor and section 14 provides detailed information as to how GMF should be used in conducting a complete system-wide evaluation.

SECTION 2. SYSTEM SUMMARY

2.1 System Application

This section explains the GMF system by logically grouping the programs into two subsystems, the General Monitor Collector (GMC) subsystem and the Resource Monitor Collector (RMC) subsystem. Overviews of the programs in both subsystems are provided. The purpose of the GMC subsystem is to provide a means of collecting detailed information on the operation of the operating system and on the flow of a job through the system. The purpose of the RMC subsystem is to provide a general view of the operation of the system. The RMC can be run on a daily basis to allow the continuous analysis of the system operation. When it appears a problem is occurring, the GMC can be run to further define and resolve the problem area.

2.2 System Operation

Both GMC and RMC monitor GCOS in real time and generate output tapes. These output tapes are then processed through a series of data reduction routines which produce histograms, graphs, and reports which allow an analyst to evaluate system performance. Both RMC and GMC have exclusive data reduction routines (i.e., a tape generated by RMC cannot be reduced by the GMC data reduction program). Differences between the GMC and RMC subsystems pertain to the methods used to collect measurement data.

The RMC is a sampling-based monitor. At 30-second intervals, the RMC enters execution, collects its data, and writes its records to the standard accounting file. Since it samples only at preselected intervals, it cannot present a complete history of what has occurred on the system. However, because the RMC is a sampling-based monitor, overhead is low. During the period of time when RMC is not in execution, it may be swapped out of memory.

The GMC is a trace-based monitor. The various monitors that are comprised in the GMC subsystem are called into execution by the occurrences of the events that are to be captured. The GMC can therefore be used to provide a complete and detailed history of system performances. Because all occurrences of a given event are retrieved, GMC overhead is higher than RMC. Unlike the RMC, GMC cannot be swapped or moved and must be locked into memory. Finally, the GMC requires its own dedicated tape drive for the writing of its generated records.

2.2.1 The Resource Monitor Collection (RMC) Subsystem. The RMC subsystem is composed of a data collector, RMCC, and three associated

data reduction routines. The RMC data collector is explained in section 3, and the RMC data reduction routines are discussed in section 4.

The RMCC is a sampling-based monitor. The RMCC will sample system queues and tables on a 30-second time interval. The data captured from these queues and tables are written to the system accounting file.

The output generated by the RMCC is input to a series of three data reduction routines that are responsible for generating all output reports. See figure 2-1 for the RMC system flowchart and figure 2-2 for the subroutines that comprise the RMC system.

2.2.2 The Generalized Monitor (GMC) Subsystem. The GMC subsystem is composed of a series of data collector programs and a related set of data reduction programs. Each data collector program consists of one or more subroutines, and each program is used to monitor a different area of system performance. The name of the program indicates the area of system performance measured. In addition, any combination of monitoring programs may be executed during a monitor session. A detailed description of the entire data collector facility is given in section 5. Figure 2-3 shows the interrelation of all programs within the GMC subsystem. Figure 2-4 shows the subroutines that comprise each data collector program and those trace types which need to be active for each subroutine.

The mechanism used by the GMC data collector for obtaining control from the operating system is that of the normal system trace. The trace records a history of the occurrence of one or more of 72 systems events, 65 of which are presently defined. This recording is done by the system executing a unique code set resident in the System Dispatcher Module (.MDISP). Execution of this code is common to all system trace events and provides the point at which the GMC obtains control. See section 5 for a detailed description of how GMC gains control from the system.

The executive routine of the GMC processes input cards for the data collection routines, determines which areas of the system are to be monitored, performs any necessary initialization, and controls all data buffering and tape writing. A detailed description of this routine is given in section 5. The associated GMC data reduction programs are described in sections 6 through 12 and section 15.

2.3 System Configuration

The GMF is designed to be run on a HIS 6000 computer system, running with WWMCCS GCOS release 6.4 or 7.2. These releases are equivalent to the HIS commercial 2H or 4JS (any level) GCOS releases. When GMC is

used on WWMCCS release 6.4, or commercial release 2H, the user must insure that the value for variable "SYS64" is changed from its current value of 0 to a value of 1. It should be noted that the CPU, TPE and TSS Monitors may be run only under a WWMCCS 7.2 release or commercial 4JS release. See subsection 2.6 for a complete description of all user requirements prior to using the GMC. The RMC can only be used with a WWMCCS 7.2 (commercial 4JS) GCOS release.

THIS PAGE LEFT INTENTIONALLY BLANK

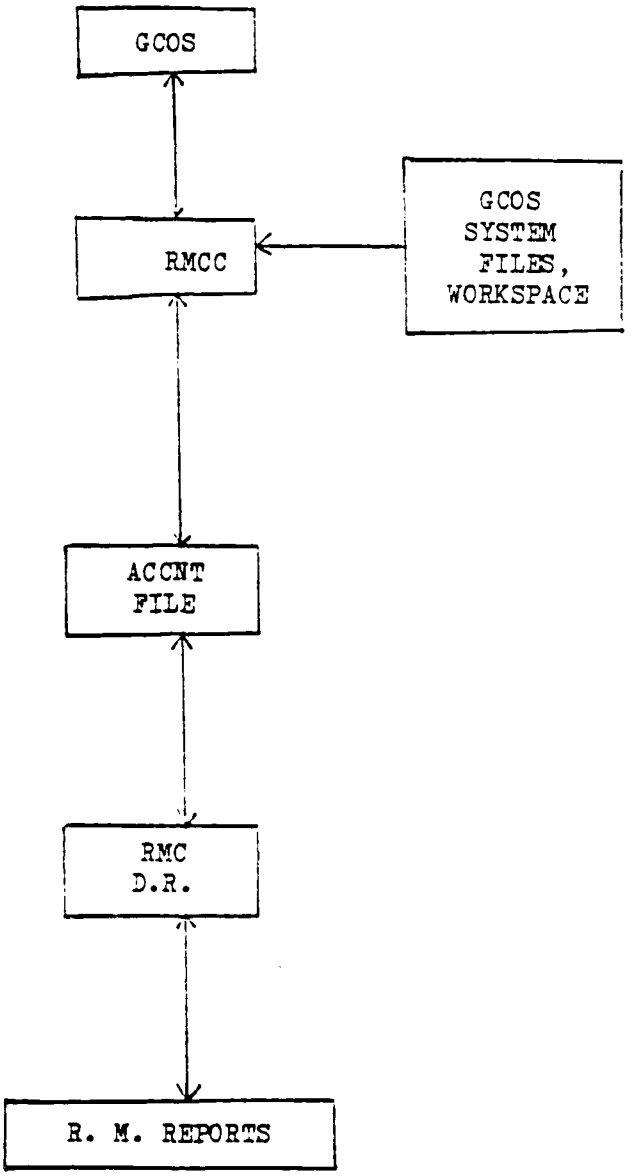


Figure 2-1. RMC Flowchart

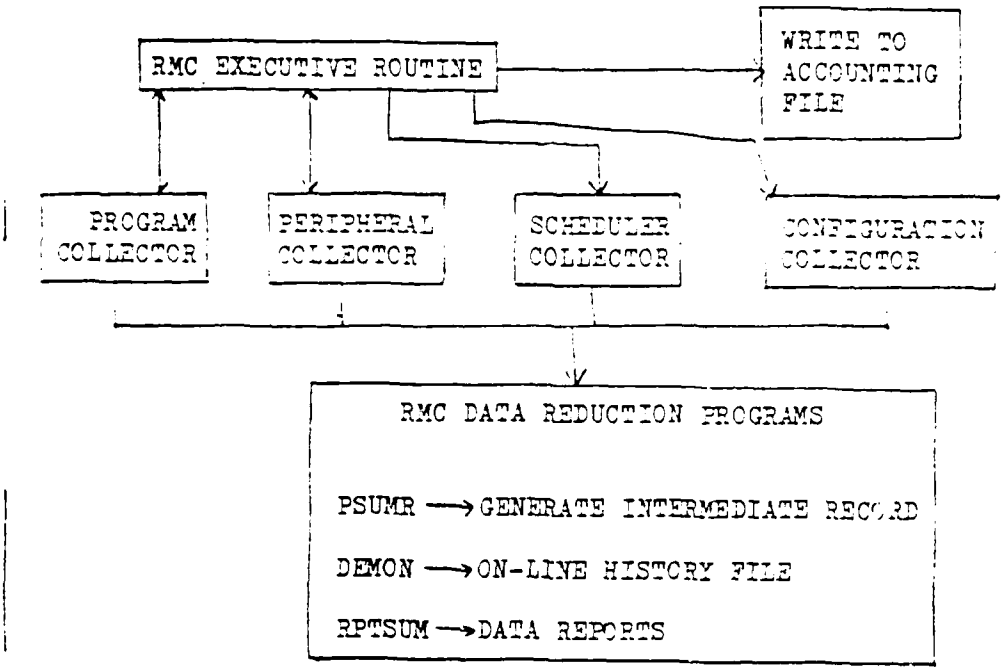


Figure 2-2. Programs in the RMC Subsystem

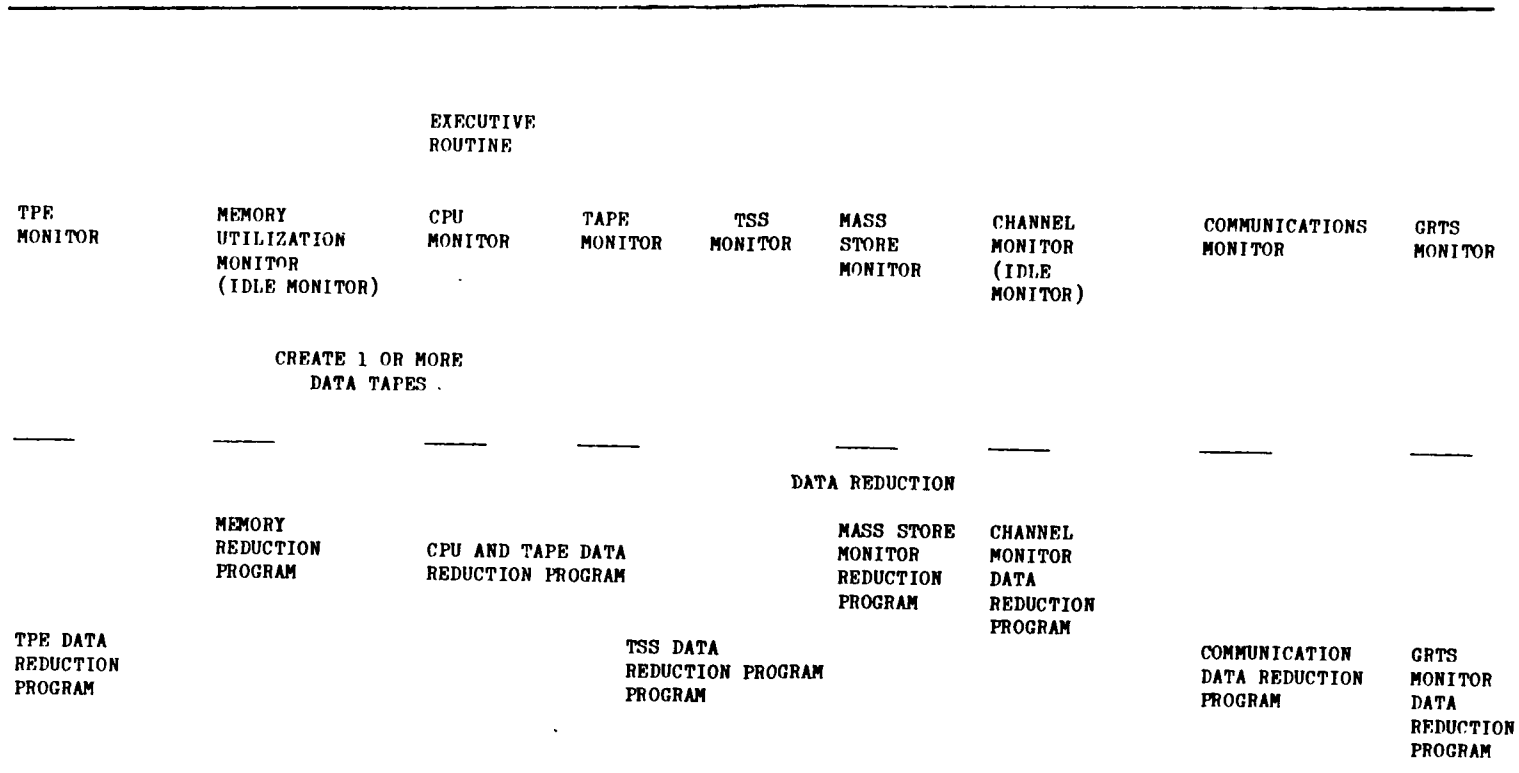


Figure 2-3. Programs in the GMC Subsystem

Data Collector Programs	Subroutines	Traces Captured (OCTAL NUMBER)
Memory Utilization Monitor	T10 T46	10,11,51 46
Idle Monitor	T21 TRCS	21 0,1,2,3,13,16,22,37
Mass Store Monitor	T7	7,15,73*,76*,77*
Channel Monitor	T4,T7,T22	4,7,15,22
Tape Monitor	T50	50,51,52
CPU Monitor	T70	23*,51,63*,70*
Communications Analysis Monitor	T14	14*,15
GRTS Monitor	T62	62*
TPE Monitor	T200	0,1,2,4,5,6,13,42,51,65,74*
TSS Monitor	T100	74*

* - Nonstandard traces generated by the particular monitor.

Figure 2-4. Subroutines and Traces in GMC Data Collector Programs

2.4 System Organization

The GMF is composed of two data collectors, GMC and RMC, and associated data reduction programs. Sections 3 through 12 and section 15 describe these programs. Figure 2-1 gives a system flowchart for the RMC. Figure 5-1, from section 5, gives a system flowchart for the GMC.

2.5 Performance

The GMF monitors the performance of a system, aids in identifying the start of system performance problems, and aids in analyzing system performance problems. The RMC requires very little system resource usage and writes all its data to the system accounting file. The GMC is a much more detailed system with the associated higher overhead. The GMC is used mainly to determine the cause of system performance problems. The GMC requires 15 to 24 thousand words of memory and one tape drive while being run. Both systems require offline data reduction.

2.6 GMF Installation

2.6.1 Creation of GMF Files. The GMF software is contained on a single user save tape. The USERID on the tape is B29IDPXO. This USERID must be created with 5474 LLINKS of file space. A user restore can then be run. B29IDPXO is subdivided into several catalogs described below:

. GMFCOL - 535 LLINKS - This subcatalog contains all the data collection software for the GMC monitoring system. All files within this subcatalog are completely described in section 5.

. SOURCE - 3040 LLINKS - This subcatalog contains Time Sharing source files for all data reduction programs contained within the GMC system. Figure 2-5 is a breakdown of the individual files within this subcatalog. Sections 6-12 and 15 describe each program in detail.

. OBJECT - 1315 LLINKS - This subcatalog contains the object decks for all the data reduction programs contained within the GMC system. Figure 2-5 is a breakdown of the individual files within this subcatalog.

. JCL - 14 LLINKS - This subcatalog contains all the JCL required to run all the data reduction programs contained within the GMC system. Figure 2-6 is a breakdown of the individual files within this subcatalog.

. NEWRMON - 570 LLINKS - This subcatalog contains all the software required to collect and reduce the data for the NMON Monitoring

FILE NAME	FUNCTION	SOURCE SIZE (LL)	OBJECT SIZE (LL)
MUM	MEMORY UTILIZATION MONITOR DATA REDUCTION PROGRAM. REFERENCED IN SECTION 6.	380	216
MSM	MASS STORE MONITOR DATA REDUCTION PROGRAM. REFERENCED IN SECTION 7.	321	188
CM	CHANNEL MONITOR DATA REDUCTION PROGRAM. REFERENCED IN SECTION 8.	258	180
CAM	COMMUNICATION MONITOR DATA REDUCTION PROGRAM. REFERENCED IN SECTION 9.	198	100
CPU-TAPE	CPU AND TAPE MONITOR DATA REDUCTION PROGRAMS. REFERENCED IN SECTION 11.	350	152
GRT	DATANET 355 MONITOR DATA REDUCTION PROGRAM. REFERENCED IN SECTION 10.	280	142
TPETG	TRANSACTION PROCESSING DATA REDUCTION PROGRAM. REFERENCED IN SECTION 12.	64	76
TPEALT	AN ALTER FILE FOR ADDING TPE TRACE CODE INTO THE TPE SUBSYSTEM (NO OBJECT FILE). REFERENCED IN SECTION 12.	620	
TPEDUMP	A PROGRAM FOR OBTAINING A FORMATTED TRACE DUMP FROM A TPE/GMF DATA TAPE. REFERENCED IN SECTION 12.	40	24
TSS	TIMESHARING MONITOR DATA REDUCTION PROGRAM. REFERENCED IN SECTION 15.	529	237

Figure 2-5. B29IDPX0/SOURCE and B29IDPX0/OBJECT Catalog Structure

FILE NAME	FUNCTION	SOURCE SIZE (LL)
MUM	JCL TO OBTAIN ALL MEMORY UTILIZATION MONITOR REPORTS	2
MSM	JCL TO OBTAIN MASS STORE MONITOR REPORTS	2
CM	JCL TO OBTAIN CHANNEL MONITOR REPORTS	2
CAM	JCL TO OBTAIN COMMUNICATION MONITOR REPORTS	2
GRT	JCL TO OBTAIN DN-355 MONITOR REPORTS	2
CPU-TAPE	JCL TO OBTAIN CPU AND TAPE MONITOR REPORTS	2
TPETG	JCL TO OBTAIN ALL REPORTS FROM THE TPE DATA REDUCTION PROGRAM	2
TSS	JCL TO OBTAIN ALL REPORTS FROM THE TSS DATA REDUCTION PROGRAM	2

Figure 2-6. B29IDPXO/JCL Catalog Structure

system. This subcatalog is further subdivided into JCL and SOURCE subcatalogs. The files within these subcatalogs are completely described in sections 3 and 4.

2.6.2 GMF Release Dependent Parameters. In order for the GMC to operate properly, it is necessary for GMC to locate certain instructions and/or words within several system programs. The user should insure that these locations are correct for the particular GCOS release, under which he is operating. Table 2-1 is a list of these dependent parameters identifying their the use and providing the approximate program source line numbers where the particular parameters are used. The list is provided for each GMC program that must checked by the user.

In order to use the GMC data reduction programs on a WW6.4 system, there is a special data card required in certain programs. This option applies to all data reduction programs except CPU-TAPE, TPETG and TSS. These programs are not designed for use under any release other than WW7.2 or commercial 4JS.

The GMC system is designed so that data collected on a WW6.4 system may be reduced under a WW6.4 system or a WW7.2 system. In addition, data collected under a WW7.2 system may likewise be reduced under a WW7.2 system, or a WW6.4 system. Whenever the data reduction programs for MUM, CM, CAM, or GRT are used on a WW6.4 system, a data card with an RN typed on it must be included in the input section of the JCL deck. It makes no difference under what release the data was collected. It is only a question of under what release the data is being reduced.

For the MSM data reduction program, there are two data cards required. The first data card always contains the letters RN. The second card is determined by the following table:

Data Collected	Data Reduced	Data Value
WW6.4	WW6.4	1
WW6.4	WW7.2	3
WW7.2	WW6.4	2
WW7.2	WW7.2	NO SPECIAL CARDS REQUIRED

Table 2-1. GMC Release Dependent Parameters
(Part 1 of 10)

Element	Program	LINE #	Variable	Explanation
1	GMF.TOP	100	SYS64	Used to control conditional assembly of GMC set=1 for W6.4(2H) release set=0 for W7.2(4J) release
2		10240-10740	-	Code in this area searches for trace processing within the dispatcher. Trace code must be within 500 octal locations of the address specified by entry point 15 decimal of the dispatcher. This entry point should contain the address of location TRACE within the dispatcher, which is where the trace processing code is located. The code being searched for is an LDAQ;STAQ;TRA 0,1. If these instructions are not found, GMF will abort with an NF abort. In order for GMF to operate, the $TRACE card in the boot deck must not specify no tracing. At a minimum, the $TRACE card must request at least a single trace to be captured. It is recommended that if a site wants to minimize trace logic overhead, but at the same time be able to run GMF without altering the boot deck, that the following trace card be used: $ TRACE 777777777777,643777777777
3		10760-11180	-	Code in this area is used to make a correction to accounting processing, if the correction has not already been made via patches. The code is searched for within 500 octal locations of .MIOS entry point. The code searched for is SBLA TRREG+7,$;ARL 12; ADLA .CRTOD,7. The ARL is changed to an ARS.
4	MUM.INIT	500-740	-	Code in this area searches the dispatcher for generation of trace types 10 and 11. The code searches

Table 2-1. (Part 2 of 10)

Element	Program	LINE #	Variable	Explanation
				for an ORA =010,DL instruction and an XED .CRTRC instruction 3 locations in front of the ORA instruction. If the XED is not found, then a NOP instruction is looked for and will be replaced with an XED instruction. If neither the XED or NOP is found, the GMF will abort with a Z0 abort. The same logic is repeated, except this time an ORA =011,DL is searched for. These instructions must be located within 1500 octal locations of entry point 15 decimal of .MDISP. This entry point should be the address of location TRACE within the dispatcher.
5	MSM.INIT CM.INIT CAM.INIT	400-600 380-570 220-420	-	Code in these areas searches .MFALT for generation of trace type 15. The code searches for an ADLA =015,DL instruction and an XEC .CRTRC+2 instruction 6 locations in front of the ADLA instruction. If the XEC instruction is not found, then a NOP instruction is looked for and will be replaced with an XEC instruction. If neither the NOP or XEC is found, the GMF will abort with a Z4 abort. These instructions must be located within 1000 octal words of the entry point of .MFALT.
6	MSM.INIT CM.INIT	610-820 580-790	-	Code in these areas searches .MIOS for generation of trace type 7. The code searches for an ORA 7,DL instruction and an XED .CRTRC instruction 7 locations in front of the ORA instruction. If the XED instruction is not found, then a NOP instruction is looked for and will be replaced with an XED instruction. If neither the NOP or XED is found, the GMF will abort with a Z1 abort. These instructions

Table 2-1. (Part 3 of 10)

Element	Program	LINE #	Variable	Explanation
				must be located within 7000 octal locations of the entry point of .MIOS.
7	CM.INIT	1690-1890	-	Code in this area searches .MIOS for generation of trace type 4. The code searches for an ARL 18 instruction and an XED instruction 3 locations in front of the ARL instruction. If the XED instruction is not found, then a NOP instruction is looked for and will be replaced with an XED instruction. If neither the NOP or XED is found, the GMF will abort with a Z3 abort. These instructions must be located within 1000 octal locations of the entry point of .MIOS.
8	CM.INIT	1900-2100	-	Code in this area searches .MIOS for generation of trace type 22. The code searches for an ARL 12 instruction and an XED instruction 3 locations in front of the ARL instruction. If the XED instruction is not found, then a NOP instruction is looked for and will be replaced with an XED instruction. If neither the NOP or XED is found, the GMF will abort with a Z2 abort. These instructions must be located within 6000 octal locations of the entry point of .MIOS.
9	CAM.INIT	1110-1520	-	Code in this area searches .MDNET to locate the CMBX0 table. The code searches for a BCI 1,C instruction within 5000 decimal locations of the entry point of .MDNET. If GMF cannot locate this instruction, GMF will abort with an NS abort. If this instruction is found, then a BCI 1,TSS instruction is searched for within 30 decimal locations of

Table 2-1. (Part 4 of 10)

Element	Program	LINE #	Variable	Explanation
				the previous instruction. If the second instruction is not found, GMF will abort with an NT abort. If this instruction is found, then an LDA T.55,2 instruction (000001235212) is searched for within 10 decimal locations of the previous instruction. If this instruction is not found, GMF will abort with an NU abort.
10	IDL.INIT	300-1620	-	Code in this area searches for the generation of a series of trace types using the same logical construct as previously described for the other INIT sections. Following is a list of instructions searched for: (1) LDX2 .CRTRC+3 and an XEC .CRTRC+2 3 instructions in front (2) LDX2 .CRTRC+3 and an XEC .CRTRC+2 3 instructions in front (3) ARL 12 and an XED .CRTRC 3 instructions in front (4) ORA =016,DL and an XEC .CRTRC+2 3 instructions in front (5) ORA =037,DL and an XED .CRTRC 4 instructions in front (6) ORQ =013,DL and an XEC .CRTRC 7 instructions in front (7) ORA =021,DL and an XEC .CRTRC+2 4 instructions in front These areas of code search for generation of trace types 0,3,22,16,37,13 and 21, respectively, and if appropriate code is not found, the following respective aborts will occur: Z6,Z6,Z7,Z8,Z9,ZA and ZB. The modules which are searched are listed below: (1) 1000 octal locations from the entry point of .MIOS (2) 1000 octal locations from the

Table 2-1. (Part 5 of 10)

Element	Program	LINE #	Variable	Explanation
				entry point of .MIOS (3) 6000 octal locations from the entry point of .MIOS (4) 1000 octal locations from the entry point of .MFALT (5)-(7) 6000 octal locations from the entry point of .MDISP
11	CPU.PAT	300-440	-	Code in this area searches for an ASA .SALT,5 instruction in the dispatcher. Code must be within 300 octal locations after the address specified by entry point 20 decimal of the dispatcher. This entry point should contain the address of location DWAIT within the dispatcher. The trace code to be inserted by GMF indicates that a job is being taken out of processing. If the ASA instruction is not found, an N1 abort will occur.
12		460-590	-	In a WW7.2 (4JS) system, code in this area searches for an STQ .QTOD,4 instruction. Search area is the same as described for line 300. This trace indicates that subdispatching has finished using the processor. If the STQ instruction is not found, an N8 abort will occur.
13		640-800	-	If the dispatcher queue option of the CPU Monitor is activated, code in this area searches for an ORSA .STATE,4 instruction followed by an LDA .STATE,4 instruction. Code must be within 100 octal locations after the address specified by entry point 7 of the dispatcher. This entry point should contain the address of location DSPQH within the dispatcher. This trace indicates that a job is being placed into the

Table 2-1. (Part 6 of 10)

Element	Program	LINE #	Variable	Explanation
				dispatcher queue. If these instructions are not found, an NA abort will occur.
14		810-970	-	If the dispatcher queue option of the CPU Monitor is activated, code in this area searches for a LCQ=010001,DL instruction followed by an ANQ .STATE,5 instruction. These instructions must be within 1500 octal locations after the address specified by entry point 1 of the dispatcher. This entry point should contain the address of location DSP within the dispatcher. This trace indicates that a job is being taken out of the dispatcher queue and placed into execution. If these instructions are not found, an NB abort will occur.
15		970-1310	-	In order to implant its special hooks, the CPU Monitor must modify the dispatcher and, therefore, it requires eight words of patch space. If the dispatcher queue option of the CPU Monitor is activated, then 16 words, instead of the normal 8, are required. This patch area must be within 200 octal locations in front of the address specified by entry point 15 decimal of the dispatcher. This entry point should contain the address of location TRACF within the dispatcher.
16		2190-2320	-	If sufficient patch space was not available in the standard patch area, the CPU Monitor will attempt to locate patch space in a specially defined user patch area. This search will take place only if bit 2 of word 0 within the dispatcher is set. This patch space should be within 200 octal locations after the

Table 2-1. (Part 7 of 10)

Element	Program	LINE #	Variable	Explanation
				address ILIST within the dispatcher. The address of ILIST is found in word 7 of the dispatcher. If sufficent patch space is not found, an N2 abort will occur.
17		1950-2360	-	Code in this area searches .MFALT for GATE LOOP timing logic and attempts to correct logic errors that exist within the code. The code searches for an ARL 12 and ASA.CROVH,7 instruction within 2000 octal locations of entry point 13 of .MFALT. This entry point should be the address of the location BOOT within .MFALT. Twelve instructions in front of the ARL instruction, a LDX6.CRPRG,7 instruction is searched for. If this instruction is found, the code will next obtain the address stored in the upper half of the word located two locations in front of the entry point of .MFALT. This should be the address of the MMECNT table. The code will add 46 decimal to this address and should now contain the address of a patch table. The code will search for nine free words, and if found will correct the logic error. If patch space is not found, the code will remain unaltered.
18 19	CAM.PAT	170-290	-	Code in this area searches for an LDQ M.LID,3 instruction, followed by an ANQ=0077777,DU instruction in module DNWW/DNET. The search for this code begins at the address specified by entry point -8 of DNWW/DNET and continues until the code is located. This entry point should contain the address of location NRQWT-DNET within DNWW/DNET. If the instructions are not found, an N3 abort will occur.

Table 2-1. (Part 8 of 10)

Element	Program	LINE #	Variable	Explanation
20		350-500 and 690-820	-	In order to implant its special hooks, the CAM must find 8 words of patch space. Even though the CAM is modifying module DNET, it will use the patch space available in the .MDISP module. This search for space is identical to that for CPU.PAT at lines 970-1310 and lines 2190-2320. If patch space is not found, an N4 abort will occur.
21	MSM.PAT	200-240	-	Code in this area searches the dispatcher for SSA cache code. If bit 4 of word 0 in the dispatcher is set, then cache is available. If this bit is not set, then no further searches are performed.
22		250-380	-	Code in this area searches the dispatcher for the location of DBASE. The address at entry point -2 is obtained. This address points to the location ILIST in the dispatcher. At location ILIST, a series of addresses are stored and MSM.PAT searches this list for the address DBASE. If DBASE is not found, an N5 abort will occur.
23		410-620	-	Code in this area searches for an AOS .CRTDL and an AOS .CRTBH instruction. This code needs to be within 300 octal locations after the address DBASE within the dispatcher. If this code is not found, an N5 or N6 abort will occur.
24		670-790 and 1120-1250	-	In order to implant its special hooks, the MSM Monitor must modify the dispatcher and, therefore, it requires 8 words of patch space. This search is identical to that for CPU.PAT at line 680 and lines 970-1310 and 2190-2310. If patch space is not found, an N7 abort will occur.

Table 2-1. (Part 9 of 10)

Element	Program	LINE #	Variable	Explanation
25	GMF.MON	1040	FMS1	Offset from entry point of .MFSIO which points to the word giving the absolute address of FMS catalog cache buffer. Used only in W7.2. Set to -13 decimal.
26		1050	FMS2	Offset from entry point of .MFSIO pointing to the word which gives the option selection for FMS catalog cache. Used only in W7.2. Set to -15 decimal.
27	MUM.T10	220	SYS64	See GMF.TOP
28		920	FIFO	Address of the FIFO buffer within PALC. It is used to search the JCT table of PALC. This includes adding in a 110 octal offset for the loading of PALC in W6.4. There is no PALC offset in W7.2.
29		5410	XPQ24	Location in CALC of the memory demand table. Set to octal 111.
30		5420	SLVSNB	Offset in slave prefix area of job SNUMB. Set to octal 36.
31		5430	MEMUSE	Offset in slave prefix area of loader memory use word. Set to octal 37.
32		5440	IDENT	Offset in slave prefix area of job IDENT. Set to octal 66.
33	CM.T07A	190	IDENT	Offset in slave prefix area of job ident. Set to octal 66.
34		210	SYS64	See GMF.TOP
35		9960-10130	-	Code in this area searches .MFSIO in order to gather statistics for FMS catalog cache analysis. All the following references are offsets from the entry point of .MFSIO:

Table 2-1. (Part 10 of 10)

Element	Program	LINE #	Variable	Explanation
				-12 # of cache hits -11 # of writes -10 # of reads 841 # of reads not in CC 842 # of 320-word reads 843 # of skips 844 # of cache clears 848 # of no hits 849 # of hits
36		10180-10250	-	Code in this area searches .MASO4 in order to gather statistics concerning the available space table utilization. All the following references are offsets from the entry point of .MASO4: -6 # of times buffer allocation attempted -5 # of times buffer busy -4 # of times AST was in memory -3 # of times AST in memory but busy
37		11440	FFCCC	Address in PALC where the file code is stored during GEFSYE processing. Set to 6177 octal in W6.4 and 13143 octal in W7.2. This includes 110 octal for loading of PALC in W6.4. There is no offset for PALC in WW7.2.
38		11450	SNUMBP	Address in PALC where the SNUMB is stored during GEFSYE processing. Set to 35012 octal in W6.4 and 2632 octal W7.2. This includes 110 octal for loading of PALC in W6.4. There is no offset for PALC in WW7.2.
39		11460	ACT	Address in PALC where the activity number is stored during GEFSYE processing. Set to 33231 octal in W6.4 and 1051 octal in W7.2. This includes 110 octal for loading of PALC in W6.4. There is no offset for PALC in W7.2.

SECTION 3. RESOURCE MONITOR COLLECTOR

3.1 Introduction

The Resource Monitor Collector (RMC) is a privileged software package which periodically samples the system programs, queues, counters and tables of GCOS. The RMC consists of initialization routines, general purpose routines and four discrete data collector routines. The four data collectors include the following:

- o Program collector
- o Peripheral collector
- o Schedular collector
- o Configuration collector

Each of the data collectors generates a unique record format in one of the common buffer areas. A .CALL to .MSCF,2 is made to write each record to the system accounting file.

3.2 RMC Input Options and JCL

There are currently no RMC input options. The JCL needed to execute the RMC is shown in figure 3-1. Since the RMC runs in master mode, it requires a $PRIVITY card, meaning the operator must GRANT the job.

3.3 Processing

The RMC requires no special tapes or disk files. All data is written to the Statistical Collection File (SCF). The RMC will swap out of core if required and produce low system overhead. The next subsections discuss the collection routine of the RMC and give the record formats.

3.3.1 Collection Routine. The RMC collection routine is a GMAP program requiring PRIVITY. As currently released, the program uses SCF record number 608 for its data record.

The type 608 record consists of 1 of 4 subtypes. The record size depends on the system configuration. The system configuration record (subtype 4) is written once per hour. The RMC writes the other 3 type 608 records every 30 seconds. The RMC initializes its tables according to the system configuration. Extended memory instructions are NOPed if required, any excess memory is released and the ** file is released. The RMC then begins normal processing writing 3 SCF records every 30 seconds.

The following procedure applies to the WWMCCS software releases. The user should ensure that the 600 class of SCF records has been turned on for his

```
$ IDENT
$ USERID
$ PROGRAM COLL
$ PRMFL **,R,R,B29IDPXO/NEWRMON/DDYN
$ LIMITS 999,10K,,500
$ PRIVITY
$ ENDJOB
```

Figure 3-1. RMC JCL

system. He must check two different places. First, the $SCFBUF card in the boot deck must contain at least a C6 to indicate collection of 600-level SCF records, or else the specific record type 608 should appear on the card.

The user must next ensure that the system MASK cards are defined correctly. The 600-class SCF records must be turned on in INIT by use of a MASK card in the boot deck. The location to be changed with the MASK card is RMAX+6 in INIT. The location should be changed with the following MASK card:

```
COL 1           8           13              73
    OCTAL       MASK        000012000000    .MINIT
    LOCATION
    OF RMAX+6
```

This card allows 10-type 600 SCF records to be collected (600-609). These same procedures should be followed by the user if he also changes the RMON SCF record number. Commercial 4Jx sites should check the System Startup Manual for procedures on defining new SCF record types. The user should then enter the command SCFRST 608 at the system console. The system should respond with: 608/AC to indicate that the records are being collected.

When used at a commercial site, the user must first execute the file B29IDPXO/NEWRMON/JCL/CUPJCL. Upon the successful completion of this step, the commercial user should then follow all standard operating procedures.

3.3.1.1 Program Data - Subtype 1. This data type contains overhead and idle times for all configured processors, size of memory, number of processors and number of TSS users. Each job in the system is examined and its status is saved. The subtype 1 record format is:

Word	Bits	Content	Source
0	0-17	Size	
	18-35	Record Type	
1-9	0-35	Standard SCF header	
10	0-29	Record sequence number	
	30-35	Subtype (=1)	
11	0-17	Current available memory (.CRPMU)	
	18-35	Number TSS users (.TCNOU)	
12	0-17	Unused	
	18-35	Number remotes connected (.CRCGT)	
13	0-35	System processor overhead (sum of all entries in .CROVH table)	
14	0-35	Processor idle time (sum of all entries in .CRIDT table)	
15	0-29	SNUMB (.CRSNB)	
	30-35	Unused	
16	0-17	Status (CRSN1)	
	18-35	Memory size if in memory (.SMSC)	

17	0-35	Accumulated processor time = -1 = not in memory (.SPRT)

Words 15-17 are repeated for all jobs in the system.

3.3.1.2 Peripheral Data - Subtype 2. This data type consists of channel use time for each channel, device status (released, assigned, dedicated, permanent, removable), and capacity of disk packs. The subtype 2 record format is:

Word	Bits	Content	Source
0	0-17	Size	
	18-35	Record type	
1-9	0-35	Standard SCF header	
10	0-29	Record sequence number	
	30-35	Record subtype (=2)	
11	0-35	Number available LLINKS for PRM devices (AST)	
12	0-35	Number available LLINKS for RMVB devices	
13	0-17	Number of two-word entries which follow	
	18-35	Not used	
14	0-5	Device type code (.CRSCT)	
	6-11	Number of removable devices on this channel (disk devices)	
	12-17	Number free	
	18-23	Number allocated	
	24-29	Number dedicated	
	30-35	Number released	
15	0-35	Channel use time (channel module + 8)	

Words 14-15 are repeated for each channel.

3.3.1.3 Scheduler Data - Subtype 3. This data type consists of system scheduler data. The subtype 3 record format is:

Word	Bits	Content	Source
0	0-17	Size	
	18-35	Record type	
1-9	0-35	Standard SCF header	
10	0-29	Record sequence number	
	30-35	Record subtype (=3)	
11	0-17	Number jobs allowed in class (CLSCAT)	
	18-35	Number jobs catalogued	
12	0	=1 = JEND on class (CLSJOB)	
	1-17	= Max number concurrent allowed	
	18-35	Number currently running	

Words 11-12 are repeated for each scheduler class

11+2n+1	0-35	= -1 fence

3.3.1.4 Configuration Data - Subtype 4. This data type consists of system configuration data. The subtype 4 record format is:

Word	Bits	Content	Source
0	0-17	Size	
	18-35	Record type	
1-9	0-35	Standard SCF header	
10	0-29	Record sequence number	
	30-35	Record subtype (=4)	
11	0-2	Number of IOMS	
	3-5	Number of processors	
	6-23	Configured memory	
	24-35	Hard core in KWDS	
12	0-8	Number tape channels	
	9-17	Number tape drives	
	18-26	Number printer channels	
	27-35	Number printers	
13	0-35	Number disk channels	
14	0-35	Number disk information entries	
15	0-5	Device type code	
	6	=1 = removable	
	7-17	Allocation unit	
	18-35	Maximum LLINKS for device	
16	0-17	Number defective LLINKS	
	18-35	Available table units	

Words 15-16 are repeated for each disk device configured.

3.4 Outputs

The RMC puts all of its output to the SCF.

3.5 Catalog Description

The RMC and data reduction files are located under the catalog B29IDPXO/NEWRMON. The resource monitor uses FILEDIT to maintain the source and object files. The programs are executed from a ** file maintained by SYSEDIT.

3.5.1 Files

3.5.1.1 FILEDIT Source File. The FILEDIT source file for all programs is:

B29IDPXO/NEWRMON/DSRC

3.5.1.2 FILEDIT Object File. The FILEDIT object file for all programs is:

B29IDPXO/NEWRMON/DOBJ

3.5.1.3 SYSEDIT Loadable File. The SYSEDIT loadable file (Q*) for all programs is:

B29IDPXO/NEWRMON/DDYN

3.5.1.4 JCL. The JCL to update and execute all programs is found under the following catalog:

B29IDPXO/NEWRMON/JCL

3.5.1.4.1 DUPJCL. File DUPJCL contains the JCL to update the source, object, and system loadable files.

3.5.1.4.2 DALT. File DALT will contain the FILEDIT directives when an update is desired. The user must place the directives on DALT prior to running DUPJCL.

3.5.1.4.3 COLL. File COLL contains the JCL for executing the collector.

3.5.1.4.4 PSUMR. File PSUMR contains the JCL for executing the intermediate record create program (PSUMR) and the daily monitor (DEMON).

3.5.1.4.5 RPTSUM. File RPTSUM contains the JCL for executing the report program (RPTSUM).

3.5.1.4.6 CUPJCL. File CUPJCL contains the JCL to update the source, object and system loadable files for commercial execution of the RMON system. This file must be processed prior to commercial use of this monitor.

3.5.2 Update Procedures. When changes are required to any of the resource monitor programs (besides those required to make RMON compatible to a commercial 4JS release), the FILEDIT directives are saved on B29IDPXO/NEWRMON/JCL/DALT. The user then runs the JCL found under file DUPJCL.

Sample TSS session to update the collector (after logon):

```
User:    Card N
System:  *
User:    10 $:MODIFY:source,object,COLL
         20 $:GMAP:::COLL
         30 $:UPDATE:LIST
         40 $:ALTER:10,10
            EOM
System:  *
User:    Resave B29IDPXO/NEWRMON/JCL/DALT
```

```
System:  Data saved - DALT
User:    OLD B29IDPXO/NEWRMON/JCL/DUPJCL
System:  *
User:    Run
System:  Card format, disposition
User:    S
System:  Tab character, settings
User:    N
System:  SNUMB = XXXXT
User:    Bye.
```

SECTION 4. RESOURCE MONITOR DATA REDUCTION

The Resource Monitor Data Reduction (RMDR) is composed of three discrete programs. The first program (PSUMR) processes the SCF records producing an intermediate record. The second program is a daily monitor (DEMON) which maintains an on-line history file. The third program (RPTSUM) produces reports and prints the required plots. Figure 4-1 is an overview of the RMDR.

4.1 PSUMR

PSUMR selects the collector data records from an SCF input file. The records are sorted by system ID, date, time, and record subtype. The records are then processed in pairs to produce an intermediate record.

4.1.1 PSUMR Inputs. PSUMR requires one input file containing SCF records. Two optional inputs are the intermediate record - old master, and a parameter file.

4.1.1.1 SCF Records (File IN). Contains resource monitor collector (RMC) SCF records.

4.1.1.2 Intermediate Record - Old Master (File OM). This optional input will be copied to the output file before new records are added from the current jobs processing.

4.1.1.3 Parameter File (File PF). PSUMR will process the following parameter language statements:

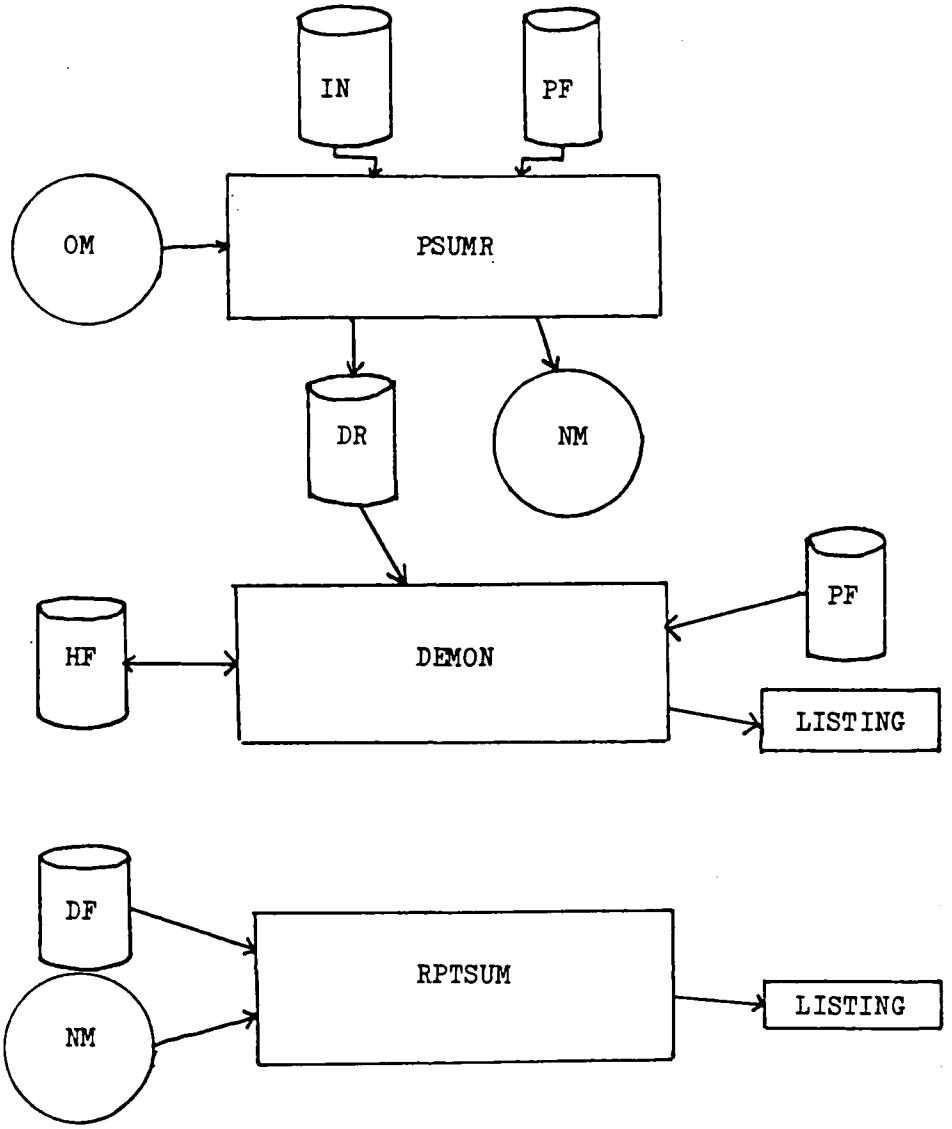
```
CONFIG
SUMMARY
SNUMB
```

The format and use of the parameter language is found in section 4.4.

4.1.2 PSUMR Output. Three outputs are available from PSUMR: (1) listing of the parameter file if present; (2) intermediate record - new master file; (3) intermediate record - current file.

4.1.2.1 Parameter File Listing (File RP). This is a listing of the contents of file PF. Any of the statements processed which were invalid will be flagged. Invalid contents will not halt the job execution, but may affect how the data is summarized in the intermediate record.

4.1.2.2 Intermediate Record New - Master File (File NM). This file contains the contents of the old master file, if present, plus the records added from the current execution of the program.



IN = SCF accounting records
PF = optional parameter file
OM = old master intermediate record
NM = new master intermediate record
DR = optional intermediate output for DEMON
HF = daily monitor history file

Figure 4-1. RMDR Overview

4.1.2.3 Intermediate Record - Current File (File DR). This file contains the intermediate records created during the current execution of the program. This file is used to pass the intermediate records to the daily monitor (DEMON), when present, in a subsequent activity.

4.1.3 PSUMR Deck Setup. The following control cards are required to execute PSUMR:

```
$     IDENT       ACCOUNTING INFORMATION
$     USERID      USERID$PASSWORD/SCC
$     PROGRAM     PSUMR,Dump
$     LIMITS      20,20K,,5K
$     PRMFL       **,R,R,B29IDPXO/NEWRMON/DDYN
$     DATA        PF            optional parameters
$     TAPE        OM            optional old master intermediate
                                record
$     TAPE        NM            new master intermediate record
$     FILE        DR            current file output
$     SYSOUT      RP            parameter file listing
$     TAPE        IN            SCF input file
$     FILE        S1,S1R,50R    sort files
$     FILE        S2,S2R,50R
$     FILE        S3,S3R,50R
$     FILE        S4,S4R,50R
$     ENDJOB
```

The sort files' size should be increased or decreased depending upon the input volume. Tape files may be replaced by permanent or temporary disk files. File DR may be null if the daily monitor does not follow in a subsequent activity.

4.2 DEMON

The daily monitor accepts the intermediate records created by PSUMR. A history file will be initialized (if current null) or updated from the data in the intermediate record. A matrix structure is built to summarize and store the data. The matrix contains 96 entries corresponding to the time period to be summed together. The entry size is variable depending upon the graphs to be produced. For DEMON, the default is all graphs. The history file contains a copy of this matrix structure. For an update run, DEMON copies the history file into core. A second matrix is built for the current day's data. Both matrices are then updated from the input. The history matrix is written to disk. Finally, the current input matrix is compared to the history matrix. If 25% of all data items are not within 25% of the history matrix value, the program generates a complete set of graphs to indicate a significant change from the history file value.

4.2.1 DEMON Inputs. DEMON has one required input and two optional inputs.

4.2.1.1 History File (File HF). The history file is required for DEMON execution. The history file is a random file which contains an image of the in-core matrix used for summarizing and comparing data.

4.2.1.2 Parameter File (File PF). The parameter file is optional. However, if PSUMR was supplied with a SUMMARY statement, the DEMON should be supplied the same summary statement for desired results. DEMON processes only a summary statement. The format and use of the parameter language statements is found in section 4.4.

4.2.1.3 Intermediate Records (File IN). This is an optional input. If present, the intermediate record data will be summarized, as described above, according to parameter setting marked on the intermediate record. The history file will be initialized or updated as it is appropriate. If file IN is not present, the history file data will be used to generate a set of graphs. This feature allows the user to get a current historical "picture" at any time.

4.2.2 DEMON Outputs. DEMON outputs are the updated history file and a listing file.

4.2.2.1 Updated History File (File HF). Refer to section 4.2.1.3.

4.2.2.2 Listing File (File RP). The output listing will contain a listing of the parameter file, if any, and any generated graphs. The graphs produced are listed and described in section 4.5..

4.2.3 DEMON Deck Setup. The following control cards are required to execute DEMON.

```
$    IDENT      ACCOUNTING INFORMATION
$    USERID     USERID$PASSWORD
$    PROGRAM    DEMON,DUMP
$    LIMITS     10,24K,,10K
$    PRMFL      **,R,R,B29IDPXO/NEWRMON/DDYN
$    DATA       PF       optional parameter file
$    TAPE       IN       optional intermediate file
$    PRMFL      HF,W,R, (history cat/file)
$    SYSOUT     RP
$    ENDJOB
```

File IN may be a permanent or temporary disk file. File HF is 24 random LLINKS.

4.3 RPTSUM

PSUMR has the ability to maintain a master file of intermediate records. RPTSUM will process this file to produce any number of reports. Each report

consists of one to eight graphs covering arbitrary time frames. The user specifies, in a parameter file, the reports required.

4.3.1 RPTSUM Inputs. RPTSUM has one required input and one optional input.

4.3.1.1 Intermediate Master Record (File IF). Typically, this required file would contain records covering an extended time (weeks or months). RPTSUM will identify the records which fall within the time limits for each report. The file is processed only once per execution, regardless of the number of reports required.

4.3.1.2 Parameter File (File PF). RPTSUM produces reports based on statements from an optional parameter file. If a parameter file is not supplied, all data on the IF file will be processed. The format and use of the parameter language is found in section 4.4.

4.3.2 RPTSUM Outputs. RPTSUM has two output listing files.

4.3.2.1 Parameter File Listing (File RP). This is a listing of the parameters found on file PF, if any.

4.3.2.2 Report Listing (File GP). This listing has the printed reports requested, including the banner page, configuration overview and plots.

4.3.3 RPTSUM Deck Setup. The following control cards are required to execute RPTSUM:

```
$     IDENT       ACCOUNTING INFORMATION
$     USERID      USERID$PASSWORD
$     PROGRAM     RPTSUM,DUMP
$     PRMFL       **,R,R,B29IDPXO/NEWRMON/DDYN
$     TAPE9       IF      intermediate record input
$     DATA        PF      optional parameter file
$     SYSOUT      RP      parameter file listing
$     SYSOUT      GP      report listing
$     LIMITS      20,24K,,10K
```

The memory requirements depend on the number of reports and number of graphs requested for each report.

4.4 Parameter Language

The parameter language developed for PSUMR, DEMON and RPTSUM consists of free format statements. Each statement has a statement identifier, followed by one or more keyword phrases. A keyword phrase consists of subkeyword and/or user-supplied parameters.

In the following discussion, a set of closed parentheses is used to identify user-supplied parameters. The parentheses are not supplied by the user.

Closed brackets indicate optional parameters. Number signs surrounding a parameter imply that a parameter is required.

4.4.1 Statement Identifiers. The language contains five statement identifiers.

FORMAT	PURPOSE	PROCESSED BY
REPORT	specify report options	RPTSUM
CONFIG	specify standard configuration	PSUMR
SUMMARY	specify default include/exclude parameters	PSUMR DEMON
SNUMB	specify SNUMB to include in nonstandard categories	PSUMR
COMPARE	specify reports to compare	RPTSUM

4.4.2 Report Statement Keywords and Parameters. The keywords for the report statement are NAME, SYSTEM, PERIOD, TIME, INCLUDE and EXCLUDE.

4.4.2.1 NAME. This is used to supply an identifier for the report.

Format

NAME (report ID)

Comments

- maximum of six characters
- default = DEFALT

4.4.2.2 SYSTEM. This is used to specify the system ID to use for the report.

Format

SYSTEM (system ID)

Comments

- maximum of six characters
- corresponds to characters taken from .CRSID
- default = ID taken from the first record on file IF

4.4.2.3 PERIOD. This is used to specify the number of days the report will span. The PERIOD phrase has a secondary keyword followed by one or more optional parameters.

Format

```
PERIOD    DAILY    ([MMDDYY])
PERIOD    WEEKLY   ([MMDDYY])
PERIOD    MONTHLY  ([MMDDYY])
PERIOD    DAYS     (#NN#) ([MMDDYY])
PERIOD    MONTHS   (#MM#) ([MMDDYY])
PERIOD    DATES    (#MMDDYY#) ([MMDDYY])
```

Comments

- [MMDDYY] is the optional start date, except for "DATES" where it is the optional stop date
- NN, MM is the number of days, months respectively
- default is all data on the file
- daily, weekly, monthly imply 1, 7, 30 days respectively.

4.4.2.4 TIME. This is used to limit the time of day covered by the report. Times are given in system time, which may not be the same as local time.

Format

```
TIME   (HHMM) (HHMM)
```

Comments

- defaults are 0001 2359

4.4.2.5 INCLUDE. The INCLUDE phrase is used to specify which plots are to be printed and to specify the disposition of periodic/nonstandard data. The INCLUDE phrase has secondary keywords followed by parameters.

Format

```
INCLUDE   GRAPH   ([1,2,...8])
INCLUDE   DATA    (#P#)
INCLUDE   DATA    (#C#)
```

Comments

- correspondence between the graph numbers and graph names is

NUMBER	GRAPH NAME
1	PROCESSOR UTILIZATION
2	MEMORY UTILIZATION
3	I/O CHANNEL TIME
4	DISK SPACE SUMMARY
5	PERIPHERAL SUMMARY
6	SCHEDULED JOBS
7	a) EXECUTION/SWAPPED SUMMARY
	b) WAITING MEMORY/PERIPHERALS
8	REMOTE STATUS
Default is	1, 2, ... 8

- P implies that data marked as periodic will be included in the report. Refer to SUMMARY statement. The default is to exclude this data.

- C implies that data marked as nonstandard configuration will be included in the report. Refer to the CONFIG statement. The default is to exclude this data.

4.4.2.6 EXCLUDE. The EXCLUDE phrase specifies plots and data to be omitted from the default list. The EXCLUDE phrase has secondary keywords followed by parameters.

Format

```
EXCLUDE GRAPH   ([1,2,...8])
EXCLUDE DATA    (#P#)
EXCLUDE DATA    (#C#)
```

Comments

- any plot numbers supplied will be excluded from the output report

- P implies that data marked as periodic will be excluded from the report

- C implies that data marked as nonstandard configuration will be excluded from the report.

4.4.2.7 Sample Parameter File. The following are three sample report statements:

```
REPORT

      NAME WEEKLY
      SYSTEM NMCC2
      PERIOD WEEKLY  082382
      TIME   1230 2030
      INCLUDE GRAPH 1,2,8

REPORT

      NAME MONTH
      SYSTEM NMCC2
      PERIOD MONTHLY  080182
      INCLUDE DATA P

REPORT

      NAME HISTRY
```

Comments

A - The weekly report will cover 7 days' data beginning on 082382 between the hours of 1230-2030. Three plots will be printed: processor utilization, memory utilization, and remote status.

B - The monitor report will cover 30 days' data beginning on 080182 between 0001-2359. Periodic data will be included. All plots will be printed.

C - The history report will summarize data for the first system ID found on the input file. Periodic data will not be included.

4.4.3 CONFIG Statement Keywords and Parameters. The keywords for the CONFIG statement are PROCESSORS, MEMORY, TAPES, DISKS, PRINTERS and CARD READERS. The CONFIG statement is processed by PSUMR to determine when data is collected under standard/nonstandard conditions. If a CONFIG statement is present, the variables will be checked against the system configuration. If the system configuration is not the same as the CONFIG statement, the intermediate record will be flagged.

4.4.3.1 PROCESSOR. PROCESSOR specifies the standard number of processors.

Format

```
PROCESSORS  (#Y#)
PROC  (#Y#)
```

Comments

- y = number processors

4.4.3.2 MEMORY. Specifies the standard memory size in K.

Format

```
MEMORY  (#YYYY#)
```

Comments

- YYYY must be divisable by 64.

4.4.3.3 TAPES. TAPES specifies the standard number of physical channels which have a tape controller configured and the number of tape drives.

Format

```
TAPE   (#X/YY#)
TAPES  (#X/YY#)
```

Comments

- X = the number of physical IOM channels
- YY = the number of tape drives
- only one TAPE phrase per CONFIG statement is allowed. Thus, X, YY should be the totals for the system.

4.4.3.4 DISKS. Specifies the device type, number of physical channels, and number of devices for disks. There should be one DISK phase for each device class configured.

Format

```
DISK   (#DDD/X/YY#)
DISKS  (#DDD/X/YY#)
```

Comments

- DDD = the device type. Permissable device types are: 167, 170, 180, 181, 190, 191, 310, 400, 450, 500, 501
- X = the number of physical IOM channels for this device
- YY = the number of disk drives of this type.

4.4.3.5 PRINTERS. Specifies the standard number of printers configured.

Format

```
PRINTERS   (#X#)
PRINT      (#X#)
```

Comments

- X = the number of printers

4.4.3.6 Sample Parameter File. The following is a sample configuration statement:

```
CONFIG

      PROC 2, MEMORY 768

      TAPES  3/18
      DISKS  191/1/19
      DISKS  450/2/27
      PRINT  4
```

4.4.4 SUMMARY Statement Keywords and Parameters. The keywords for the SUMMARY statement are NOSUM, SUM. This statement is used to indicate what data is to be included/excluded from standard reports.

4.4.4.1 NOSUM. NOSUM specifies a sequence of intervals for which data will be included or omitted (marked "periodic").

Format

```
NOSUM  (#MMDDYY/I/O#)
```

Comments

- I = the number of days to include
- O = the number of days to exclude

Example

```
NOSUM  082382/5/2
```

This statement says "beginning with 23 August 1982, include 5 days' data, then omit 2 days." The initial include/omit intervals would be

INCLUDE	OMIT
8/23/82-8/27/82	8/28/82-8/29/82
8/30/82-9/3/82	9/4/82-9/5/82
9/6/82-9/10/82	9/11/82-9/12/82

REPORT IDENTIFICATION = DAILY

SYSTEM IDENTIFICATION = NMCC2

INCLUSIVE DATES = 01-10-83 01-10-83

INCLUSIVE TIMES = 00:01 23:59

GRAPH LIST

NUMBER	DESCRIPTION
1	PROCESSOR UTILIZATION
2	MEMORY UTILIZATION
3	I/O CHANNEL TIME
4	DISK SPACE SUMMARY
5	PERIPHERAL SUMMARY
6	SCHEDULED JOBS
7	EXECUTION/SWAPPED SUMMARY
8	WAITING MEMORY/PERIPHERALS
9	REMOTE STATUS

CONFIGURATION SUMMARY

PROCESSORS = 2

MEMORY = 0768K

TAPE CHANNELS = 03

TAPE DRIVES = 18

DISK CHANNELS = 03

DISK DRIVES = 45

PRINT CHANNELS = 06

PRINTERS = 06

DISK SUMMARY

DEVICE	NUMBER
191	27
450	18

Figure 4-3. Daily Resource Monitor Overview

DAILY 15 01 11-83 01.040 DAILY RESOURCE MONITOR PAGE 21

7 EXECUTION/SWAPPED SUMMARY

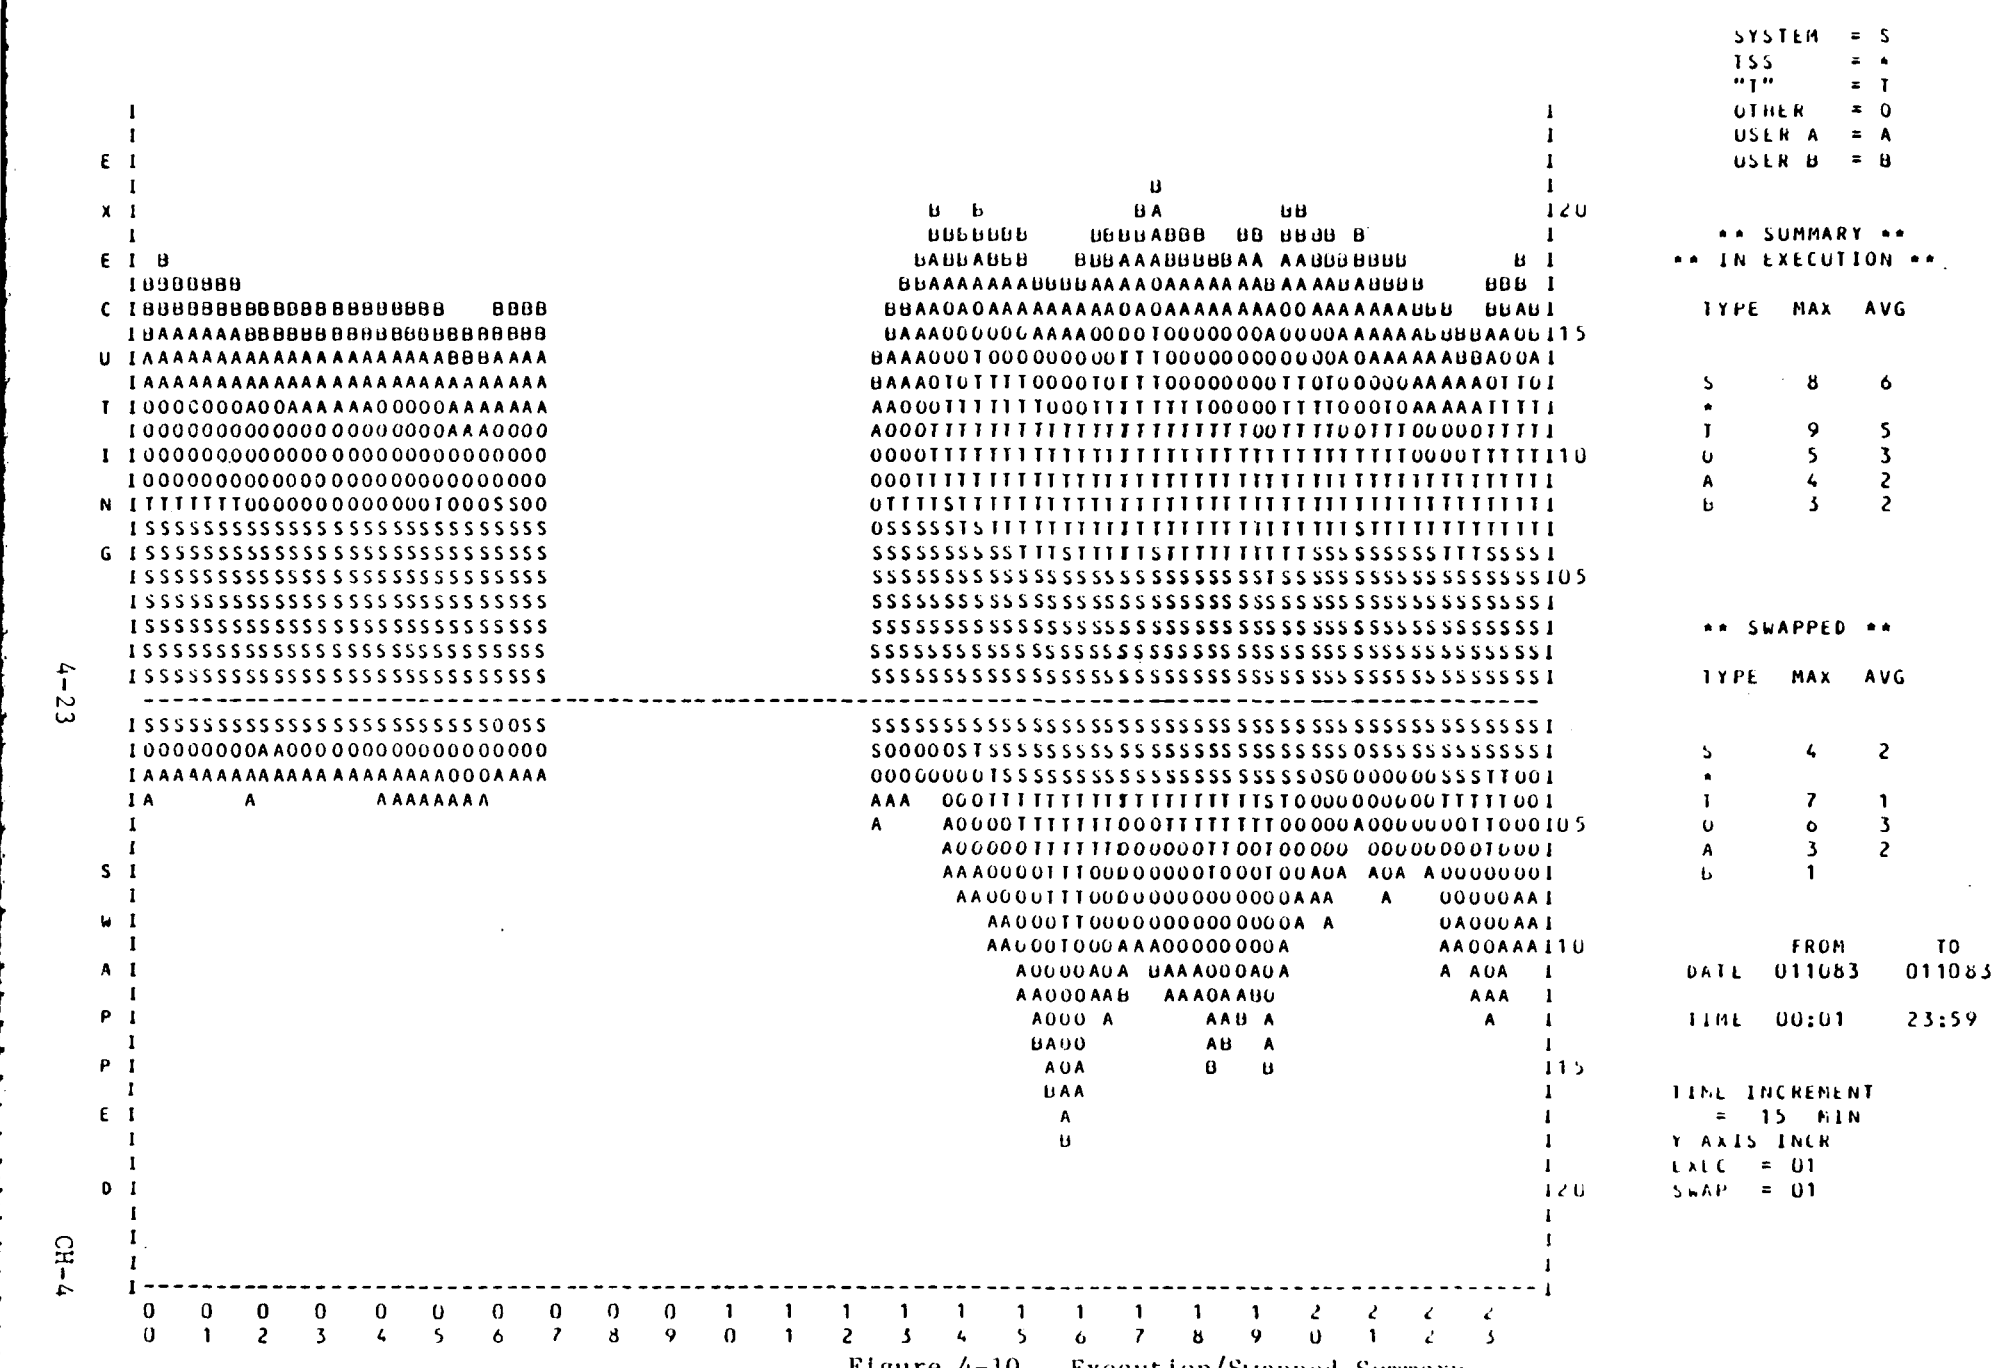
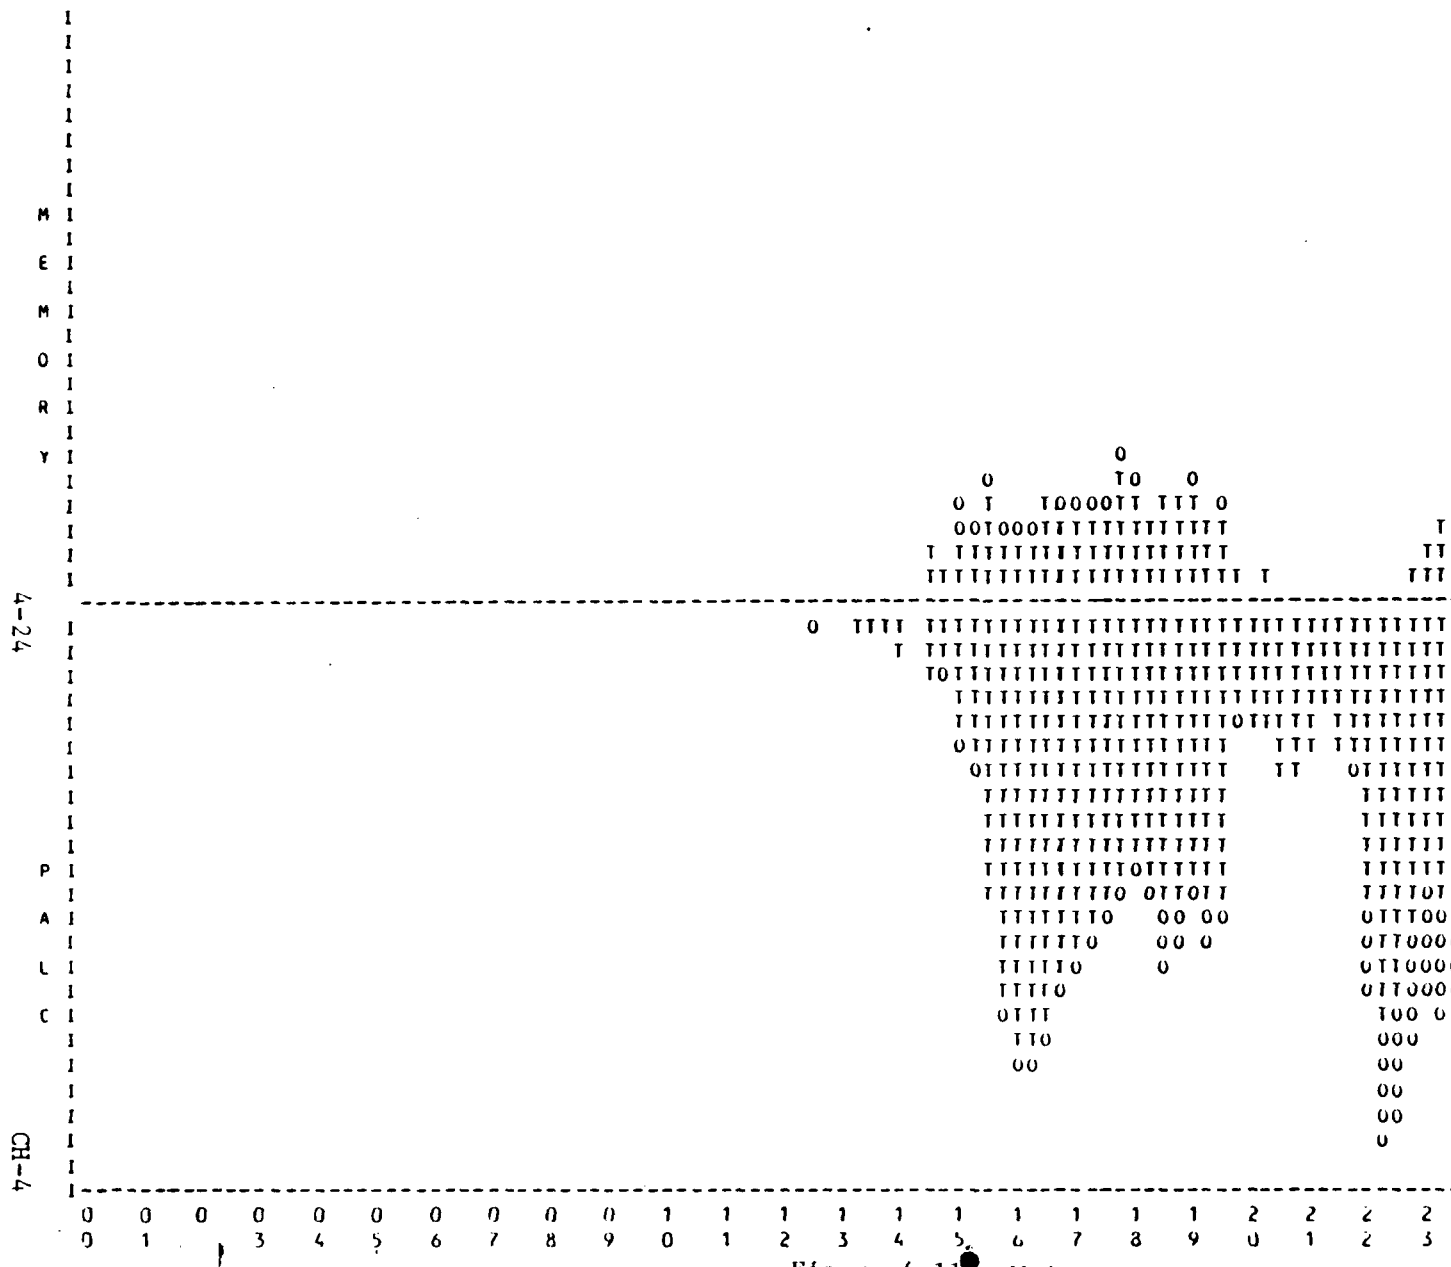


Figure 4-10. Execution/Swapped Summary

8 WAITING MEMORY/PERIPHERALS



USAGE CODE

SYSTEM = S
TSS = *
"I" = T
OTHER = O
USER A = A
USER B = B

** SUMMARY **

** WAITING MEMORY **

TYPE	MAX	AVG
S		
*		
T	5	2
O	2	
A		
B		

** WAITING PALC **

TYPE	MAX	AVG
S		
*		
T	18	10
O	5	2
A		
B		

	FROM	TO
DATE	011083	011083
TIME	00:01	23:59

TIME INCREMENT
= 15 MIN
Y AXIS INCR
PALC = 01
MEM = 01

Figure 4-11. Waiting Memory/Peripherals

9 REMOTE STATUS

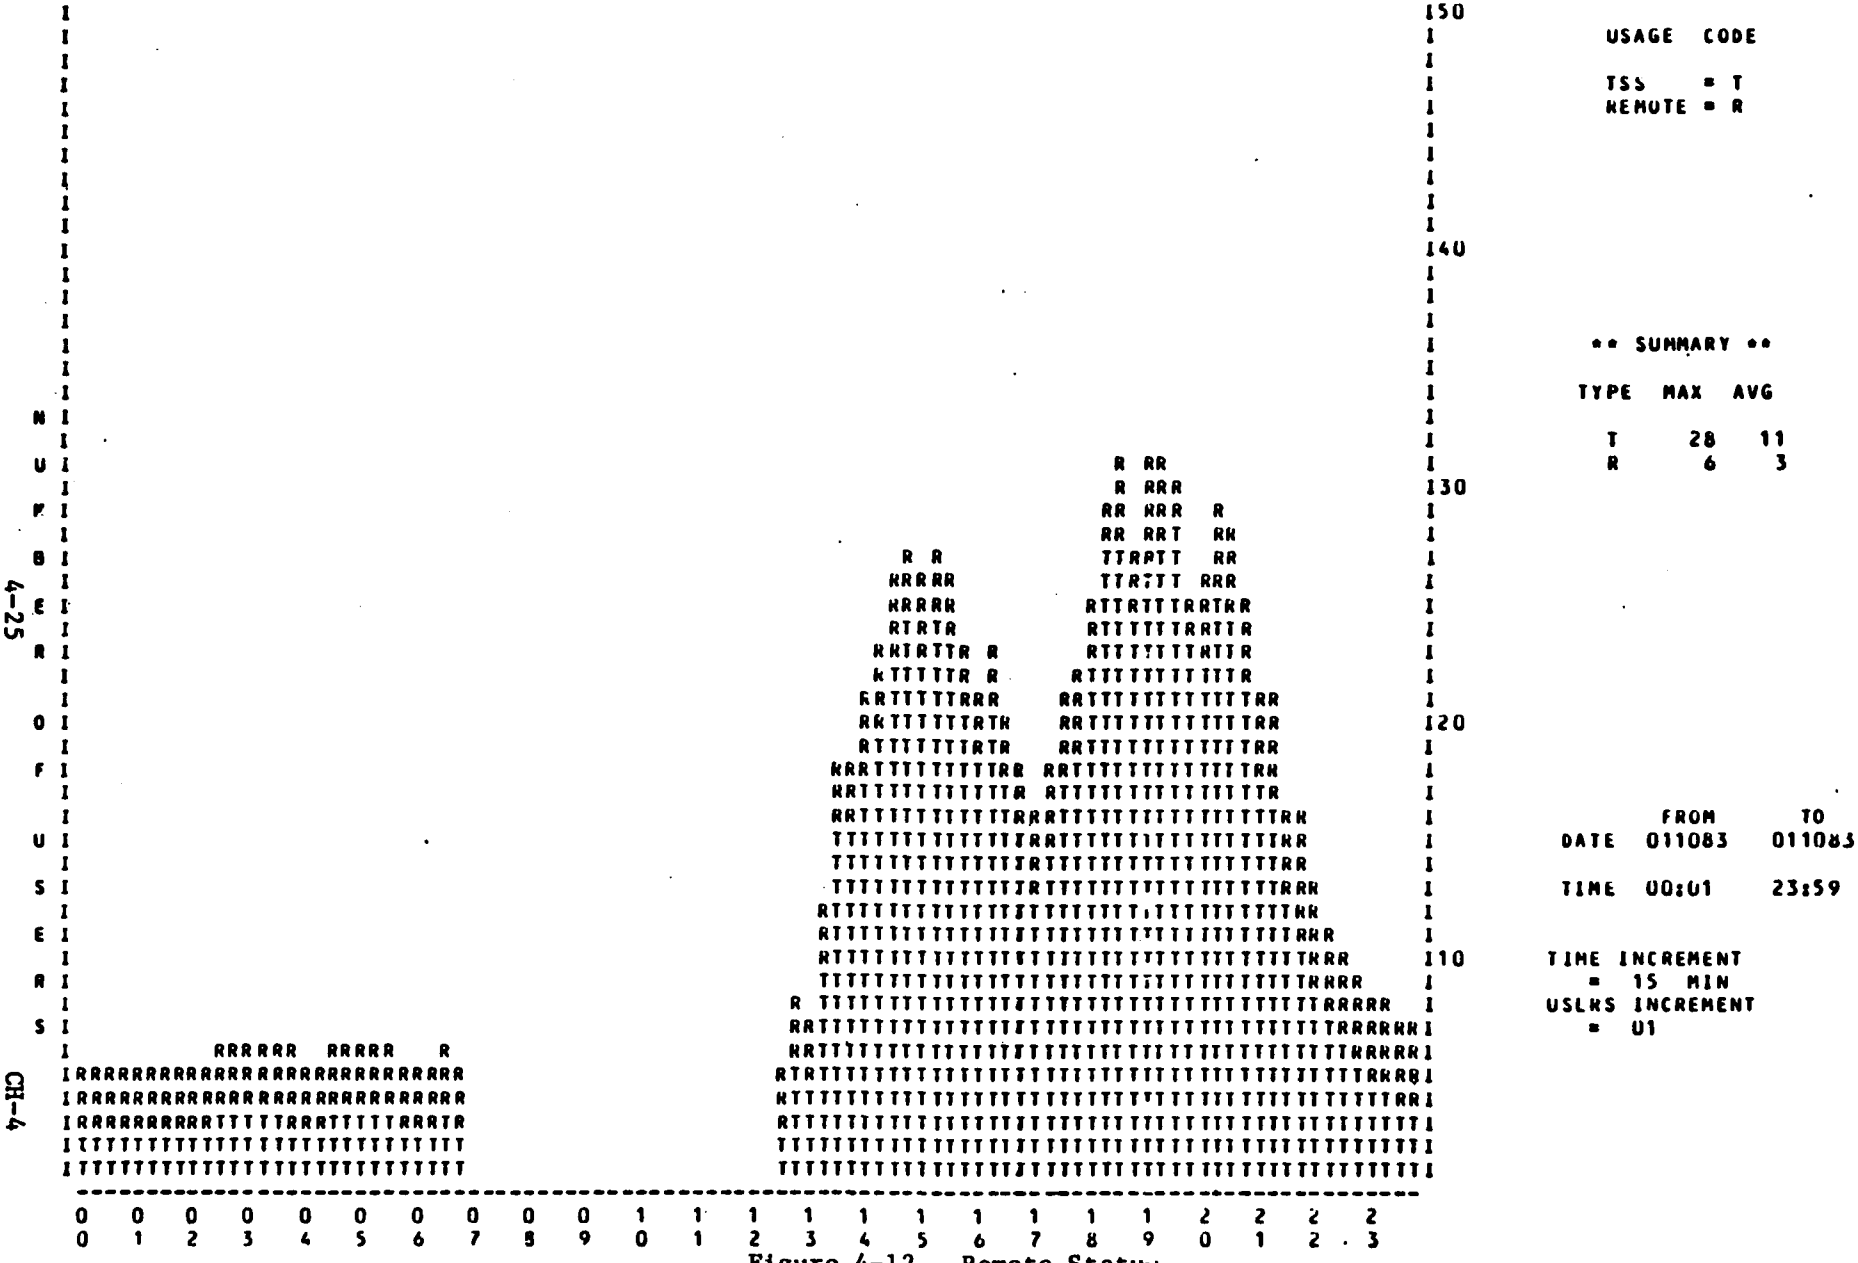


Figure 4-12. Remote Status

4.6 Usage Concepts

The WWMCCS resource monitor is designed to provide a means of tracking the utilization of major system components over extended time periods and to daily monitor the system usage automatically. By using the data supplied from the graphs, a PLOT of the system utilization can be developed. Such a plot would provide useful trending information. The daily monitoring capability provides a means of determining when utilization has varied from the established norm.

To use the resource monitor as designed, the following steps are required:

1. Turn on collection of SCF 608 records (refer to section 3).

2. Run the collector while the system is operational.

3. At the end of the system clock day (2400Z for WWMCCS sites), close the SCF file.

4. Run a PSUMR/DEMON job to create intermediate records and update the history file. The input is the day's SCFCLO tapes (or a daily tape made from the SCFCLO tapes). The intermediate records should be accumulated on a disk file if a weekly report or trending analysis will be done. The parameters of the history file should be set for prime time. This is accomplished by having the PSUMR/DEMON parameter file contain the following SUMMARY statement:

```
SUMMARY
      SUM     1300 2100
      NOSUM   011083/5/2
```

   Washington is ZULU +5 so prime time is 0800 + 5 = 1300. Hawaii is +10; hence, prime time begins at 1800. The NOSUM statement will exclude weekend data from the history file.

5. Daily Monitoring - Once the history file is initialzed, the graphs will not be generated unless the current day's data differs from the history file data by 25 percent for 25 percent of the data items. When the graphs are generated, it is an "automatic" signal that this day's utilization varied from the norm. The current status of the history file can be obtained (i.e., a set of graphs generated) by running DEMON and omitting file code for input.

6. Trending Analysis - In order to detect trends (increases or decreases) in utilization, the analyst must develop a plot of the values he wishes to "track." Until more experience is gained with the system, we recommend a weekly plot because a daily plot will probably have to much variance, while monthly might not be sufficient to detect trends. The weekly figures to plot are

obtained by generating desired graphs using RPTSUM and the saved intermediated record file. For example: suppose the site wishes to track TSS processor utilization and the number of TSS users during prime and non-prime time. This would be accomplished by running RPTSUM with the following parameters (assume +5 for time):

```
REPORT

     NAME PRIME
     SYSTEM NMCC
     TIME   1300 2100
     INCLUDE GRAPH 1,8

REPORT

     NAME NPRIME
     SYSTEM NMCC
     TIME   2101 1259
     INCLUDE GRAPH 1,8
```

These parameters will generate two processor utilization graphs and two remote status graphs: one for prime and one for non-prime time. Using graph paper, record the average (or maximum) TSS processor value and number of TSS users (see figure 4-13). If a site is currently experiencing a change in utilization, several weeks' data should show a trend (increase/decrease).

4.7 Abort Codes

Table 4-1 shows all abort codes associated with the Resource Monitor. The table lists the module, abort code and abort reason.

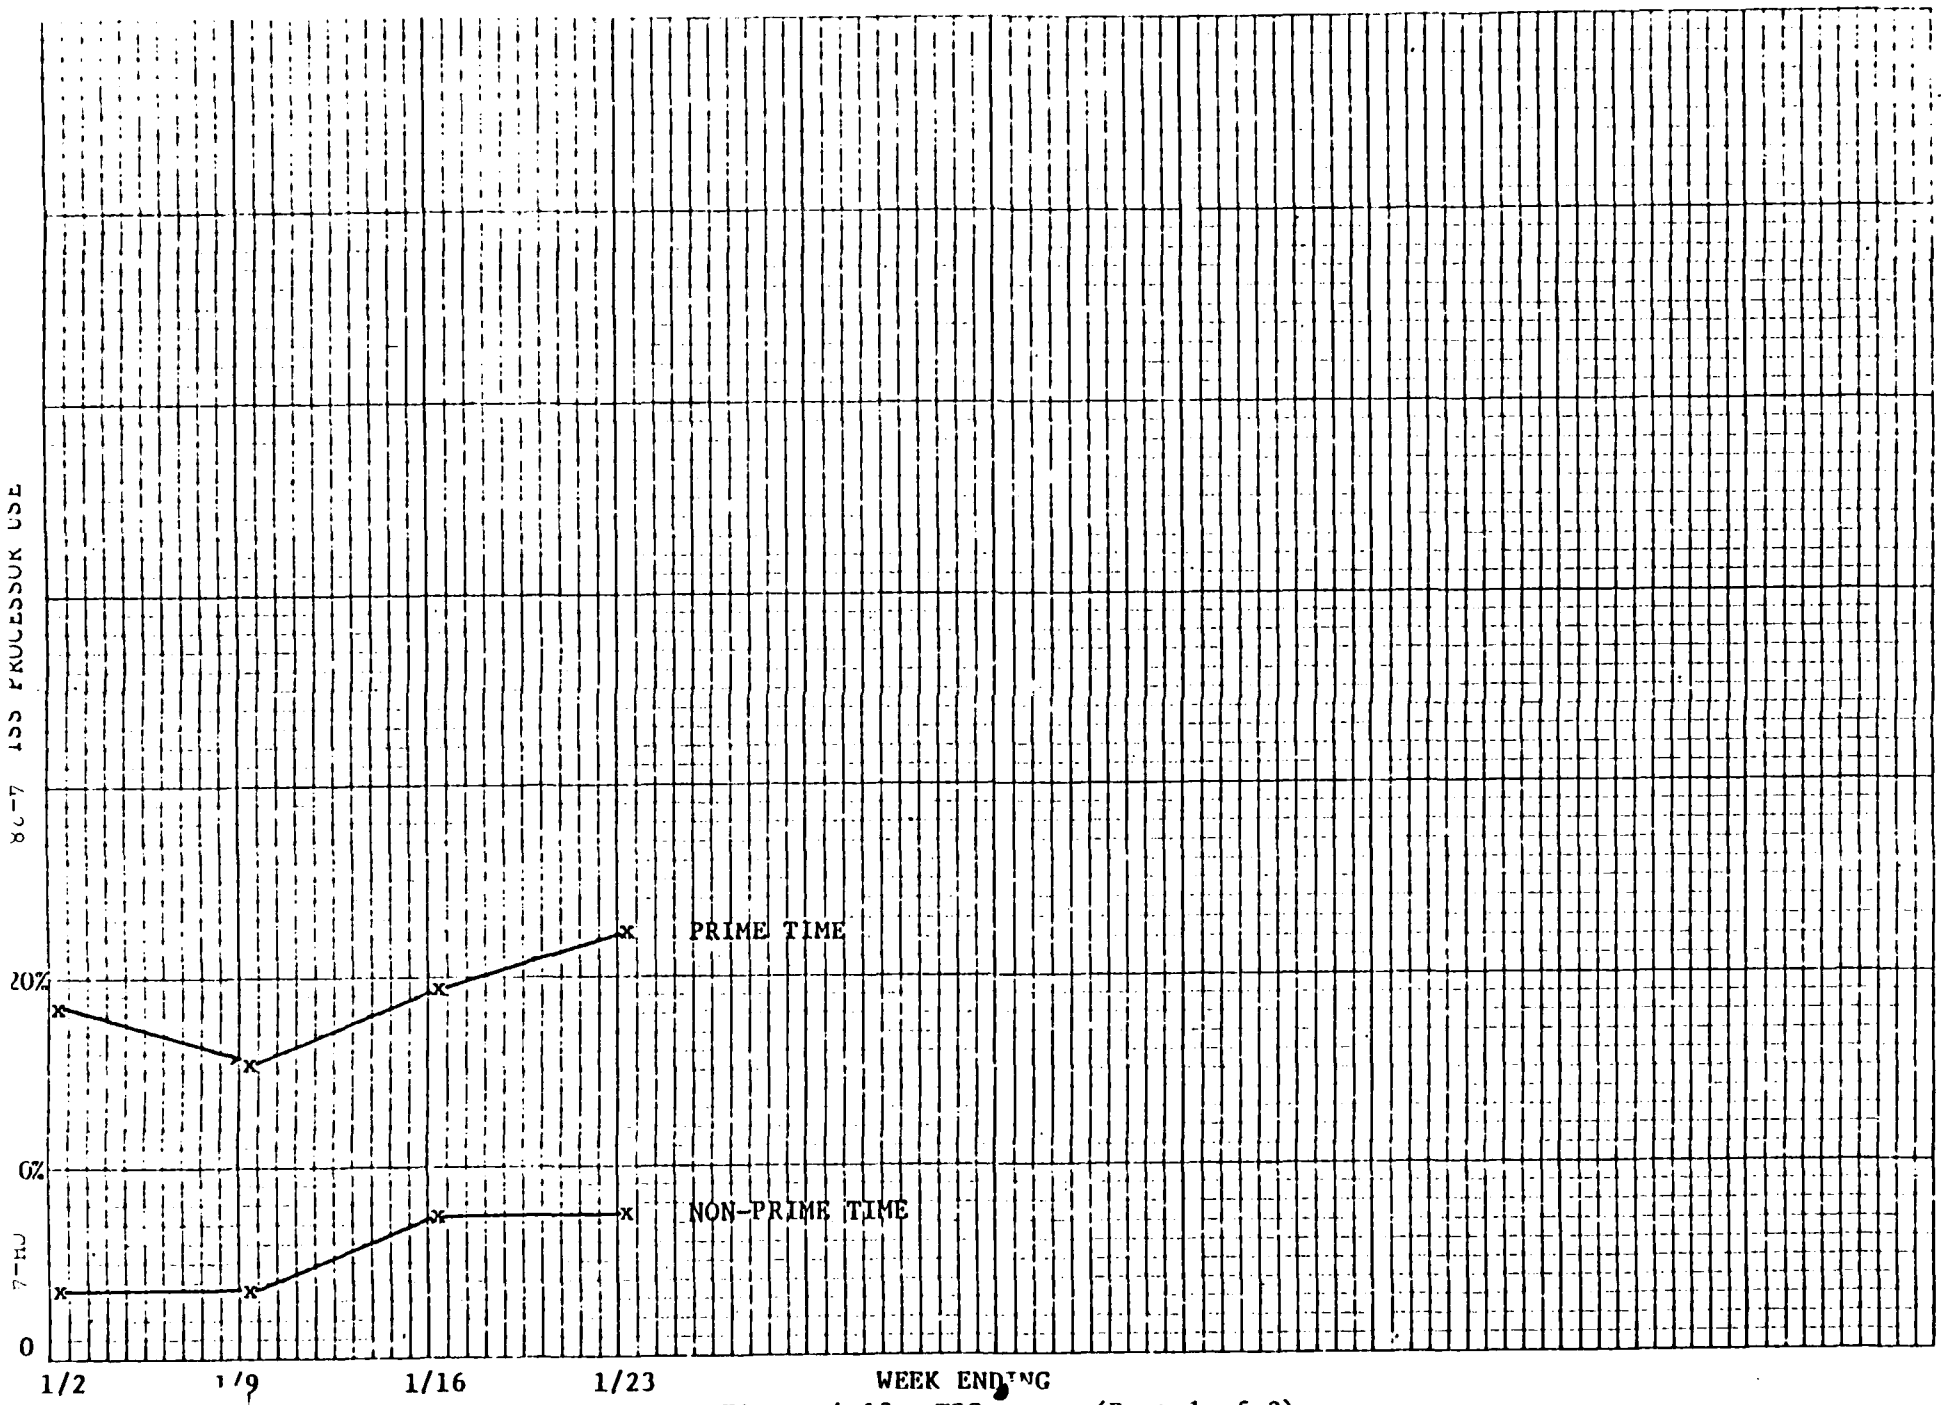


Figure 4-13. TSS [illegible] (Part 1 of 2)



Figure 4-13. (Part 2 of 2)

Table 4-1. Resource Monitor Abort Codes
(Part 1 of 3)

MODULE	CODE	REASON
DEMON	DM	Buffer denial for report buffers
	HF	BTAB processing error
	HB	Bad I/O status on history file
	HC	BTAB processing error. Ensure that the SUMMARY statement SUM parameters are the same for both PSUMR and DEMON.
	HD	Bad I/O status on history file
BLDMTX	BØ	Buffer denial processing error
	B2	Buffer size greater than 64 words is illegal
	B3	Disk file allocation denial from MTX10
MTX10	XØ	More sectors allocated than required
	X1	No disk space to allocate
	X2	No disk space large enough for request
	X3	Bad read I/O status
	X4	Bad write I/O status
	X5	DCW list overflow
	X6	Random file allocate denial
	X7	Read required write but no write possible -- processing error
	X8	No disk space to allocate
MTXADD	AØ	Current time less than start
	A1	Current time greater than stop
	A2	No graphs specified in report buffer -- processing error

Table 4-1. (Part 2 of 3)

MODULE	CODE	REASON
	A3	Required buffer not in core
PSUMR	PZ	I/O error on parameter file
	PØ	Buffer denial
	P1	Duplicate SCF record (sort should eliminate this)
	P2	Subtype 4 sequence error
FSUMR	F1	No subtype 1 record
	F2	No addr for subtype
	F3	Date check error
	F4	Subtype 4 records out of sequence
	F7	Negative time increment
	F8	Bad DF1G processing (lost data/system down flag)
	F9	Bad input data
	FA	Invalid input parameters
	FB	Interval processing error
	TS	Table sizes disagree
	F5	Too many disk types specified
	F6	Overflow
	FC	System ID change
GRAPH	G1	BINBCD conversion exceeded limits
	XM	Character move routine had bad character count

Table 4-1. (Part 3 of 3)

MODULE	CODE	REASON
SUPRTN	FA	Current time less than start
	FB	Current time greater than stop
	FC	No graphs specified
	BD	Core request denied
	OF	Overflow where not permitted
	01	Processing error in overflow correction
	02	Processing error in overflow correction
SBMR	BN	Release buffer not in valid range
	BM	Release same buffer
	BO	Pointer problem
CONIO	WT	I/O to console problem

SECTION 5. THE GENERAL MONITOR COLLECTOR - DATA COLLECTION PROGRAM

5.1 Introduction

The General Monitor Collector (GMC) data collection program is a privileged slave program which processes GCOS system trace data, organizes the data in GMC records, and writes the collected GMC data records on its output magnetic tape. The general concept of operation for the GMC facility is shown in figure 5-1.

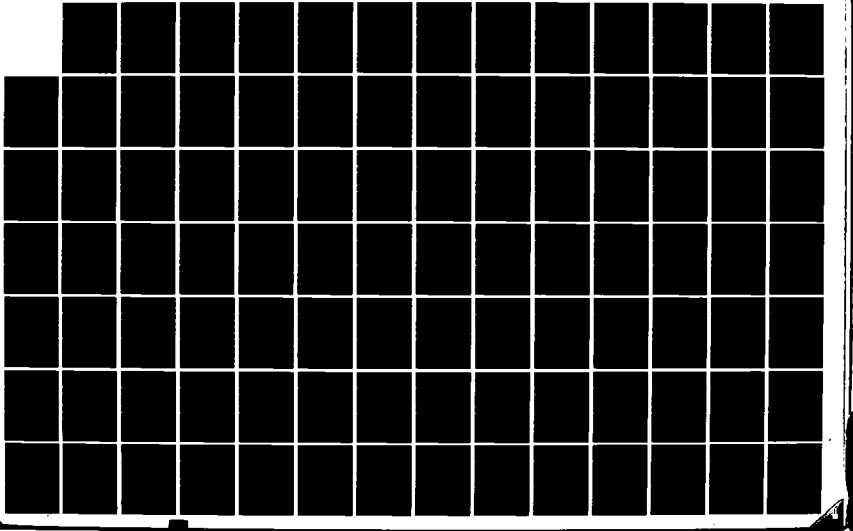
As a privileged slave program, the monitor data collector requires the permission of the system operator to run and must execute in master mode. The master mode capability allows the GMC to access all of the system main memory. The areas of interest are the system Communication Region (CR) and the individual job Slave Service Areas (SSAs).

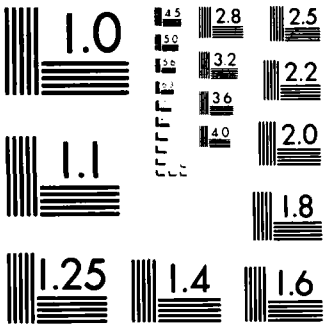
The GMC is actually a series of independent data collection monitors which are controlled via the central Executive Routine (ER) (described in subsection 5.3.1), and which use a common buffering routine for writing collected data to a common tape. The current GMC is comprised of 10 different monitors, which can be executed independently, or in any combination, except that TPE and TSS monitors cannot be run simultaneously. The monitors are described in this section. Each of the monitors has a dedicated data reduction program that produces formatted reports. These data reduction programs are described in sections 6 through 12 and 15.

The GMC can obtain control from the system in one of three ways. First, the standard manner is for the GMC to obtain control via the normal system trace. Second and third, GMC may also gain control in two nonstandard manners. These are via internally created system sub-traces (referred to as IT traces) and direct transfer patches (referred to as DT traces).

In the first way, the standard mechanism used by the monitor for obtaining control from the operating system is the normal GCOS system trace. The system trace capability records a history of the occurrence of as many as 72 system events, 64 of which are presently defined. This recording takes place in a circular table in the Communication Region of memory and is accomplished by execution of a unique code set resident in the system Dispatcher Module (.MDISP). Execution of this code is common to all system trace events and provides the point at which the GMC obtains control. The initialization portion of the GMC locates the system trace code set and implants a transfer instruction to the GMC executive. Thus, whenever a system trace event occurs anywhere within the system, the GMC executive will obtain control.

AD-A131 954 GENERALIZED MONITORING FACILITY USERS MANUAL CHANGE 6 2/8
(U) COMMAND AND CONTROL TECHNICAL CENTER WASHINGTON DC
B M WALLACK ET AL. 01 JUL 83 CCTC-CSM-UM-246-82-CHG-6
UNCLASSIFIED DCA/DF-83/001B F/G 9/2 NL





MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS 1963 A

There are three different options of system tracing permitted on the H6000. The extent to which tracing will occur is controlled via the $TRACE card. The format of this card, and its correct placement in the boot deck, is fully described in the GCOS System Startup Manual. Option 1 is to completely inhibit all system tracing. This would be accomplished with a $TRACE card containing the following pattern:

```
777777777777,777777777777
```

If this $TRACE card is present in the boot deck, the .MDISP module will eliminate all trace logic from the system during initialization. Because of this removal of trace logic, it is physically impossible to activate tracing at this point. If the GMF is started in this environment, it will abort with an NF abort.

Option 2 is to activate all traces. This would be accomplished with a $TRACE card containing the following pattern:

```
0,0
```

This option will result in significant CPU overhead being generated by the operating system. However, if this option is selected, the GMF can be run at any time with no changes required in the boot deck.

Option 3 is to activate only a subset of all traces. During initialization, the .MDISP module will "NOP" out all calls to the trace module for those traces not requested on the $TRACE card. Using this option, a site can minimize the CPU overhead generated by the trace logic, and at the same time, ensure that GMF can be executed without requiring that the boot deck be altered. If this option is in effect, GMF will determine which traces it requires. It will then proceed to turn on all required traces and, if necessary, remove any NOP instructions and replace them with the proper call to the trace module. When GMF terminates, the system will be restored to the way it was prior to the execution of GMF. GMF can force activate all of its required traces, except for trace types 46, 50, 51 and 52. Trace types 46 and 51 are required by the Memory Monitor, trace type 51 is also required by the CPU Monitor, and trace types 50, 51 and 52 are required by the Tape Monitor. In order to ensure that GMF can be activated at any time, and at the same time minimizing trace logic overhead, the following is the suggested minimum trace card for use by any site running GMF:

```
$       TRACE       777777777777,643777777777
```

THIS PAGE LEFT INTENTIONALLY BLANK

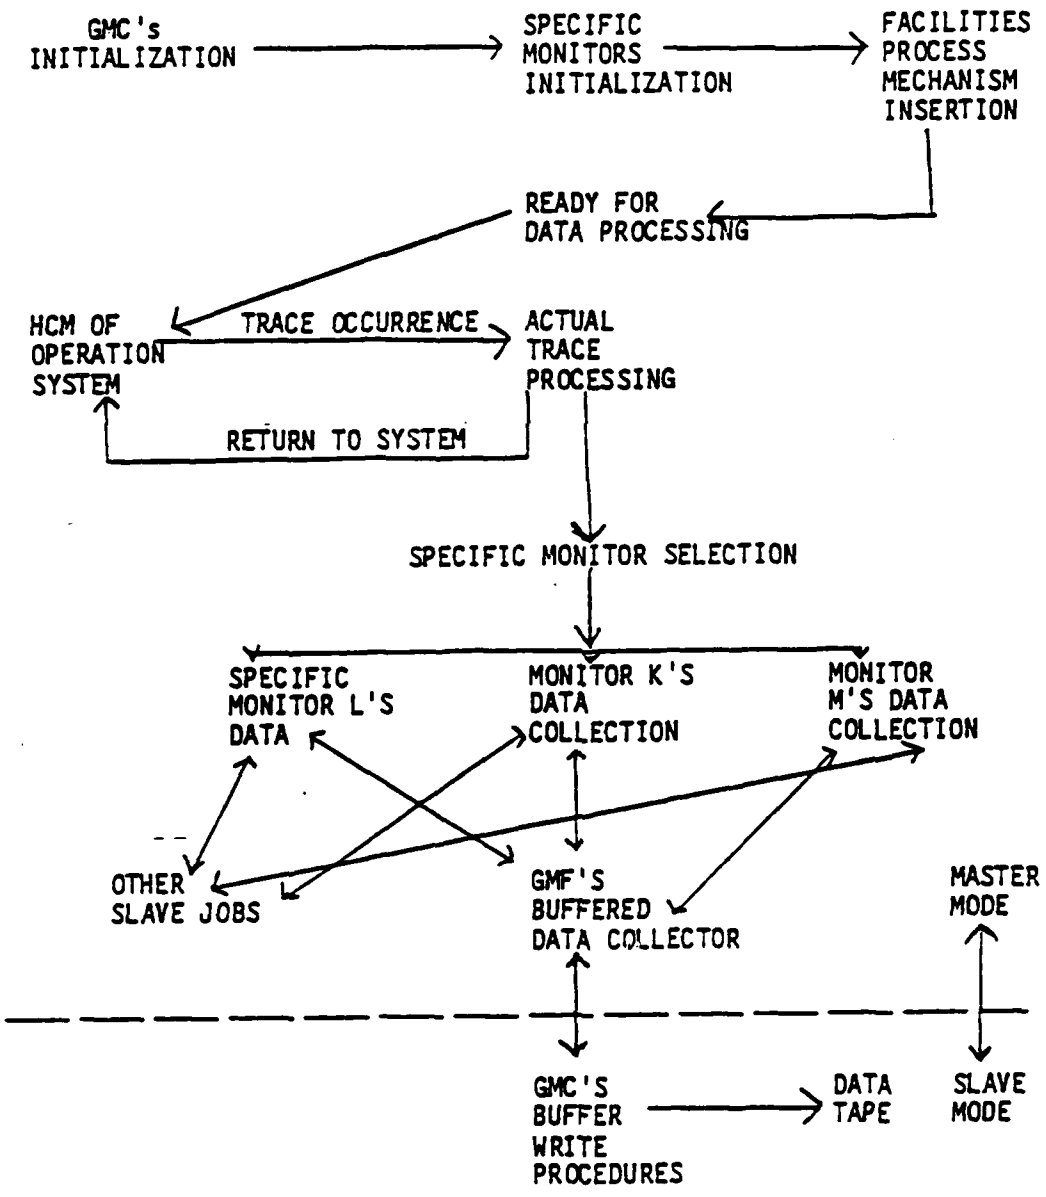


Figure 5-1. GMC Concept

After obtaining control, the system trace recording is completed using normal system procedures and then the trace is investigated. If an executing GMF monitor requires this trace type, control is given to that monitor. The monitor then collects data of interest to itself and requests the GMC executive to buffer the data. When the monitor completes its task, it returns control to the GMC executive. The GMC executive then transfers control back to the system trace processing routine within the GCOS dispatcher. The activities mentioned above take place in master mode under the guise of an extension of normal system trace procedures.

The second method used by GMC to gain control is for GMC to create its own system traces. The GMC will search a given GCOS module for a known line of code. It will replace this line of code with a transfer to a patch area. In the patch area, the monitor will insert code to create a new GMC system trace. At this point, the execution of this code will be processed just as all other system traces are processed. This procedure is used when the Communication Analysis Monitor, TSS Monitor and GRT Monitor are selected for execution. (See subsection 5.2.6 for a complete description of this procedure.)

The third method used by GMC to gain control is via a direct transfer from a GCOS module. In this case, GMC will search a given GCOS module for a known line of code. It will replace this line of code with a direct transfer into the GMC. This can be accomplished since GMC is locked in core during its entire execution. The benefit of this procedure is that the overhead of the system trace is eliminated. This procedure is used only when the Mass Store Monitor or the CPU Monitor is selected for execution. (See subsections 5.2.2 and 5.2.3 for a complete description of this procedure.)

As the GMC ER buffers the collected monitor data, it will determine when one of the internal buffers are filled and must be written to tape. At this point, it will establish a normal slave dispatch to its tape writing facility. The tape writing facility will then write the internal buffer to tape and signal the executive that this buffer may be reused.

5.2 GMC Monitor Subroutines

In this subsection, each of the ten monitor subsystems will be addressed. See table 5-1 for a quick reference of required trace types for each monitor and refer to table 5-2 for all GMC abort codes.

5.2.1 Memory Utilization Monitor. The Memory Utilization Monitor (MUM) is used to measure memory utilization. For a detailed description of reports containing data collected by this monitor, see section 6.

When MUM is active it is essential that GCOS system trace types (octal) 46 and 51 are enabled in the computer system boot deck $ TRACE card. MUM collects data upon the occurrence of traces 10, 11, 46 and 51 and builds its records which are then passed to the Executive Routine (ER) for buffering. A separate discussion of the format of the MUM collected data records is contained in subsection 5.4.2. MUM requires that at least the following segment numbers from table 5-3 be used to generate the GMC R* file:1, 2, 11, 19, 23, 26 and 27. The complete process for generating an R* file is described in subsection 5.6. It should be noted that this version of MUM will not collect any idle trace information and therefore not all MUM reports will be produced during data reduction. See section 6 for description of reports that will not be produced.

If the user wants all the MUM reports produced, then the Idle Monitor must be made active, along with the Memory Monitor. To do this, GCOS trace types (octal) 0, 1, 2, 3, 10, 11, 13, 16, 21, 22, 37, 46 and 51 will be enabled by GMF. In generating the R* file the following segment numbers should be used: 1, 2, 10, 11, 19, 23, 24, 25, 26 and 27. It should be noted that when the Idle Monitor is active the amount of data collected will be about double that collected when the Idle Monitor is off. The user should carefully evaluate the need for those reports produced by the Idle Monitor.

The MUM is designed so that it can accurately report all changes to the memory subsystem. It does this by processing all trace type 10's (.CALL) and all trace type 11's (.GO TO). Upon receiving these traces a further check is made to see whether a memory allocation process is being requested i.e. the use of SSA modules: .MPOP3,.MPOQ4, or MPOR5. The MUM will collect data only if one of these modules has been requested.

The first item of information reported by the MUM is the current status of all jobs waiting in the Peripheral Allocator's Queue. This information is reported so as to be better able to represent the true core demand being made by the current workload. When a GCOS system has a large number of jobs waiting core, a Core Damper Switch is set. This switch is used to prevent jobs from being sent from the Peripheral Allocator to the Core Allocator. Therefore, the Peripheral Allocator's queue may contain many jobs that would normally be in the Core Allocator's queue, were it not for the Core Damper Switch. This

Table 5-1. Required Trace Type for Each Monitor

Monitor #	Monitor	Required Trace Type (Octal)
0	Memory Utilization Monitor (MUM)	10, 11, 46, 51, (Idle Monitor traces optional)
1	Mass Storage Monitor (MSM)	7, 15, 73*, 76*
2	CPU Monitor (CPUM)	23*,51,63*,70*
3	Tape Monitor (TM)	50, 51, 52
4	Channel Monitor (CM)	4, 7, 15, 22 (Idle Monitor traces optional)
5	Communications Analysis Monitor (CAM)	14*, 15
6	GRTS Monitor (GRTM)	62*
7	Transaction Proc-cessing System Monitor (TPSM)	0, 1, 2, 4, 5, 6, 13, 42, 51, 65, 74*
8	Idle Monitor (IDLEM)	0, 1, 2, 3, 13, 16, 21, 22, 37
A	TSS Monitor (TSSM)	74*

*These are not standard traces. They are specially created by the GMC or by an editing of the GCOS TPE Subsystem in the case of trace type 74. Trace types 23, 63, 70, 73 and 76 are direct transfers into the GMC and as such are not required to be active via the $ TRACE card in the system boot deck. Trace types 14, 62 and 74 do use the System Trace Function.

Table 5-2. Abort Codes (Part 1 of 4)

B2 - Illegal SNUMB on MSM data card (more than 5 characters).

B3 - More than 5 SNUMBs for MSM SNUMB option.

BC - Illegal request on limited connect option.

BK - Backspace of the full data tape was bad. Multireel will not be collected. Check for tape drive problems.

BS - Bad tape status. Check condition of tape and rerun job.

C1 - CPU Monitor turned off but SNUMB input requested on the data card.

C2 - Illegal SNUMB (more then five characters) on data card for CPU SNUMB option.

C3 - More then five SNUMBs for CPU Monitor on data card.

CD - Illegal character in CAM special option.

CE - Console message garbled. Check console sheet and check with operator.

CM - Cannot move out of the growth range of TSS.

CO - CAM turned off but special option requested.

DK - No multireel disk file was present. Use a $ FILE card in the JCL or use the M9 option to turn off multireel capability.

DR - Disk read-in. End-of-reel processing was bad.

DS - Bad status of the disk write.

ER - Wrapup record could not be written.

ET - More than two terminals requested for CAM special option.

FN - The IOS accounting modification could not be found. Call CCTC.

GC - No GRTS control card.

GD - No FEP I/O can be performed.

GM - Needed memory for GRTS Monitor denied. Increase $LIMIT card.

GO - GRTS Monitor illegal data card.

GS - Extra SSA is not available for GRTS Monitor. Check $ LIMIT card.

Table 5-2. (Part 2 of 4)

M0-M8,MA - Monitor was not turned off and not present on the R* file. Any monitor not contained on the R* file must be turned off via the data card option. The number following the M is the monitor that was not turned off.

MM - The dispatcher hook has already been inserted. Another version of GMC must already be in execution.

N1 - The CPU Monitor hook code could not be found. See subsection 5.2.3 and table 2-1, element 11.

N2 - Sufficient patch space is not available in .MDISP to run the CPU Monitor. See subsection 5.2.3 and table 2-1, element 16.

N3 - DNWW/MDNET patch location could not be found. See subsection 5.2.6 and table 2-1, element 18.

N4 - Sufficient patch space is not available in DNWW/MDNET to run the Communications Analysis Monitor. See subsection 5.2.6 and table 2-1, element 20.

N5 - MSM patch for SSA cache analysis not found (AOS .CRTDL). See subsection 5.2.2. and table 2-1, elements 22 and 23.

N6 - MSM patch for SSA cache analysis not found (AOS .CRTBH). See subsection 5.2.2 and table 2-1, elements 22 and 23.

N7 - MSM patch space in .MDISP not sufficient for monitoring SSA cache. See subsection 5.2.2. and table 2-1, element 24.

N8 - CPU Monitor hook code for subdispatch could not be found. See subsection 5.2.3 and table 2-1, element 12.

NA - CPU Monitor hook code for dispatcher queuing could not be found. See subsection 5.2.3 and table 2-1, element 13.

NB - CPU Monitor hook code for removing jobs from dispatcher queue could not be found. See subsection 5.2.3 and table 2-1, element 14.

NF - The Dispatcher hook code could not be found. See subsection 5.1 and table 2-1, element 2.

NL - Number of LLINKS specified on DK card 0 or greater than 16000.

NS,NT,NU - Problems with Communication Monitor initialization. See table 2-1, element 9.

OE - An error in an off option was encountered. Check the data cards. There is either an illegal character on the data card or a monitor which was not compiled in the R* file (see assembly listing) has not been turned off.

Table 5-2. (Part 3 of 4)

OK - All went correctly.

OL - Overlap of disk file. Increase size of disk file. Check if operator failed to respond to tape mount message during multiprocessing.

OV - A tally overflow occurred in the MUM.T10 subroutine. Increase the size of the data area within subroutine MUM.T10, variable SIZBUF. If this action is required, the user must insure that the variable CRSIZE is changed via a global edit in all data reduction routines. The new value of CRSIZE should equal the new value of SIZBUF +50.

RS - Routine depth requested exceeded table length.

RW - Error in initial rewind. Check tape and drive.

SB - End-of-reel processing was bad. Check tape and drive.

SD - Error setting of density.

SF - Limited reel option termed successfully.

SQ - Sequence error in the physical records.

S1 - Subroutine MUM.T10 missing

S2 - Subroutine MUM.T46 missing

S3 - Subroutine CM.T07A missing

S4 - Subroutine CPU.T70 missing

S5 - Subroutine CM.T04A missing

S6 - Subroutine CM.T22A missing

S7 - Subroutine TM.T50 missing

S8 - Subroutine CAM.T14 missing

S9 - Subroutine GRT.T62 missing

SA - Subroutine IDL.TRCS missing

SC - Subroutine IDL.T21 missing

SD - Subroutine TPE200 missing

SG - Subroutine TSS.COL missing

Table 5-2. (Part 4 of 4)

TE - The start/stop times appear improper. Check data card.

TL - Trailer record write was bad. Check tape and drive.

TS - An OK abort directed by a time to stop command.

TW - The tally word has been garbled.

T0 - Trace type 46 or 51 not active on $TRACE card.

T2 - Trace type 51 not active on $TRACE card.

T3 - Trace type 50, 51 or 52 not active on $TRACE card.

T7 - TPE Monitor required traces not active on $TRACE card.

Z0 - Cannot find trace 10 or 11 logic. See table 2-1, element 4.

Z1 - Cannot find trace 7 logic. See table 2-1, element 6.

Z2 - Cannot find trace 22 logic. See table 2-1, element 8.

Z3 - Cannot find trace 4 logic. See table 2-1, element 7.

Z4 - Cannot find trace 15 logic. See table 2-1, element 5.

Z6-Z9, ZA, ZB - Cannot find trace logic required for Idle Monitor. See table 2-1, element 10.

Table 5-3. GMC Catalog and File Index (Part 1 of 4)

SEGMENT #	FILE	REQUIRED	FUNCTION
1	GMF.TOP	Y	Read data card, initialization, find hook in dispatcher, and create initial record
2	MUM.INIT		Initialize Memory Monitor
3	MSM.INIT		Initialize Mass Store Monitor
4	CPU.INIT		Initialize CPU Monitor
5	CAM.INIT		Initialize Communications Analysis Monitor
6	CM.INIT		Initialize Channel Monitor
7	TM.INIT		Initialize Tape Monitor
8	GRT.INIT		Initialize DN-355 Monitor
9	TP.INIT		Initialize TPE Monitor
10	IDL.INIT		Initialize Idle Monitor
10A	TSS.INIT		Initialize Timesharing Monitor
11	GMF.MID	Y	Ensure at least one active monitor
12	CAM.PAT		Preparation for patching DNWW/MDNET for Communications Analysis Monitor
13	CPU.PAT		Preparation for patching dispatcher for CPU Monitor
14	MSM.PAT		Preparation for patching dispatcher for MSM Cache Analysis
15	PATLOOK		Searches for patch space for CPUM, CAM, MSM
16	CPUDOIT		Patch the dispatcher for CPU Monitor
17	CAMDOIT		Patch DNWW for Communications Analysis Monitor
18	MSMDOIT		Patch dispatcher for MSM Cache Analysis

Table 5-3. (Part 2 of 4)

SEGMENT #	FILE	REQUIRED	FUNCTION
19	GMF.MON	Y	Insert the trace hook for GMC traces
20	CPU.REMO		Remove CPU Patches to dispatcher
21	CAM.REMO		Remove CAM patches to DNWW/MDNET
22	MSM.REMO		Remove MSM patches to dispatcher
23	GMF.BTM	Y	Remove the trace hook, do all IO control
24	IDL.TRCS		Processes traces (octal) 0,1,2,3,13,16,22,37 for Idle Monitor
25	IDL.T21		Processes trace (octal) 21 for Idle Monitor
26	MUM.T10		Processes traces (octal) 10,11,51 for Memory Monitor
27	MUM.T46		Processes trace (octal) 46 for Memory Monitor
28	CPU.T70		Processes traces (octal) 23,51,63,70 for CPU Monitor**
29	TM.T50		Processes traces (octal) 50,51,52 for Tape Monitor
30	CAM.T14		Processes traces (octal) 14 and 15 for CAM*
31	CM.TO4A		Processes trace (octal) 4 for Channel Monitor
32	CM.T22A		Processes trace (octal) 22 for Channel Monitor
33	CM.TO7A		Processes traces (octal) 7,15,73,76 for Channel Monitor and Mass Store Monitor**

Table 5-3. (Part 3 of 4)

SEGMENT #	FILE	REQUIRED	FUNCTION
34	GRT.T62		Processes trace (octal) 62 for GRTS Monitor*
35	GRT.COL		Interfaces with the DN-355
36	TPE200		Processes traces (octal) 0,1,2,4,5,6,13,42,51,65 and 74 for TPE Monitor*
36A	TSS.COL		Captures trace (octal) 74 for TSS Monitor*
37	RUN.GMF		JCL to run a GMC R *
38	GMF.OBJ		File to contain a GMC R *
	The following set of files are a series of JCL files under the catalog B29IDPXO/GMFCOL/MAKE used to create different GMF R* monitors. The numbers in the name are the corresponding monitor number (see subsection 5.5.1):		
39A	MO2		Memory and CPU Monitors
39B	ALL		Total GMC
39C	MUM		Memory Monitor
39D	CPU		CPU Monitor
39E	TM		Tape Monitor
39F	MSM		Mass Store Monitor
39G	MO258		Memory, CPU, Communications and Idle Monitors
39H	M148		Mass Store, Channel and Idle Monitors
39I	CAM		Communications Analysis Monitor
39J	CM		Channel Monitor
39K	GRT		DATANET-355 Monitor

Table 5-3. (Part 4 of 4)

SEGMENT #	FILE	REQUIRED	FUNCTION
39L	M48		Channel and Idle Monitors
39M	M14		Mass Store and Channel Monitors
39N	M56		Communications and DATANET Monitors
39O	TPE		TPE Monitor
39P	TSS		TSS Monitor
39Q	MO25		Memory, CPU and Communications Monitors
39R	MO1245		Memory, Mass Store, CPU, Channel and Communications Monitors

*Trace types 14,62 and 74, are not standard. They are internally generated (IT) traces.

**Trace types 23,63,70,73 and 76 are not standard. They are direct transfer (DT) traces.

information from the Peripheral Allocator is reported only when the Peripheral Allocator is in memory and a Memory Monitor trace is about to be generated. For this reason, not all Peripheral Allocator queue changes will be reported. In order to reduce the amount of information being collected, a job's status in the Peripheral Allocator's queue is reported only for new jobs, when a job has changed activity, or when its status has changed.

After reporting any Peripheral Allocator status information, the MUM will next report the status of every job waiting for or currently using memory. Information such as the SNUMB, IDENT, USERID, Activity Number, memory demands, current memory address, whether the job is in memory or waiting for memory, and whether the job is a system program or user program is collected. This information is reported for each job only if a change has occurred from previous information that was reported. In addition, the current amount of CPU and IO time used by a job is reported in every MUM trace that is generated.

The MUM will consider a job to be a system job if it has a program number less then octal 10, or if it has no J* file and requires privity. Since the user may want to consider other jobs to be system jobs, such as HEALS or VIDEO, the data reduction program allows the user to extend this definition of system jobs (see section 6).

5.2.2 Mass Storage Monitor. The Mass Storage Monitor (MSM) is used to collect data on usage of peripheral resources. For a detailed description of reports containing data collected by this monitor, see section 7.

MSM processes trace types 7, 15, 73, and 76 to build its own records which are passed to ER. A separate discussion of the format of the MSM collected data records is contained in subsection 5.4.3. As has been stated earlier, trace types 73 and 76 are direct transfer traces created by the GMC, and as such need not be defined on the $ TRACE card. The MSM requires that at least the following segment numbers from table 5-3 be used to generate the GMC R* file: 1, 3, 11, 14, 15, 18, 19, 22, 23, and 33. The complete process for generating an R* file is described in subsection 5.6. If the system being monitored by the Mass Store Monitor contains SSA Cache Core, two new direct transfer traces, are created by the Mass Store Monitor in order to collect sufficient data to be able to analyze the operation of SSA Cache Core. These traces are created only if SSA Cache Core is configured. The Mass Store Monitor searches the dispatcher for a AOS .CRTDL instruction and then inserts code to make a direct transfer into the GMF. In addition, an AOS .CRTBH instruction is also searched for so that another direct transfer into the GMC can be generated. The first instruction locates the area of the dispatcher where it has been determined that the required SSA module is not in the SSA Cache Buffer and needs to be loaded from disk. The second instruction

locates that area of the dispatcher where it has been determined that the required SSA module is in the SSA Cache Buffer. Table 2-1, elements 21-23, completely describe the procedure for locating these instructions. If GMC cannot find these instructions between these locations, it will abort with an N5 or an N6 abort. If this problem occurs, the dispatcher code should be examined to see if this instruction has been moved or modified. If so, the code in GMC will need to be altered.

Upon finding the above sets of instructions, GMC searches the dispatcher area for 8 free locations in which to create two new direct transfer traces. Table 2-1, element 24, completely describes the procedure for locating the patch area. If 8 words of free space are not found, an N7 abort will occur. In this case, the user should examine the patch deck and a listing of the patches on the total edit tape to see if a large number of patches have been made to the dispatcher. If this is the case, the dispatcher code will need to be reassembled in order to remove these patches or else the Monitor will not be able to be utilized. The user does have another alternative. This alternative involves patching word 0 of the dispatcher in order to generate a user patch area. The patch involves the setting of bit 2 to a 1 in word 0 of the dispatcher. No other modification by the user is necessary.

The MSM collects sufficient information so as to be able to completely monitor the usage of the entire disk subsystem, the usage of SSA Cache core and the usage of FMS catalog cache, when active. When either the MSM or CM is active, a record containing device names and addresses is written at the beginning of the GMC run and periodically afterward if device names change. This is done only for mass store devices. Every time the system issues a connect request to a tape drive or disk drive, sufficient information is collected so as to be able to identify who is issuing the connect, the file being connected to, the pack upon which the file is located, the parameter types for the file being connected to and the reason for the connect, i.e read, write, write verify, etc.

Whenever a MME is issued the MSM will check whether it is a system job issuing a MME GEFSYE. For purposes of this check a system job is considered to be any job with a program number less then octal 15. If the Peripheral Allocator is issuing the GEFSYE, information is collected as to the SNUMB the Peripheral Allocator is working for, and the file code that is being GEFSYE'd. If the GEFSYE type is a 2, 3, 4, 5, 8, 9, 10, 11, 18, 19, 20, 22, 23, 27, 28, or a 29 then additional information is collected so as to be able to determine the CAT/FILE string of the file being GEFSYE'd. This information will be used by the data reduction program to correlate file codes used by jobs to the actual CAT/FILE string being referenced by a job. Also, sufficient information is collected so as to be able to report how

many FILSYS connects are required in order for the system to be able to allocate and deallocate files requested by a job.

If FMS catalog cache is active or available space tables are being buffered in memory then the MSM will generate a record type octal 77 with sufficient data as to be able to monitor the effect of FMS catalog cache and available space table buffering. This record is generated once, for every 5000 connects issued by the system. This is not a physical trace that is being generated and, as such, does not need to be present on the $ TRACE card. Rather, it is merely a data record that is being written to tape and given the unique number of octal 77. The data record consists of a dump of some internal performance parameters maintained by the GCOS system within modules .MFSIO and .MASO4.

5.2.3 CPU Monitor. The CPU Monitor (CPUM) is used to collect data on CPU utilization. This monitor can be used only on a WW7.2 or commercial 4JS system. For a detailed description of reports containing data collected by this monitor, see section 11. When the user desires that the CPUM be active, GCOS trace type (octal) 51 must be enabled in the computer system boot deck on the $ TRACE card. CPUM processes trace types 51, 23, 63 and 70 to build its records which are passed to ER. A separate discussion of the format of the CPUM collected data records is contained in subsection 5.4.4. Trace types 23, 63 and 70 are direct transfer traces, created by the GMC, and as such, need not be defined on the $ TRACE card. The CPUM requires that at least the following segment numbers from table 5-3 be used to generate the GMC R* file: 1, 4, 11, 13, 15, 16, 19, 20, 23, and 28. The complete process for generating an R* file is described in subsection 5.6. The CPU Monitor searches the dispatcher for an ASA .SALT,5 instruction and then inserts code to generate a direct transfer trace into GMC. In order to capture subdispatch processor time, it also searches for a STQ .QTOD,4 instruction and then inserts code to make a direct transfer into GMC. In the T70 capture routine, the time increment will be negative for a regular dispatch and positive for a subdispatch. In order to monitor dispatcher queuing, the CPU Monitor will search the dispatcher for an LDA .STATE,4 instruction and then insert code to generate a direct transfer trace into GMC. Finally, in order to monitor the removal of jobs from the dispatcher queue, the monitor will search the dispatcher for an ANQ .STATE,5 instruction and then insert code to generate a direct transfer trace into GMC. Table 2-1, elements 11-14, completely describe the procedure for locating these groups of instructions.

If GMC cannot find the ASA .SALT,5 instruction, it will abort with an N1 abort; if it cannot find the STQ instruction it will abort with an N8 abort; if it cannot find the LDA instruction, it will abort with an NA abort; and if it cannot find the ANQ instruction, it will abort with an NB abort. If these aborts occur, the dispatcher code should be examined to determine if the instruction has been modified, moved, or patched. If so, the code in GMC will need to be modified.

Upon finding these instructions, GMC searches the dispatcher patch area(s) for either 8 or 16 free locations in which to create a direct transfer trace into the GMC. Table 2-1, elements 15 and 16, completely describe the procedure for locating the patch area. If patch space is not found, an N2 abort will occur. See subsection 5.2.2 for a description of an alternate procedure in case the search for patch space fails.

The CPU Monitor tracks the CPU usage of all system programs and accumulates CPU usage of slave jobs into a single value (see subsection 5.4.4). If the user desires, he can specify up to five slave jobs for which he wants the CPU monitor to track CPU usage, just as it does for system jobs. Subsection 5.5.5. describes this user option.

The CPU Monitor can be operated in one of two modes. Under the standard default mode, the monitor will capture data for reporting on CPU utilization, as well as CPU dispatcher queuing. The user should refer to subsection 5.4.4 for a description of the data collected and section 11 for a description of the reports produced by this monitor. Under this mode of operation, the monitor will use approximately 3-4% of the available processor power. The user has the option of disabling the CPU dispatcher queuing portion of the monitor. Under this mode of operation, the monitor will only require about 1% of the available processor power, but the user will receive no reports concerning dispatcher queuing, lengths of queues or amount of time in queue. See subsection 5.5.5 for a description of this user option.

5.2.4 Tape Monitor. The Tape Monitor (TM) is used to collect utilization statistics on magnetic tape drive activity. A separate discussion of the format of the TM collected data records is contained in subsection 5.4.5. Reports containing data collected by this monitor are described in section 11.

When the user desires that the TM be active, GCOS trace types (octal) 50, 51, and 52 should be enabled in the computer system boot deck on the $ TRACE card. TM processes these trace types to build its records which are passed to the ER. The TM requires that at least the following segment numbers from table 5-3 be used to generate the GMC R* file: 1, 7, 11, 19, 23, and 29. The complete process for generating an R* file is described in subsection 5.6.

5.2.5 Channel Monitor. The Channel Monitor (CM) is used to measure I/O channel activity over tape and disk channels and contention to disk devices. A separate discussion of the format of the CM collected data records is contained in subsection 5.4.6. See section 8 for a description of reports containing data collected by this monitor.

CM processes these trace types 4, 7, 15 and 22 to build its records, which are passed to the ER. The CM requires that at least the

following segment numbers from table 5-3 be used to generate the GMC R* file : 1, 6, 11, 19, 23, 31, 32, and 33. The complete process for generating an R* file is described in subsection 5.6.

Actually, when the CM is active, sufficient data is processed for obtaining reports not only from the Channel Monitor but also from the Mass Store Monitor. The only Mass Store Monitor data that cannot be collected would be the data needed to analyze Cache Memory. If the user also wants this data to be collected, he should create an R* file from the following segments (see table 5-3): 1, 3, 6, 11, 14, 15, 18, 19, 22, 23, 31, 32, and 33. In addition, the Mass Store Monitor must be made active. There is an additional option available with the Channel Monitor. This option allows the Channel Monitor Data Reduction Program to produce a CPU Idle/IO Active Report. This report is described in section 8. To obtain this report, the Idle Monitor must be included in the R* file. The following segments are required to generate the R* file: 1, 6, 10, 11, 19, 23, 24, 25, 31, 32, and 33.

5.2.6 Communications Analysis Monitor. The Communications Analysis Monitor (CAM) is used to measure machine and user response time and terminal usage. A separate discussion of the format of the CAM collected data records is contained in subsection 5.4.7. The complete process for generating an R* file is described in subsection 5.6. The output reports, containing data collected by CAM, are described in section 9. CAM patches the DNWW (MDNET in W7.2) module, looking for a LDQ M.LID,3 instruction followed by an ANQ =0077777,DU instruction. When the monitor is terminated, CAM removes these patches from the system. The CAM requires that at least the following segment numbers from table 5-3 be used to generate the GMC R* file: 1, 5, 11, 12, 15, 17, 19, 21, 23, and 30.

If CAM cannot find these instructions, GMC will abort with an N3 abort. See table 2-1, element 18, for procedure used to locate these instructions. If this problem occurs, the DNWW/MDNET code should be examined to see if this instruction has been moved or modified. If so, the code in CAM.PAT will need to be altered.

Upon finding these instructions, CAM then searches DNWW/MDNET patch area for 10 free locations in which to create a new system trace type 14. Table 2-1, element 20, completely describes the procedure for locating this patch area. If no space is found by this search, an N4

abort will occur. In this case, the user should examine the patch deck to see if a large number of patches have been made to DNWW/MDNET. If this is the case, DNWW/MDNET will need to be re-edited in order to remove these patches or else the CAM will not be able to be utilized. In addition to the above patching, CAM.INIT also searches DNWW/MDNET for certain instructions. See table 2-1, element 9, for a complete description of this search. If a specific terminal's traffic is to be monitored (see subsection 5.5.3), the CAM will insure that no more than two terminal IDs have been included. Invalid terminal IDs, extra terminal IDs or terminal IDs without the CAM input option request will cause the GMC to abort with a CD, CO, or ET abort. See table 5-2 for the meaning of these aborts.

The CAM also uses the GCOS trace type 15 (octal) to check for any JDAC processing or any other line switching which may occur.

5.2.7 GRTS Monitor. The purpose of the GRTS Monitor (GRTM) is to collect statistical data to be used in evaluating the performance of the DATANET 355 Front-End Processor (FEP). This data includes CPU Utilization, Response Time, and Terminal Performance. The collected information is sent to the GMC within the H6000, which writes the data to tape. A separate discussion of the format of the GRTM collected data records is contained in subsection 5.4.8. This tape, containing GRTM performance data and possibly data from other monitors, is used as input to a data reduction program used to produce statistical reports (see section 10). This monitor can only be used on a WWMCCS system. It cannot be used on a commercial system.

5.2.7.1 Configuration Requirements. The GRTM will execute on H6000 system with up to eight FEPs. The monitor is designed to run on the GRTS II Phase IIA (GRTS 1.2) FEP software.

5.2.7.2 H6000 Configuration Requirements. The GRTM requires that at least the following segment numbers from table 5-3 be used to generate the GMC R* file: 1, 8, 11, 19, 23, 34, and 35. The complete process for generating an R* file is described in subsection 5.6.

5.2.7.3 Altering of Phase II-A Software. To use the GRTM, the user must modify the standard GRTS software. It should be noted that in Release WW7.2 the alter cards needed to modify the standard GRTS software are included within the standard release. The user must check the FMAC module to insure that variable FMON has been set to 1. The setting of this variable will ensure that the required alters are assembled into the GRTS system. The FMAC module must be recompiled and the macro library reloaded. Finally, all the GRTS modules should be recompiled.

5.2.7.4 FEP Configuration Requirements. The modified GRTS II software will produce performance data. These modifications are then enabled in the software through the use of the following control card during FEP system initialization:

```
CC1

1       8       16
$       GOPT    RCSMON
```

When the monitor is not running, the FEP will function normally. Execution of the GRTM is initiated by the H6000 as it connects to each FEP to be monitored. At that time, instructional parameters are sent to each FEP to be used in determining the amount of buffer space needed for the collection of the statistical data.

The GRTM software when configured will require an additional 1,000 (decimal) words of DATANET main memory to execute. This core requirement can grow to as much as 2,500 decimal words depending upon the input options selected. The main portion of the monitor code will be resident within the GRTS II FSUB module with additional patches being incorporated within the FCCP, FEXC, FCIP, FINT, FICM, and FHCB modules. It should be noted that when the monitor is not assembled, the 1-2K of core required for the monitor will be available for buffer space. This should be kept in mind when reviewing the output reports.

5.2.7.5 Interface Requirements. During its initialization phase, the GRTM software will attempt to log onto the H6000 system via a pseudo-terminal that is used in its interaction with the host-resident GMC program. Due to .MSECR requirements, the pseudoterminal attempting to gain access to the system must be on a .MSECR configured SYSTEM HSLA subchannel. The Physical Terminal Address (PTA) used by the pseudo terminal is unique and therefore cannot be configured in the GRTS II configuration deck for any other purpose.

The format of this PTA word is as follows:

BITS	MEANING
0-3	IOM CHANNEL NUMBER OF HSLA
4-5	HSLA NUMBER (1,2,3)
6-10	HSLA SUBCHANNEL NUMBER (0-31)
11-15	POLLED SCREEN NUMBER
16-17	MUST BE ZERO

As an example, a PTA value of 317200 would be used to identify the pseudo terminal as being on:

```
IOM CHANNEL NUMBER OF HSLA = 6
HSLA NUMBER                = 1
HSLA SUBCHANNEL            = 29
POLLED SCREEN NUMBER       = 0
MUST BE ZERO
```

This is the same format used to describe subchannels within .MSECR.

A user must insure the setting of the PTA value as circumstances demand. By using the format shown above, the symbolic location PTA located in the FSUB module may be altered to reflect the user's own PTA value. The FSUB module must then be reassembled prior to bootloading the DATANET.

5.2.7.6 Abort Codes. The general abort codes listed in table 5-2 apply to specific options. Abort codes created specifically for the GRTM are listed below.

GS = SSA processing error caused by missing or invalid $LIMITS card

GC = Invalid GRTM options on data card file (I*)

GC = Missing GRTM option card

GD = No FEP I/O can be performed

GM = GEMORE unsuccessful in getting needed buffers

NOTE: GM aborts occur if another program occupies continuous memory just above that of the GRTM when buffer space is requested using MME GEMORE. Increasing the $LIMITS memory value or loading the GRTM immediately after booting the system will enhance the chances of getting DATANET buffers.

5.2.7.7 DATANET Monitor Software Description. The GRTS II system operating within the DATANET will be used in the collection and recording of various internal resource information which is then sent

to the GMC executing in the host (6000) system. Monitoring functions within the DATANET have been separated into three basic monitoring categories:

a. CPU and Resource Monitoring

b. Host/FEP Response Time

c. HSLA Subchannel Monitoring

Monitor information will be passed between the DATANET and the GMC program executing in the host using the normal FICM interface.

5.2.7.7.1 DATANET-HOST Interface. The GRTM executing in the DATANET will be in the form of a pseudoterminal connected to the DATANET. Once initialized, the pseudomonitor terminal will be as any other remote terminal connected to a Direct Access (DAC) program in the host.

As part of GRTM initialization, the pseudo terminal will issue a connect-to-slave request to the DAC GMC program in the host. The DAC connect name requested will be the special monitor name of GRTMN'X' with the 'X' being a number from zero through seven which corresponds to the DATANET (0-7) that issues the connect request.

The connect-to-slave request will remain outstanding until it is answered by an inquiry issued by the GMC program. Once the request has been answered, the GRTM/GMC program connection will be made.

Since the pseudo monitor terminal will not be physically connected to the DATANET (i.e., on an HSLA S/C) the need for a special monitor "Terminal Control Block (TCB)" becomes necessary.

The special monitor TCB is necessary in order to utilize the normal FICM interface in the passing of monitor information to the GMC program.

5.2.7.7.2 Monitoring of CPU. The DATANET Monitor will collect CPU and various resource information by the placement of conditionally assembled instructions at appropriate points in the ICM and EXC modules of the GRTS II software. The information to be collected at the CPU level is independent of the HSLA subchannel and is sent to the host based upon a predefined time sampling increment for writing buffers to the host collector program. For a given buffer sent to the host, the following information is included:

a. Time Idle - The amount of time in milliseconds spent in the exec idle loop since the last buffer was sent to the host.

b. Buffer Denials - A cumulative count of buffer denials. This is a count of the number of times that the GRTS II software was unable to get a buffer for a user.

c. Buffer Availability - The number of 18-bit words currently available for buffers.

d. Number of Users - Count of the number of users currently logged on the system.

e. Number of Transactions - Two counts are maintained and sent to the host:

(1) The accumulated number of transactions sent to the host, and

(2) The accumulated number of transactions received from the host.

f. Size of Transactions - There are two counts maintained and sent to the host:

(1) A cumulative count of the number of 36-bit words sent to the host, and

(2) A cumulative count of the number 36-bit words received from the host.

g. Host RSVPs Count - A cumulative count of the number of host RSVPs received by the DATANET.

h. Buffer Requests - A cumulative count of the number of times the buffer allocation routine was called.

5.2.7.7.3 Host/DATANET Response Time. The GRTS II Monitoring of the FEP/Host Response Time will be measured on a program name basis. For this monitoring, conditional coding will be added to the FICM module to detect the various requests to the Host. Each time that either a "Connect-to-Slave" or "Disconnect" request is detected, the following formatted entry will be made into the response time buffer area:

a. Function Code

b. Type of Device/Line ID

c. Time Stamp

d. DAC Name (for connect-to-slave requests only)

The Function Code (i.e., connect-to-slave, disconnect) will identify the type of request with the line number entry identifying the logical line number of the device.

The GRTM will capture the following requests:

a. Accept Direct Input (355 asking H6CCO to accept data)

b. Input Accepted (input received by the H6CCO)

c. Send Output (355 requesting continuation of output)

d. Output Received (355 has received data from H6CCO)

e. Output Started (355 has begun transmitting data to terminal)

f. Output Complete (355 has completed transmission of data to terminal)

g. Accept Direct Output (H6CCO has told 355 it has data to send)

h. Accept Direct Output, Then Input (H6CCO has told the 355 it has data to send and expects input)

Each time one of the above requests or responses is detected, a response time buffer entry of the following format is made:

a. Function Code

b. Type of Device/Line ID

c. Time Stamp

In placing both the line number and the time stamp in every collector buffer entry, response times between the various request cycles of each DAC program executing the host will be effectively monitored.

5.2.7.7.4 Terminal Monitoring. GRTS II Terminal Monitoring will be on a HSLA subchannel (S/C) basis. Every monitored HSLA S/C will be allocated a four word (18 bit) record area within the output buffer where the various monitor information for the S/C is accumulated. Each word within the S/Cs record will be a "predefined" location where the various counts for the S/C are held. The first word of each S/C record will contain information as to the HSLA number and the S/C number to which the record belongs.

The GRTM will update these various counts dynamically as they occur within the DATANET.

5.2.8 Idle Monitor. The Idle Monitor (IDLEM) is used to collect data concerning CPU activity. This monitor can only be used in conjunction with the MUM or the CM. It should not be activated if one of the two aforementioned monitors is not active. If the Idle Monitor is present on the R* file and active and if, in addition, the MUM or CM is not active, then the IDLEM will automatically be turned off. The user should read subsections 5.2.1 and 5.2.5 for information concerning the use of the IDLEM. A separate discussion of the format of the collected data records is contained in subsection 5.4.9.

This monitor generates an excessive number of traces and significant overhead. It should be activated only after the user ensures that the reports produced because of its presence are an absolute necessity for the evaluation. In most cases, this monitor should not be required.

5.2.9 Transaction Processing System Monitor. The GMC Transaction Processing System Monitor (TPSM) is used to collect data on the performance of the Navy Transaction Processing Executive (NTPE) System. A separate discussion of the format of the TPSM collected data records is contained in subsection 5.4.10. The reports containing data collected by TPSM are described in section 12.

When TPSM is active, the required traces must be enabled in the computer system boot deck on the $ TRACE card (see table 5-1). This is the only GMF monitor for which the minimum $TRACE card format described in subsection 5.1 is not sufficient for proper execution. A sample of the reports and run time procedures for the data reduction program can be found in the Transaction Processing System (section 12). The TPSM requires that at least the following segment numbers from table 5-3 be used to generate the GMC R* file: 1, 9, 11, 19, 23, and 36. The complete process for generating an R* file is described in subsection 5.6.

NOTE: The TPSM cannot be run concurrently with the TSSM and should only be used on a WWMCCS WW7.2 system. It should not be used on a commercial release.

5.2.9.1 TPS Trace Collection. The TPSM is unlike most other GMC monitors in that monitoring of the Transaction Processing System is controlled via the operator console. Prior to collecting data, the user must alter the NTPE (see subsection 5.2.9.2) and must also create a usable GMC R* file (see subsection 5.6). Once these actions are performed and a GMC execution is started, the user must still perform one additional action before data collection can begin. The TPSM is turned on or off by the console operator via the TP MESS command. The operator must request "TP MESS". When the console responds "TP MESS?", the message "TRACE ON" is entered to start trace generation, or "TRACE OFF" to suspend trace generation. This procedure can be repeated as often as desired. The TPSM and the TSSM are the only GMC monitors that can be turned on or off while the GMC is physically executing.

5.2.9.2 Modifying the Transaction Processing System. To use the TPSM, the user must alter the Navy Transaction Processing System. NTPE is delivered as a set of alters to the commercial version of TPE (these alters are in "*C" form for use by the SCED utility program). Embedded in these alters are the modifications to place the GMC trace points in the executive modules (see figure 12-21). All sections of code for the GMC support are under control of conditional assembly parameters. To enable the GMC code, the following steps must be performed immediately prior to STEP 15 of the NTPE installation procedures. This procedure is in chapter 3 of the Navy Transaction Processing Executive (NTPE) Site Manual (N7506-SM April 1981). From TSS, enter:

```
*CONV TRAX-EXEC/Wxxx/STAR-C:V
FILE STAR-C CONTAINED 12984 CARD RECORDS ...
*EDIT
-RVS:/GMF    SET     0/;*:/GMF    SET     1/

GMF     SET     1                                          ***GMF**
GMF     SET     1                                          ***GMF**
GMF     SET     1                                          ***GMF**
GMF     SET     1                                          ***GMF**
GMF     SET     1                                          ***GMF**

end of file - request executed   5 times

-DONE
*CONV *=TRAX-EXEC/Wxxx/STAR-C:CARD,V
THE CURRENT FILE CONTAINED 12984 ASCII RECORDS ...
```

At this point, the user may resume with STEP 15, preparation of the NTPE object library.

5.2.10 Timesharing Subsystem Monitor. The Timesharing Subsystem Monitor (TSSM) is used to measure TSS performance. Section 15 details those reports available from the data collected by this monitor. This monitor should be used only on a WWMCCS WW7.2 system. If desired to be run on a commercial release, care should be exercised to ensure that all alters are located correctly (see subsection 5.2.10.1).

The TSSM causes the trace to be taken from many points in TS1; the collector builds its records which are then passed to the ER for buffering. An example of the record format appears in subsection 5.4.11. TSSM requires the following segment numbers from table 5-3 be used to generate the GMC R* file: 1, 10a, 11, 19, 23 and 36a. Subsection 5.6 shows how to generate an R* file.

THIS PAGE LEFT INTENTIONALLY BLANK

If all TSSM reports are wanted, then both the CPU Monitor and Channel Monitor are additionally required. Restricting Channel Monitor activity to SNUMB TS1 will conserve tape and GMC overhead. The TSSM and TPSM cannot be run concurrently.

The TSS Monitor consists of two major pieces of software: a data generation program loaded into TSS and a data capture routine loaded along with GMC routines. The TSS Monitor is used to analyze periods of poor TSS response to determine which users are affected, when and for how long they are affected, and what the possible causes of the poor response are. To collect information for such analysis, probes are inserted throughout TSS to gather individual pieces of information which, when combined by a data reduction program, yield the desired result. When active, the TSS Monitor is an area of code entered via transfer instructions planted throughout TSS by an initialization routine. After the TSS Monitor is entered from a probe point, it makes an entry in the GCOS trace table and then returns to the TSS process which was executing. The places where transfers are made fall into the following categories:

- Session identification and correlation - retrieval of line ID, DRUN, ID, USERID and User Status Table (UST) address (the unique identifier which is used to correlate traces coming from an individual session)

- Type of user interaction - retrieval of subsystem names, indication of build mode, memory sizes, derail codes

- Terminal I/O - a single location where I/O is started and a number of locations where courtesy calls for I/O complete are paid

- Disk I/O to user files - a single location where I/O is started and a number of locations where courtesy calls are paid for I/O complete

- FMS service - pairs of locations, one where TSS issues a request for service, and the other where the service is complete; the second location also gives an FMS status to tell whether the service was actually performed or why it was denied

- Processor allocation - locations for entering the processor allocation process, placing an entry into the subdispatch ready queue, and removing an entry from the subdispatch fault queue

- Memory allocation - retrieval of subsystem size; recording of the flow of control in the memory allocation and swap processes

- Errors and denials - retrieval of numeric error codes, recording of individual denial events

- Intervals of no TSS service to individual users - locations for recording line switches, DRL TASK jobs, wait disposition for batch jobs, console interaction for master user

- Intervals and events for the TSS executive - release of processor to allow subdispatch to start, execution of MME GEWAKE when there is no work to do, processing of TSS queue entries including TRACE ON and OFF, processing of terminal log-on requests prior to UST assignment

- Denials seen by the TSS executive - swap space denial, memory allocation denials, deadlock with other batch jobs at the remote inquiry MME GEROUT

- Unusual conditions - places where TSS recognizes that an error has occurred, errors such as bad terminal status, inability to assign a UST within 15 seconds, maximum number of users or VIPs logged on

5.2.10.1 Locations of TSS Trace Points. The following is a list of trace numbers, modules and symbolic locations provided with the TSS Monitor source code. The monitor has these locations coded as offsets relative to the beginning of the respective modules. During execution it retrieves the TSS load map and relocates these offsets. The current monitor uses offsets for W7.2.0 level 4. In future levels and releases, the symbolic locations will probably remain constant, but some octal offsets into the modules may change. Commercial users should ensure that all offsets are correct for the release they are operating under.

Trace	Module	Location	Instruction	Description
1	TSSB	.EV7+2	LDA .TSTOD	Log-on complete

Data: UST address, line ID
Use: First correlation of UST address with line ID (UST address used in all future traces for this session)

2	TSSC	DRM998+1	TRA 0,1	Swap space denied

Data: UST address
Use: Trigger resource denial report entry

3	TSSD	PATDP+3	LREG REGSAV	PAT denial

Data: UST address, flag (=0, no PAT space; ≠0, duplicate file name)
Use: Trigger resource denial report entry

```
  4    TSSD    PD621+18   LDA 2,DL              GEFSYE deallocate

Data:  UST address
Use:   Open interval closed by trace 5

  5    TSSD    PDCC+5     CMPXL =04046CC,DU     GEFSYE deallocate
                                                complete

Data:  UST address, status
Use:   Trigger entry in FMS report if status is octal 4046CC (SMC busy)
       SMC wait interval is FMS status trace, 30 or 31 (wait), FMS
       status trace, ..., until status not octal 4046CC

  6    TSSG    PNTTSX+9   LDAQ .LFNAM,2         Print error message

Data:  UST address, error code (binary)
Use:   Trigger entry in error report, possibly conditionally based on
       error number

 82    TSSH    CCMCL+4    CMPA =04CCCC,DU       Scan command list

Data:  UST address, command text
Use:   Connect command (e.g. RUN) with subsystem loaded (e.g. CDIN)

 64    TSSH    INTRP1+1   CANQ .FBT23,DL        Interpret primitive

Data:  UST address, current program stack tally, flag for 4C file
         permission, primitive number (1-11)
Use:   If flag set, explain return to top of stack; trace primitives
       not recorded by traces 8, 9, 10, 11, and 85 (e.g. XCALL,
       combination of CALLP, EXEC, and BIN)

  7    TSSH    POPPR1+2   CMPXC POPPR1          POPUP primitive

Data:  UST address, current program stack tally, subsystem name (ASCII)
Use:   Cause pointer in stack of subsystem names for user to be backed
       up; verify subsystem name in table against name in trace

 85    TSSH    POPPR1+21  LDQ 2,AU              DRL CALLSS complete

Data:  UST address, subsystem name (ASCII)
Use:   Cause pointer in stack of subsystem names for user to be backed
       up; indicate that subsystem issued DRL CALLSS and must be
       swapped back in

  8    TSSH    CAL1C+2    LDA 2,DL              CALLP primitive

Data:  UST address, program stack tally, subsystem name (ASCII)
Use:   Cause pointer in stack of subsystem names for user to be pushed
       down and name recorded (tally used to verify depth; no pushdown
       for DRL T.GOTO or DRL CALLSS immediately followed by DRL RETURN)
```

```
 9    TSSH    RSPRM+6    LDA -1,DU              Enter build mode
```

Data: UST address, program stack tally
Use: Cause entry to be made in stack of subsystem names; remove entry from stack when trace 13 is received

```
 10   TSSH    LODPRM+9   TNC SCERR6             EXEC primitive
```

Data: UST address, program stack tally, subsystem name (ASCII)
Use: Set flag to indicate actual code executed (as opposed to subsystems CARD, FORT, MDQ, etc.)

```
 11   TSSH    SYSPRM+4   LDA .LBUF,2            SYSTM primitive
```

Data: UST address
Use: Clear stack of subsystem names

```
 12   TSSI    CU1CF+7    EAX1 .LCALS,2          Log-off
```

Data: UST address
Use: Terminate processing of session, release UST address

```
 13   TSSI    STARTP+3   LDA 0,3                Command received
```

Data: UST address, command text (ASCII)
Use: Record last user build mode command for session snapshot reports; end build mode; clear subsystem name stack if command is "NONE" and does not come from DRL CALLSS or DRL T.GOTO

```
 14   TSSJ    LOGON+1    CANA =07777,DL         Periodic check
```

Data: Line ID if new user waiting
Use: Produce report entry if time interval between trace 102 and this one too great

```
 15   TSSJ    LOG11+3    STQ BYTIME             GEWAKE - no users
```

Data: Length of GEWAKE (clock pulses)
Use: Force clearing of tables in data reduction program; opens interval closed by trace 61

```
 16   TSSJ    LN100+2    STA .LKST,2            Break or disconnect
```

Data: UST address, GEROUT status (binary)
Use: Set flag for reconnect if status 2 (disconnect); simulate input complete if break and if I/O complete trace has been processed (user initiates request for service with break as in DJST)

17	TSSJ	.SSDSP+7	TNC TYTSS	GEWAKE until subdispatch done

Data: Length of GEWAKE (clock pulses)
Use: Formatted dump or accounting for GEWAKEs; opens interval closed by trace 61

18	TSSJ	TYTSS	STQ SLEEP	GEWAKE with subdispatch busy

Data: Length of GEWAKE (clock pulses)
Use: Formatted dump or accounting for GEWAKEs; opens interval closed by trace 61

19	TSSJ	QSPEC	EAA SPEC	All Points Bulletin or remote I/O courtesy call

Data: UST address, flags (.LFLG2) for CRUN/DRUN, GEROUT status
Use: Completes terminal I/O started by trace 100

20	TSSJ	STGCC+5	CMPA 4,DL	Build mode input received

Data: UST address, flags (.LFLG2), GEROUT status
Use: Completes terminal I/O started by trace 100

21	TSSJ	DSTAT+4	LDA .LIOST,2	SY** I/O complete

Data: UST address
Use: Completes disk I/O started by trace 24

22	TSSJ	RECON2+1	CANA .FK19,DL	Place user in reconnect mode

Data: UST address, flag for data in transmission
Use: Positive assurance that session is being put in hold (not given by trace 16)

59	TSSJ	PPTCC2+2	STZ .LCC,2	Disk I/O for tape mode complete

Data: UST address
Use: Completes disk I/O started by trace 24

104	TSSK	RETSBS+1	SZN .LSFLG	DRL processing complete

Data: UST address, A-register status if DRL T.SYOT or DRL TASK, Q-register status if DRL SPAWN or DRL PASFLR, DRL number if status reported
Use: End DRL processing state; report unusual status conditions

23 TSSK EXDRL2+2 CWL 0,3 Process DRL

Data: UST address, flag for data in transmission, DRL number, instruction counter
Use: Make entry into circular table for trace of last user activity; indicate roadblock for terminal I/O if restricted number; clear subsystem name stack if DRL SYSRET

24 TSSK LINK+3 LDA TSXSW Request file I/O

Data: UST address, flags for user I/O, trace made at courtesy call level, return address within range, module number for return address or actual return address, I/O queue address
Use: Common point for issuing I/O to files in a user's AFT; traces for I/O complete are 21, 25, 33, 35, 39, 43, 51, 53, 54, 59; queue address can be cross referenced to channel monitor traces; return address information can be used to ignore traces for which there will be no termination trace (calls from TSSN, module number=14 for CRUN I/O, for example)

25 TSSK DIOCC+1 STZ -1,3 Disk I/O complete

Data: UST address
Use: Completes disk I/O started by trace 24

26 TSSK DFRET1 EAQ 0 Denial from DRL DEFIL

Data: UST address, denial code
Use: Trigger resource denial report entry based on denial code value

27 TSSK FILC+7 STA 3,7 Issue MME GEFSYE

Data: UST address, FMS function code
Use: Open interval closed by trace 29 .

28 TSSK FRETC+1 STX3 0,1 GEFSYE denied - bad FILACT parameter

Data: UST address, denial code
Use: Enter into table of events for individual sessions

29 TSSK FILCC+22 LDX4 0,7 Courtesy call for trace 27

Data: UST address, FMS status
Use: Closes interval opened by trace 27; triggers entry in FMS report if status is octal 404600 (SMC busy); opens SMC wait interval if status is octal 404600 until status changes

30	TSSK	DELY1	STQ .LTEMP,2	DRL delayed 200 ms

Data: UST address
Use: Verifies SMC wait interval

31	TSSK	DELY2+3	LDQ 2,DL	DRL delayed 2 seconds with DRL GWAKE

Data: UST address
Use: Indicates following DRL GWAKE is for SMC wait and is not coded in user program

32	TSSK	USER+1	STAQ .LID,2	Store USERID in UST

Data: UST address, DRUN ID or zero, USERID (obtained by GMC capture routine)
Use: First time USERID is available for a session; DRUN ID not available when log-on trace is made

33	TSSK	PBIOCC+2	LDA .LIOST,2	SY** I/O (PASUST, code -1) complete

Data: UST address
Use: Completes disk I/O started by trace 24

34	TSSK	MOR2+1	LDX1 RETFLG	DRL MORLNK error

Data: UST address, error code
Use: Trigger resource denial report entry or enter into table of events for user

35	TSSK	CC.1+2	STZ .LCC,2	I/O for DRL SPAWN or PASFLR complete

Data: UST address
Use: Completes disk I/O started by trace 24

36	TSSK	PASGD+4	ANA -1,DL	Batch job submitted

Data: UST address, SNUMB, wait disposition flag
Use: Enter into table of events for user; explain delay if trace 60 follows

37	TSSK	PSFLX+1	STQ 5,1	Error in DRL SPAWN

Data: UST address, error code
Use: Enter into table of events for user

38	TSSK	WRTABT+1	CANA .FP19,DL	Enter routine to write dump to ABRT

Data: UST address
Use: Provide reason for following disk I/O

39	TSSK	ABT70+1	STZ .LSWAP+1,2	ABRT file I/O complete

Data: UST address
Use: Completes disk I/O started by trace 24

40	TSSK	RESTOR+12	AOS 5,1	Increment number of executions for subsystem load via DRL RESTOR

Data: UST address, subsystem name (ASCII)
Use: Provide information on what user is doing

41	TSSK	RES20+12	STA .LSWAP+1,2	Load subsystem with DRL RESTOR

Data: UST address
Use: Open interval for disk I/O not covered by trace 24, closed by trace 42

42	TSSK	RESCC+4	STZ .LSWAP+1,2	Subsystem load complete

Data: UST address
Use: Close interval opened by trace 41

43	TSSK	RESCCX	STZ .LSWAP+1,2	Permanent file I/O from DRL RESTOR complete

Data: UST address
Use: Completes disk I/O started by trace 24

44	TSSK	CGROT1+4	STA 1,3	Start line switch

Data: UST address, JDAC switch (wait), BCD name
Use: Start roadblock interval if switch set (indicated by trace 60); provide information about what user is doing

45	TSSK	GROW01+2	SXL5 .LTEMP,2	Call .MFS19 for file grow

Data: UST address
Use: Open interval closed by trace 46

```
 46    TSSK    GROWCC+4   CMPX4 =04046CC,DU     File grow complete

Data:  UST address, FMS status
Use:   Close interval from trace 45; open SMC wait interval if status
       is octal 4046CC

 47    TSSK    CONO15+2   LDA 4,DL              Start console I/O for
                                                DRL CONSOL

Data:  UST address
Use:   Open interval closed by trace 60

 48    TSSK    JSNUM4     LDA 2,3               Return from DRL JOUT

Data:  UST address, status
Use:   Report along with other return codes if -1 (batch system full)
       or 1 (lost PAT)

 49    TSSK    JSNUM7+6   LDA =0000200,DU       Output busy status
                                                for DRL JOUT

Data:  UST address
Use:   Report along with other return codes

 50    TSSK    GWAKE+4    SXL1 .LBACK,2         Start DRL GWAKE

Data:  UST address, time (seconds)
Use:   Explain part of interval involving no user interaction

 51    TSSK    TSKCC+2    STZ .LCC,2            I/O for *J write for
                                                DRL TASK complete

Data:  UST address
Use:   Completes disk I/O started by trace 24

 52    TSSK    TSK325+5   LDA 4,DL              Start DRL TASK

Data:  UST address, SNUMB
Use:   Opens interval for user roadblock indicated by trace 60

 53    TSSK    TSKCC1+2   STZ .LCC,2            I/O for *J read for
                                                DRL TASK complete

Data:  UST address
Use:   Completes disk I/O started by trace 24

 54    TSSK    TRUECC+3   STZ .LSWAP+1,2        I/O for DRL SAVE
                                                complete
```

Data: UST address
Use: Completes disk I/O started by trace 24

55	TSSK	ATTRI1+5	STZ 0,0	Call .MFS13 for IDS attributes

Data: UST address
Use: Opens interval closed by trace 56

56	TSSK	ATTRCC+4	CMPX4 =0404600,DU	Status from IDS attributes

Data: UST address, FMS status
Use: Closes interval opened by trace 55; starts SMC wait if status is octal 404600

57	TSSK	SYTE1	STA 4,0	Denial from DRL T.SYOT

Data: UST address, denial code
Use: Report along with other denial codes

58	TSSK	DC1.0+29	LDA .LRTLL,3	UST address switch with DRL T.CONN

Data: Old, new UST addresses
Use: Informs data reduction program to perform same switching function

60	TSSL	ALLOC	AOS A.TRCT	Allocator services

Data: UST address, function code, flags (.LFLAG)
Use: Closes intervals for some DRL processing; opens and closes intervals for non-TSS processes

61	TSSL	LINSRV	LDX3 .QFHT+.TSSDQ	Enter processor allocation

Data: Contents of index register 2 (last UST address)
Use: Closes intervals opened by traces 15, 17, and 18; if after trace 23 for DRL TASK, no pushdown file available; ends DRL processing state for some DRLs

62	TSSL	MAP+10	SZN A.URWT	Enter memory allocation

Data: Flags for urgent user present, fence up too long
Use: Provide information in dump for analyzing memory allocation flow

77	TSSL	MAP.4+2	EAA 0,3	Attempt memory allocation

Data: UST address, program size
Use: Provide information on current TSS memory demand, initiate memory wait for user

63	TSSL	SDP+1	LDA .LSIZE,2	Enter swap decision processing

Data: None
Use: Compute memory demand ratio of number of trace 63 versus number of trace 62

65	TSSL	SDP.4+6	CMPA .TASCF	Consider TSS size increase

Data: Indicator register (carry bit)
Use: Provide information for swap flow of control

66	TSSL	SDP.4C+3	AOS .TSIRC	Initiate size increase for urgent user

Data: UST address for most urgent user, new TSS size
Use: Count occurrences of trace; provide information for swap flow of control

67	TSSL	SDP.5A+6	STA A.LUTM	Set up fence for urgent user

Data: UST address
Use: Count occurrences of trace; provide information for swap flow of control

68	TSSL	SDP.5C+4	CMPX1 A.URMR	Set up urgent user class memory reserve

Data: UST address, size of urgent user class memory reserve
Use: Count occurrences of trace; provide information for swap flow of control

69	TSSL	SDP.7+2	LDX0 1,DU	Force swap

Data: UST address for user force swapped
Use: Count occurrences of trace; provide information for swap flow of control

70	TSSL	SDP.8	LCA .FL34+1,DL	Terminate swap process

Data: UST address of user which could not be swapped
Use: Count occurrences of trace, give TSS snapshot ("system is in trouble")

71 TSSL SPMACT+10 ASA 1,3 UST area increase by 1K

Data: None
Use: Deduct 1K from memory available for user programs

72 TSSL SPM.2A+11 ARL 10 Issue GEMORE for memory

Data: Number of K words requested
Use: Provide flow of control information; give explanation if trace 75 follows immediately

73 TSSL SPM.2B+3 ALS 1 GEMORE successful

Data: Number of 512-word blocks added
Use: Completes interval opened by trace 72; adds to memory available for user programs

74 TSSL SPM.3+22 QLS 9 Memory release

Data: Number of words released
Use: Deduct from memory available for user programs

75 TSSL SPM.5 LDA .TSTOD GEMORE refused or reduction not possible

Data: None
Use: Count occurrences of trace; indicates GEMORE refusal if trace 72 precedes

78 TSSL MMV.1A+6 ANQ =0777000,DL Memory map change

Data: UST address, subsystem size and LAL, extra buffer memory LAL
Use: Completes memory wait for user or indicates memory release

79 TSSL SWOUT+3 ALS 18 Swap user program

Data: UST address, program size
Use: Provide flow of control information (may be able to delete trace if 80 sufficient)

80 TSSL .ATCHG+12 LDA -1,3 Change time types

Data: UST address, old and new time types (range 61-66 for non-useful memory residence time, swap/load time, useful memory residence time, out of memory time, waiting for normal memory allocation time, waiting memory allocation after forced swap time)
Use: Maintain state of user interaction

81	TSSL	STB.2+4	LDA BUFSIZ,DU	Request extra buffer memory

Data: UST address
Use: Open interval closed bvy trace 78 or 83

83	TSSL	STB.4	ORSA .LPQF,2	EBM refused

Data: UST address, flags (bit 20=1, no memory; bit 21=1, buffer full)
Use: Ignore if bit 21 of flag=1; count and terminate interval from trace 81 if bit 20=1

84	TSSL	KICC+2	SZN .TCCMT	Terminal input complete

Data: UST address, flags (.LFLG2), GEROUT status
Use: Completes terminal I/O from trace 100; indicates whether user is receiving output on terminal

86	TSSL	APBCC+3	CANA 2,DL	All points bulletin complete

Data: UST address, flags (.LFLG2), GEROUT status
Use: Completes terminal I/O from trace 100

87	TSSL	KSTAT+1	EAX2 -.LKCC-1,3	Terminal I/O complete

Data: UST address, flags (.LFLG2), GEROUT status
Use: Completes terminal I/O from trace 100

88	TSSM	EXENTR+30	CANQ =0200000,DL	Remove entry from sub-dispatch fault queue

Data: UST address, fault type, processor time used (clock pulses), I/O interrupt flag
Use: Completes interval opened by trace 91; provides information for formatted dump

89	TSSM	EXTSQ	STX1 EXTSR,$	Enter routine to process queue

Data: None
Use: Maintain intervals between these traces; report if time excessive

90	TSSM	EXTVRB+1	STA EXTS6F+5	Process unrecognized console verb

Data: If TRACE ON, UST address, USERID, line ID, DRUN ID, program stack (subsystem names), subsystem BAR and LAL, time type, flags (.LFLG2, .LPQF, .LFLAG), log-on time, EBM address, swap area size for all TSS users (obtained by GMC capture routine;

TSS passes flag to GMC to write initialization records or GMC writes initialization records after it has recovered from a lost data condition)
If TRACE OFF, flag (=0)
If bad console input, no trace

Use: Initialize tables in data reduction program

91	TSSM	RETS3X+6	CMPX1 1,DU	Make subdispatch entry

Data: UST address
Use: Open interval closed by trace 88

92	TSSN	GUST1+2	CMPA .LTSRI,DL	Process log-on request

Data: Line ID (octal 2020 if deferred), reject flags (.TLFLG bit 32 or 35)
Use: Open interval closed by trace 1

93	TSSN	GUST21+15	TRA GUSTR-1	Reject user - bad line status

Data: Line ID
Use: Report along with log-on rejections (this is an exotic line condition trace)

94	TSSN	GUST2A+4	CMPQ =030000,DU	Check for VIP as terminal type

Data: Terminal ID, type, number of VIPs allowed (.T760), number logged on (.TL760)
Use: Explain trace 96 if it occurs

95	TSSN	GUST4+4	CMPQ GUSTT	Check UST wait time

Data: Line ID, flag if wait greater than 16 seconds
Use: If flag set, report as log-on reject

96	TSSN	GUSM+5	STA 2,3	Reject user

Data: Line ID
Use: Report trace if not already done

97	TSSN	MUST3+2	AWDX ,2,0	UST compression

Data: Old, new UST addresses
Use: Update UST-oriented tables with new address

98	TSSN	ASGCC1+4	ASQ 1,0	UST area increase by 1K

Data: None
Use: Decrement amount of user memory available

99	TSSN	RELWC3+18	ASA 1,C	UST area decrease by 1K

Data: None
Use: Increment amount of user memory available

100	TSSN	RIO.1A+2	LDA .LBUF,2	Terminal I/O request

Data: UST address, GEROUT opcode, CRUN flags (.LFLG2)
Use: Open interval closed by trace 19, 20, 84, 86, 87, or 103

101	TSSN	R.SFOK+2	LXL1 -1,6	Process command file $*$ function

Data: UST address, text (ASCII)
Use: Provide information to explain changes in .LFLG2 or to explain trace 6

76	TSSN	TRMCF+7	EAQ 0,AU	Cancel CRUN mode

Data: UST address, flags (.LFLG2)
Use: Force clearing of subsystem name stack

102	TSSN	IRINQ+4	LDX1 =0050000,DU	Issue remote inquiry GEROUT

Data: None
Use: Open interval closed by next trace; if long time, indicates roadblock of TSS in .MROUT because a table is full

103	TSS0	760CC+3	CMPA 4,DL	VIP input complete

Data: UST address, flags (.LFLG2), GEROUT status
Use: Completes terminal I/O started by trace 100

5.2.10.2 Formats of TSS Traces. The collection routine executing within TSS passes two-word traces to GMC through the dispatcher trace mechanism. In all but four trace subtypes, GMC stores the A and Q registers after a record control word and RSCR time to form a four-word logical record. For subtype 24, GMC edits the Q register before making a trace. For subtype 32, GMC additionally retrieves the USERID and appends it to form a 6-word logical record. Subtype 105 is generated internally by GMC when it senses a period of no TSS activity. Logical records generated as a result of receipt of a subtype 90 trace may have 3 different formats. The first format is written whenever TSS passes a subtype 90 trace to GMC indicating that the TSS traces have been turned off. GMC passes this trace to the data tape in the standard format described above. The following 2 formats apply to a group of logical records written by GMC whenever it

recovers from a lost data condition or whenever TSS passes a subtype 90 trace indicating that TSS traces have been turned back on. In the latter case, GMC formats a series of logical records: one 16-word record is written for each user logged onto TSS to give such attributes as UST address, line ID, USERID, and current program stack. A four-word record which indicates the TSS swap area limits is the last record written by GMC in the sequence of records. If no users are logged on to TSS and if TSRI is not running, only the four-word record is written. The Monitor source code gives the formats of all TSS record subtypes.

5.2.10.3 Installation Procedures

5.2.10.3.1 Description of Monitor Software. The TSS Monitor program element executes as a TSS master subsystem. When the master user logs on and requests the program via "SYST GMF", it first performs address relocations needed to convert relative offsets into actual slave addresses. Then it verifies that the instructions at the trace points match those in the monitor coding (the originals have transfer instructions patched over them when the monitor is executing). Any mismatches found are reported on the master terminal and verification continues. If any mismatches are found, the subsystem terminates without making any modifications to TSS. If verification succeeds, the master subsystem copies part of itself (about 1K memory) into an area of module TSS0 which was reserved by means of patches on the TSS INIT file. The origin of this area is the UST address origin which TSS would have used were the patch not applied. Next, the master subsystem applies Execute Double (XED) instructions to most trace points to save the return address and indicator register. In a few rare cases, unconditional transfer instructions are used. The master subsystem terminates at this point. As TSS continues execution, traces will be formatted, but a switch prevents GCOS traces from being written. If the console operator enters "TS1 TRACE ON", GCOS traces will be written from TSS, and an additional overhead will be imposed on TSS. The console operator can stop generation of TSS traces with "TS1 TRACE OFF". The traces can be turned on and off multiple times.

5.2.10.3.2 Software Installation. When TSS is loading the monitor subsystem, it must be able to access the program element using a MME GECALL to a BCD name of .TSGMF. The program element may reside either on a system file defined in the EDIT and FILES sections of the boot deck or on a permanent file (B29IDPXO/GMFCOL/TSS/TSSMON) dynamically accessed during TSS startup. Use of a permanent file requires additional patches in the TSS INIT file. This is the current method of implementation because it does not require changes to the GCOS startup deck. The job stream used to create the program element is located on file B29IDPXO/GMFCOL/TSS/TSGMF.

If the system file option is to be used, the $ PRMFL Q* card on file TSGMF must be replaced with a $ TAPE Q*,X2D,,,,TSSMON card. The startup file must be defined in the EDIT section of the boot deck as follows:

```
$     FILDEF   ST1,TSSMON,12/0,SYS,1T1
```

The tape drive name 1T1 must be appropriate to the hardware configuration. If this startup file is appended onto an existing edit tape, replace "1T1" with "*". In the FILES section, insert the following card in front of existing $ SYSTEM cards:

```
$     SYSTEM   TSSMON
```

If a permanent file is used, no changes need to be made to the boot deck, and no system interruption occurs during installation of the TSS Monitor.

5.2.10.3.3 Software Activation. The purpose of this section is to describe how TSS builds tables so that the master user can find a subsystem named "GMF".

5.2.10.3.3.1 Overview of TSS INIT File Changes. The TSS INIT file is a quick access permanent file named TS1 normally residing under USERID OPNSUTIL, the default USERID for TSS and patchable as symbolic location .TUSER. File TS1 is read as soon as TSS starts in order to pass parameters for the current loading of TSS or to allow symbolic specification of site option patches such as the maximum number of concurrent users. The TS1 file has two sections: $INFO to symbolically define site option parameters, and $PATCH to apply patches to TSS beyond those already in the PATCH section of the boot deck. Installation of the TSS monitor requires that patches be placed at the end of the TSS INIT file. These patches are located on file B29IDPX0/GMFCOL/TSS/TSS.PAT.

5.2.10.3.3.2 Definition of the Master Subsystem Name. The following three patch cards, included in file TSS.PAT, overlay an unused program descriptor in TSSA to define a subsystem named GMF with an edit name of .TSGMF and attributes .BPRIV (can execute privileged DRLs, not currently used), .BMAST (master subsystem), and .BMASX (permission to alter TSS executive with a DRL instruction, not currently used):

```
6724   OCTAL   147155146040   GMF                    .MTIMS
6725   OCTAL   336362274426   .TSGMF                 .MTIMS
6726   OCTAL   40003          MASTER SUBSYSTEM       .MTIMS
```

5.2.10.3.3.3 Definition of the UST Origin. The following two patches, included in file TSS.PAT, disable the check for overlaying VIP code in TSS0 when no VIPs are configured. They define the UST

origin to be at location 3760 with respect to the beginning of TSS0. The first patch NOP's out a conditional transfer instruction around an instruction which loads a UST origin corresponding to the case when VIPs are configured. The letter "0" in these patches means that the patch locations are with respect to the origin of TSS0, not to slave address 0 of TSS as in the subsystem attribute patches. The letter "R" before patch content indicates that the address field of the patch must be relocated to the beginning of the module (TSS0) identified in column 7 of the patch. The symbolic addresses for these patches are IN7630+2 and IN7630+3. The following patches complete the TSS INIT file when the TSS Monitor is loaded on a startup file:

```
5271  00CTAL   11007           DON'T TRANSFER IF NO VIPS      .MTIMS
5272  00CTAL   R3760235007     ALLOW 1K FOR GMF               .MTIMS
```

5.2.10.3.3.4 Installation from a Permanent File. The principle of this method is that if a job has a file code ** active, any MME GECALL will cause that file to be searched before the system files are searched. Patches on the TSS INIT file must access a permanent file before subsystem loading begins and release it after subsystem loading finishes. The TSS User Derail Loader (TUDL, SDN K79005) provides an example of how to access and release a system loadable file from within the TSS executive. The patches on the TSS INIT file for the TSS Monitor are compatible with TUDL because TUDL starts its work after the TSS Monitor has finished its work. TUDL may further relocate the UST origin upward, but TUDL uses the address stored by the TSS Monitor patches. A list of an entire $ PATCH section is given in the TSSM source code. The first four patches match ones described in earlier sections.

The bulk of the patches used to access a permanent file are placed in an area of TSS0 which is reserved for UST space when the UST origin is not adjusted upward (actually, this space was needed in releases prior to W7.2.0 to prevent the generation of a UST for TSR1 from destroying the code at the end of TSS startup; with the TSS INIT file feature, enough code intervenes so that this buffer is not necessary and so that the UST origin can be moved up and not destroy the code). The first transfer into these patches is at offset 4673 in TSS0 (symbolic offset DMYERR+2, instruction EAX5 1). Here, the USERID in the file structure defined in the patches is stored into the SSA of TSS at location .SUID. Then a MME GEMORE accesses the file. Use of .SUID makes FMS think that the file is being accessed by its owner and thus, no special permissions are needed. If the MME GEMORE is denied, a flag is set to 0 before control returns to the original TSS0 coding from offset 2013 in the patches.

The second transfer, at octal offset 5272 (symbolic location IN7630+3) replaces the instruction defining the UST origin. The patch at octal offset 5271 remains intact. In the second group of patches, the

permanent file must be released, the number of PATs decremented by one, and the USERID in .SUID set to zero. If the GEMORE in the first set of patches is denied, then the number of PATs is not decremented. This alteration of word .SNPAT in the SSA is necessary because releasing the file does not remove the file code from the SSA. The patch for defining the new UST origin is moved to octal offset 2031 in the patches, just before a transfer back to TSS0.

5.2.10.4 Production Use of the Monitor. The following steps must be taken to enable data capture:

1. Append patches on file B29IDPX0/GMFC0L/TSS/TSS.PAT to Timesharing INIT File (normally OPNSUTIL/TS1).

2. Run the file B29IDPX0/GMFC0L/TSS/TSGMF to create the new TSS subsystem file.

3. Start TS1.

4. Start a copy of GMC with the TSS Monitor active.

5. Log on to a master USERID and enter "SYST GMF" at the "*" prompt following log-on to install the hooks into TSS code. Traces will not start at this time. The master USERIDs assembled in TSS are MASA and MASB. Unless these are patched or redefined in the TSS INIT file, only a terminal ID designated as master in .MSECR may be used for this step.

6. Enter "TS1 TRACE ON" from the system console to start generation of traces from TSS.

7. Enter "TS1 TRACE OFF" from the system console to suspend generation of traces.

Steps (5), (6) and (7) execute independently of (4); however, use of step (6) without use of step (4) will cause unnecessary TSS overhead if traces are being generated and lost due to GMC not being in execution. Steps (6) and (7) may be repeated multiple times if traces are to be captured for specific periods of the day. (The companion data reduction program, described in section 15 cannot process GMC sessions longer than 9 hours).

5.2.10.5 Monitor Limitations. To obviate lockup fault, initial trace generation (at "TS1 TRACE ON", and following lost data) is inhibited until the number of TSS users falls below 50. Further, all trace generation is inhibited until the number of users falls below 100. (The companion data reduction program is limited, by parameter, to 50 active USTs).

5.3 Processing

The GMC requires one tape drive, 15K-24K words of storage, and, if the multireel option will be used, 300 links of disk storage to execute. The actual memory required will depend on the number of monitors selected for execution. The GMC will lock itself into memory, assuring that it cannot be swapped or moved during the run time. An initialization procedure will attempt to relocate GMC out of the memory growth range of Time-Sharing (TS1) and, in addition, relocate GMC to the high or low end of a quadrant of memory. Due to this feature, GMC can be started at any time of the day, without fear of causing memory fragmentation or degraded TSS response.

Immediately prior to beginning collection of data, GMC will attempt to relocate out of the TSS (TS1) memory growth range. On a system with

more than 256K of memory, the growth range of TSS (TS1) is assumed to be 180K, while on a system with less than 256K of memory the growth range is considered to be 100K. During this relocation procedure GMC might grow in size; however, the user need not be concerned, since GMC will release all unneeded memory after the relocation operation is completed. If GMC cannot relocate out of the growth range of TSS (TS1), the GMC will abort with a CM abort. The user can override this abort with a data card option, but he should do so with great care since TSS response might be severely degraded if GMC prevents TSS (TS1) from growing in size. The procedures for overriding the abort can be found in subsection 5.5.4.

After GMC succeeds in relocating out of the growth range of Time-Sharing, GMC will attempt to relocate itself to the high or low end of the memory quadrant. It will relocate as far away as it can, but it will not abort if it is unsuccessful. Therefore, slight memory fragmentation is still possible, but this should cause little if any problem to the operation of the computer.

If multiple copies of TSS are present, the current version of GMC will move out of the growth range of only TS1. Therefore, the possibility does exist that GMC will be within the growth range of other copies of TSS. The user can check this by doing a List Core Console command. If GMC is within the growth range of some copy of TSS, the user may want to abort the GMC.

5.3.1 Executive Routine. The Executive Routine (ER) controls the processing of the GMC. The ER controls all inputs, outputs, start and stop time, and all required error processing. The ER also hooks the GMC into the dispatcher.

The ER begins processing by reading the input parameters on the data card. The procedure for determining the required data card parameters for GMC operation is fully described in subsection 5.5. Using these parameters, ER determines which monitors will be active during the run. If a start time is specified, ER will remain idle until the start time is reached. During this time, it is possible that GMC will be swapped. The ER will cause the GMC to terminate at a specified stop time if the option is present. Multireel output is the standard default for GMC; however, the user can request GMC to terminate after one or more reels of data have been collected. The ER controls all error aborts within GMC. Abort codes and reasons for the abort are shown in table 5-2. When ER begins processing, it will query the operator for the first tape number with the message:

*FOR XXXXX WHAT IS THE MOUNTED TAPE'S NUMBER?

where XXXXX is the SNUMB assigned GMC. The operator must enter the reel number of the tape to be used for GMC output. If the operator

Preceding page blank

does not reply to the message, the question will be repeated. It is important that all tapes requested by GMC contain a write ring and be mounted as quickly as possible.

The GMC will execute a GEWAKE statement (and may be swapped) if it is not desired immediately (user has requested a time option). When the ER determines that it is time to begin monitoring, it will move the GMC out of the growth range of TSS (TS1). Upon successfully moving GMC, the ER will issue a message to the operator indicating which monitors have been selected for execution. The monitors will be referenced by the mnemonic code shown in table 5-1.

For example if the Mass Store Monitor, Channel Monitor, and Communications Analysis Monitor are active, the following message will be printed:

```
*MONITORS MSM CM CAM
```

If this message does not appear within 5 minutes of the monitoring start time, the GMC has encountered a severe problem in its attempt to move. Such a problem should be a rare occurrence. If this problem does occur, GMC should be aborted and restarted.

If the multireel option (see subsection 5.5.2) is specified (i.e., an M9 or M91 does not appear on the GMC data card), the GMC will issue the following message to the operator when it reaches the end of a reel:

```
*FOR XXXXX MOUNT Y REEL ON ZZZZZZ
 FIRST TYPE NEW TAPE # -
```

where XXXXX represents the SNUMB of GMC, Y represents the sequence number of the next data tape, and ZZZZZZ is the tape drive address. Should the operator be distracted for more than 30 seconds, GMC will reissue the message.

Once the operator gives the GMC a tape number, the ER rewinds the data tape and waits one minute for the tape to rewind. The ER then checks the tape status, and if the next tape has not been mounted, the ER issues the message:

```
**AGAIN MOUNT TAPE * XXXXX FOR MONITOR YYYYY ON ZZZZZZ
```

where XXXXX is the reel number, YYYYY is the SNUMB assigned to GMC, and ZZZZZZ is the tape drive address.

This procedure of waiting one minute and reissuing the message is repeated until the tape is mounted. During this timeframe, the ER is writing any full buffers collected to its overflow disk file. Since

this file can fill and cause an abort, the operator should ensure that the data tape is mounted as soon as possible.

After the GMC processes the data card, individual monitor initialization is accomplished. After initialization is completed, the memory used for the initialization code is used for buffering the collected monitor data prior to it being written to tape. GMC is coded for extended memory systems with the exception of nonextended memory. When nonextended memory is sensed, the extended memory instructions throughout the program are changed. If the initialization is accomplished successfully, GMC will "hook" itself into the dispatcher and become an extension of the dispatcher. The normal system trace procedures consist of the following three unique instructions:

```
LDAQ **,7*
STAQ **,2
TRA  0,1
```

The GMC modifies this normal sequence in the case of extended memory. Each trace event will preserve the current MBA value, set the MBA value for the GMC, and transfer to the GMC. The resulting interface includes the following:

```
SMBA  MBASTR,7
LMBA  GMCSMBA,DU
TRA   GMC.
```

The interface is completed by execution of the

```
LDAQ  **,7*
STAQ  **,2
```

in GMC, which completes the normal trace procedures. In a system with nonextended memory, the interface is less complex, resulting in one instruction change in the trace procedures. The

```
LDAQ  **,7*
STAQ  **,2
TRA   0,1    of the trace
```

is altered to look like

```
LDAQ  **,7*
STAQ  **,2
TRA   GMC.
```

Whenever the dispatcher executes a trace after this point, the GMC ER will gain control and pass the system control over to the proper monitor for data collection. Control is then returned to the ER for return to the dispatcher.

In the process of collecting data, a monitor may desire to save a logical record on the collection tape. In order to sequence all logical records to the tape file, a common buffering system has been provided by the GMC. Any monitor can pass control to symbol BUFCTL for buffering a logical record to be saved on the tape file. This code is executed in master mode as an extension of the operating system function where control has been taken. Writing the buffered data to tape is a GMC slave function which requires controls between the buffering system and the writing system. The buffering system indicates, via software to the slave, when an individual buffer is to be written. After the slave has written the buffer, it is returned to the buffering system. The slave code for writing the tape is started at symbol BUFCHK. This slave portion is normally in GEWAKE mode and is awakened either by GMC master mode code or by a normal dispatch. The slave function will check for a full buffer and write it before going back to sleep. This process is repeated until GMC has been terminated, at which time the slave portion will execute the proper abort code via a MME GEBORT. The wrap-up procedure then performs all termination requirements. GMC will terminate via direction from the time option on the data card, limited reel option on the data card or else via console command. If no time option or limited reel option has been selected, the GMC will continue to process until it is ABORTED via the console.

At termination, the wrap-up location is given control and the termination record is processed. This is the slave termination of GMC. Also associated with termination is the return of the dispatcher code to its original state. This is accomplished by GMC checking its own .STATE word for a termination code. When the termination code is sensed, the dispatcher hook is removed, the original code is replaced, and the slave termination process is completed. Upon termination, ER reads the communication region table, writes the data to tape with a termination record, removes any patches it has placed in the system, and proceeds to wrap up.

5.3.2 Output. Output of GMC is always a magnetic tape file. The following is a guideline for the number of tapes generated by each monitor (800 BPI, 2400 feet tape reels):

MUM - one tape every 6-8 hours. With Idle Monitor, one tape every 3 hours.
MSM - one tape every hour

CPUM - 1 tape every 24 hours
TM - 1 tape every 24 hours
CM - 1 tape every half hour
TPSM - 1 tape every 4 hours
CAM - 1 tape every 8 hours
GRTM - 1 tape every 24 hours
TSSM - 1 tape every 2 hours

Analysis of this tape(s) is accomplished by a series of dedicated data reduction programs which are discussed in sections 6 through 12 and section 15. It is essential that the proper data reduction program is run against data created by its associated monitor routine.

The GMC is usually terminated by operator request. However, there are times when the monitor may terminate itself. In most of these cases, the operator will receive a console message indicating the reason for the abort (see table 5-2). However, under some circumstances the message the operator sees is that GMC terminated with "NO REASON SPECIFIED." In this case, if the job listing is examined, a reason will be given under the WRAPUP activity information as shown below.

```
* ACTY-01 $ CARD #0008* GELOAD  03/16/78  SW=000000000000
NO REASON SPECIFIED          AT 030440 I=5060 SW=000000000000

*WRAPUP BEGUN
(USERS BS MME GEBORT)        AT 031413 I=0000 SW=000000000000
```

5.4 GMC Data Records

5.4.1 GMC Executive. The GMC Executive produces an initialization record that must be read by every data reduction program. This record contains information on the system configuration and the status of various queues at the time the monitor started. The length of the record is dependent on the size of the configuration. The Executive also writes a termination record whenever it terminates normally.

Initialization Record

Word	Bits	Information
1	0-35	Block Control
2	0-35	Zero
3	0-17	Year (.CRJCD)
	18-35	Julian day (.CRJCD)
4	Not used	
5	0-35	Current date (.CRDAT)
6	0-35	Current time (.CRTOD)
7	0-35	Reg A of RSCR 32

8	0-35	Reg Q of RSCR 32
9	0-35	Software release (.CRSR)
10	0-17	Size of HCM (.CRHCL)
	18-35	Not used
11	0-35	I/O Channel limits (.CRICL)
12	0-35	Hardware IOM type .CRIOC
13	0-17	Not used
	18-35	Memory size (.CRMSZ)
14	0-17	Tape sieve parameter (.CRMTP)
	18-35	Link sieve parameter (.CRMTP)
15	0-17	Number of IOM's configured (.CRNIC)
	18-35	Length of system file table (2 words per file)
16	0-17	Number of processors configured (.CRNPC)
	18-35	Shared System Number (.CRSSN) This value applicable only if bit 2 in word 19 is set.
17	0-35	CPU type (RSW 2 instruction)
18	0-35	BCD System ID (.CRSID)
19	0-35	System configuration (.CRFIG)
20	0-17	1st user program number (15 for WWMCCS, set by FSTSLV)
	18-35	first program number for GEIN (64 - number of card readers)
21	0-35	GMF LAL and full size (.SALIM+1)
22	0-17	GMF program number
	18-35	Zero
23	0-35	Idle time for processor 0 (.CRIDT)
24	0-35	Overhead time for processor 0 (above 2 words are repeated for each configured processor (word 16 bits 0-17))
25+	0-17	SCT address for system file 1
	18-23	Device type for system file 1
	24-29	Device number for system file 1
	30-35	Not used
26+	0-17	Starting llink number on device for system file 1
	18-35	Length of system file 1 in llinks (word 25+ and 26+ are repeated for each system file. Total number of words for this portion of record is found in word 15 bits 18-35.)
27+	0-35	Module directory (MDD) table address and length (.CRMDD)
28+	0-35	GECALL name (GMD) table address and length (.CRGMD)
29+		MDD table-one word per entry, plus one extra word if length is odd

30+		GMD table-one word per entry, plus one extra word if length is odd
31+	0-35	SCT table address (.CRSCT)
32+	0-17	Number of devices described below
	18-35	Not used
33+	0-35	SCT information for each channel on each IOM. Since both MTS0600 and MTS0610 tape drives have the same device type, bit 25 in the corresponding .CRCTl word is used. If this bit is set, then 610 tape drives are indicated by device types 03 and 04 for 7- and 9-track tape drives. The length of this portion of the initialization record is 48 times the number of IOM's configured (word 15 bits 0-17) plus the contents of word (32+) bits 0-17.
34+	0-35	GMF tape number
35+	0-35	Monitors active indicator - Bit i set corresponds to monitor i active

Following the GMC initialization record is a special record describing the CPU configuration, channel availability, and memory utilization. Depending on the monitors selected for execution and on the particular GCOS trace which first passes through the dispatcher "hook," a few traces may precede this record (e.g. during their initialization routines, the mass storage, channel, communications, and GRTS monitors write records). During the execution of GMC, this record is written again every five minutes. This is done in order to capture any releasing or assigning of memory that may have occurred. Since this record is generated every five minutes, it is possible that the correct memory configuration will not be available to the data reduction program for at most a 5-minute period of time. This will result in only a minor distortion of reports and will prove insignificant. If a central processing unit is dropped or added, the collector will recognize this fact immediately and generate this reconfiguration record immediately. If a channel or IOM is dropped or added, this will be reported by the reconfiguration record that is generated every five minutes. It should be noted that if multiple changes occur within a 5-minute period, only the final change will be reported.

The following communication region cells are constantly monitored for changes, since a processor reconfiguration may generate numerous traces while it is being performed:

1. .CRNPC - Lower half specifies number of processors released

2. .CRXPC - Upper specifies processor dedicated to T&D

3. .CRCMC - Six-word processor assignment table maps logical processor number to a physical port on the SCU; table is also updated if a program sets its .STATE word to force execution on a specific processor.

4. .CRMDB - Specifies information about SSA cache; may change during a memory release operation.

Because of the number of events involved in reconfiguration, multiple records may be written for the same reconfiguration. System programs (e.g. T&D) may cause a record to be written if they withdraw a processor from normal dispatching. They may also be written if a program sets its .STATE word to restrict it to or from a specific processor. Such a program must also update a counter within the .CRCMC table.

The format of the reconfiguration record is as follows:

Word	Bits	Information
1	0-17	Number of words following (record control word - RCW)
	18-35	Octal 500100 (Special record)
2	0-35	RSCR time
3*	0-17	Number of processors configured (.CRNPC)
	18-35	Number of processor released
4*	0-17	Processor dedicated to T&D (.CRXPC)
	18-35	*Reconfiguration Flag*
5-10	0-35	Processor assignments (.CRCMC)
11	0-17	MBA of SSA cache (.CRMBA) (zero if non-extended memory)
	18-35	Not used
12	0-17	LAL of SSA cache (.CRMDB)
	18-29	Size of SSA cache (# of buffers)
	30-35	Not used
13*	0-8	Option word in .MFSIO
	9-17	Number of 320-word buffers for catalog cache
	18-35	Not used
14*	0-17	MBA of catalog cache (entry point -9 of .MFSIO) (zero if nonextended memory)
	18-35	Absolute address of catalog cache

15	0	Flag (=1 - Record written after 100 .EXIT traces processed without match on triggering .CALL trace)
	1-17	Code for event triggering this record 1 Initialization record 2 5 minute record 10 Change in .CRNPC or .CRXPC (Not monitored in WW6.4)* 11 Change in SSA cache (.CRMBA or .CRMDB) If a RLSEC request references the SSA cache buffer, then two reconfiguration records may be generated: one for deallocation of SSA cache (.CRMDB) and a second for the actual memory release (update of .CRPMU memory map). A further reconfiguration record will be generated if the SSA cache buffer is relocated. 12 Change in processor assignments (.CRCMC) 13 Change in FMS options or catalog cache (Not monitored in WW6.4)*
	18-35	Number of words following for channel information, one word per configured channel (.CRCT1)
16+	0-35	.CRCT1
17+	0-35	Memory map (.CRPMU)
		The length of this part of the special record is equal to the length contained in the RCW minus 14 minus the contents of bits 18-35 of word 15.

*This data is reported as zeros in WW6.4.

When GMC terminates, it writes a record containing information read from communication region cells and information about the amount of processor time used by the GMC collector in processing traces and also by the slave portion of GMC in performing its initialization and tape I/O. The format of the complete termination record is as follows.

Termination Record

Word	Bits	Infomation
1	0-17	Size of record (=23)
	18-35	Octal 300100

2	0-35	Memory size .CRMSZ
3	0-17	Number of IOM's (.CRNIC)
	18-35	Not used
4	0-17	Number of processors (.CRNPC)
	18-35	Not used
5	0-35	Current date (.CRDAT)
6	0-35	Current time (.CRTOD)
7	0-35	System configuration .CRFIG
8	0-35	Monitor slave processor time
9	0-35	A register from RSCR 32
10	0-35	Q register from RSCR 32
11	0-35	Idle time processor 0
12	0-35	Overhead time for processor 0 (words 11 and 12 are repeated for each configured processor)
13+	0-35	Number of times through GMC
14+	0-35	CPU time in GMC executive
15-25+	0-35	CPU time in MUM, MSM, CPUM, TAPE, CM, CAM, GRTM, TPE, IDLEM, NAME, FMS GEFSY
26+		Not Used

5.4.2 MUM. The MUM processes GCOS system trace types 10, 11, 46 and 51 by creating its own data collection records to describe the effect of these system events.

5.4.2.1 Trace Type 10. This MUM generated record contains sufficient information to describe the status of every job known to the Peripheral Allocator and Core Allocator. Each job is described in program number sequence. The format of a trace type 10 GMC data record is shown below. The data in this record is obtained from the occurrence of a system trace 10, 11 or 51 event.

Word	Bits	Information
1	0-17	Number of words following (record control word)
	18-26	Not used
	27-35	Octal 10 (trace number)
2	0-17	# (N) of 3 word entries in PALC data - if both bits 24 and 25 are set to 1
	18-23	not used
	24,25	If both set to 1, PALC record exists as described below - else PALC data not present
	26-35	not used

3-(3*N+2)	0-35	SNUMB
	0-11	Program status
	12-17	Program #
	18-26	Memory needed
	27-28	Not used
	29-35	Activity #
	0-5	Device type needed
	6	0 = fixed device 1 = removable device
	7-17	# of devices needed
	18-35	Not used

The following occur after the PALC data or after word 1 if there is no PALC data (bits 24 or 25 set to 0). This information is variable with 1 or more words per program. When no PALC data exists, word 2 as described above does not pertain.

One word of data for any activity not in memory:

0	Flag (=1 - activity is terminating)
1-8	Activity number from .CRSN1
9-17	Demand in 512-word blocks
18-23	Memory allocator urgency
24	Flag (=1 - job is waiting memory)
25	Must be zero
26	Privity indicator. Assume system program if program number is greater than 7, activity number is zero, and privity bit is set
27	Flag (=1 - if new job or new activity) (copy of bit 29, .CRLAL)
28-29	Must be zero
30-35	Program number

Multiple-word entries for any activity in memory where the first word contains:

0	Flag (=1 - activity is terminating)
1-8	GEMORE demand in 512-word blocks
9-17	Number of 512-word blocks allocated
18-23	Number of SSA's
24	Must be zero
25-27	Job status
28-29	Not used
30-35	Program number

The following information is dependent on job status (bits 25-27 of word 1)

Job Status	Information Collected
0	No data follows
1	New memory address
2	Snumb and activity
3	Snumb, activity, new address
4	10 ident words, 2 userid words
5	Ident, userid, new address
6	Ident, userid,snumb, activity
7	Ident, userid, snumb, activity, address

The Memory address word contains the following:

0-17	MBA in 512 word blocks
18-26	LAL in 512 word blocks
27-35	Not used

The Activity word contains the following:

0-8	Not used
9-17	Activity number
18-35	Not used

All multiple word entries contain the following three words:

Current CPU time
Current I/O Time
Memory Use Word

The Memory Use Word contains the following:

0-17	Memory used
18-29	Termination Code
30-35	Job Urgency from .SURG of a job

At the end of all T10 records the following two words will appear:

0-35	Number of 512-word blocks Available (word 0, .CRPMU table)
0-35	RSCR Time

If a RLSEC request is delayed (e.g. a request to release memory in which TSS is loaded must wait until TSS is swapped), then the number of blocks available may include nonallocatable memory (that portion of the RLSEC beyond the size of TSS). A delayed RLSEC request may be

indicated by consecutive MUM records which show no changes in memory(the MUM writes its record every time that .MPC04 is called; in a delayed RLSEC request, .MPC04 takes its denial exit without changing the memory state).

5.4.2.2 Trace Type 46. This GCOS system trace is generated every time a job is given a program number and results in a trace type 46 MUM record being generated. The format of the GMC trace type 46 record is shown below.

Word	Bits	Information
1	0-17	Size of record (=3)
	18-26	Not used
	27-35	Octal 46 (trace number)
2	0-29	SNUMB
	30-35	Octal 46
3	0-11	Not used
	12-17	Program number
	18-35	Not used
4	0-35	Time stamp

5.4.3 MSM. The MSM processes GCOS system trace types 7 and 15 by creating its own data collection records to describe the effect of these events. It also processes the specially-generated GMC trace events 73 and 76.

5.4.3.1 Trace Type 7. This GCOS system trace is generated every time a connect is issued and will result in the generation of a GMC trace type 7 record. The format of the GMC trace type 7 record is shown below.

Word	Bits	Information
1	0-17	Size of record (normal = 14, special = 21)
	18-23	Special file code description flags
	18	Permanent (=0), temporary (=1)
	19	Random (=0), sequential (=1)
	20	Not catalogued (=0), catalogued (=1)
	21	Removable (=0), fixed (=1)
	22	Flag (=1 - TSS user file (no file code))
	23	Flag (=1 - SCF File I/O (trace 22 is missing))
	24-26	Processor #
	27-35	Octal 7 (trace number)
2	0-17	I/O's word count
	18-23	Program number
	24-35	File code
3	0-35	System controller time of day

4	0-29	CP time usage
	30-35	IOM command
5	0-5	Not used
	6-35	SNUMB
6	0-5	Device command Special note: A command of octal 72 to a permanent disk pack indicates that a pack exchange is in progress. The .MGP66 module issues another standby command to the device to which the permanent pack is to be moved. A special device name record should appear either in the current block on tape or at the beginning of the next block to confirm the pack exchange.
	6-17	DCW length
	18-35	File origin block number
7	0-13	File size
	14	Sysout flag
	15	Seek flag
	16-35	Seek address
8	0-5	Device command
	6-11	Device number
	12-17	IOM PUB number
	18-23	I/O Command
	24-26	Not used
	27-29	IOM number
	30-35	Record count
9	0-17	MBA of job issuing connect or zero if nonextended memory
	18-35	I/O queue address in SSA (absolute address if nonextended memory or if value less than 64K; relative address (to MBA) if extended memory and if value greater then 64K)
10	0-17	Activity number
	18	Flag (=1 - I/O status is stopped)
	19-28	Not used
	29-34	I/O status (I/O queue word 0)
	35	Flag (=1 - system job)
11	0-35	.CRCOM (use only bits 18-29)
12	0-35	.CRCOM+2 (use only bits 18-29)
13	0-35	.SGCPT (use only bits 18-35)
14	0-35	.SRQCT (use only bits 0-17)
15	0-35	.SNIO (use only bits 18-35)
(words 11-15 are not generated if the IDENT for a job is reported)		
11-20		IDENT (words 11-14 described above are not collected)
21-22		USERID

5.4.3.2 MSM Special Record. During the execution of MSM a special record is written at preselected times during the monitoring session. These records are used to analyze SSA cache core, when configured. The format of this record is shown below. This record is based on data collected during the processing of GMC trace events 73 or 76.

Word	Bits	Information
1	0-17	Size of record (=517)
	18-26	Not used
	27-35	Octal 7 (trace number)
2-516	0-35	a. Number of times each GCOS module 1-515 was in the SSA cache buffer b. Number of times each GCOS module 1-515 was loaded by an I/O because it was not in the SSA cache buffer.
517	0-35	Not used
518	0-35	Flag (=2 - case a. above) (=3 - case b. above)

5.4.3.3 Device Name Record. If either MSM or CM is active, the GMC writes a record which correlates device names to device addresses. The System Configuration Name table is processed sequentially as this record is formatted. Names for all disk devices are reported. In order to detect exchanges of mass storage devices, GMC periodically examines the device name table. If any changes have occurred, then another device name record is written. This record is variable in size, and recognized by the special format of the second word.

Word	Bits	Information
1	0-17	Size of record
	18-26	Not used
	27-35	Octal 7 (trace number)
2	0-35	Octal 535353535353
3-n	0	Flag (=1 - fixed device if mass storage) (set if bits 13 and 14 of the second word of the SCT entry for any mass storage device are both zero; in Shared Mass Storage environment, shared devices must be fixed)
	1-5	IOM number
	6-11	Channel number
	12-17	Device number
	18-35	Device name in BCD

5.4.3.4 FILSYS Catalog Structure Record. During the execution of the Mass Store Monitor, certain MME GEFSYEs GCOS Trace 15 data are collected concerning the catalog file string that is being referenced. The purpose of this data is to try and determine how many connects are being made because of the particular structure of a given catalog or file. This data is also used to provide the catalog file string name associated with the various user file codes that are reported by the Mass Store Monitor. MME GEFSYE traces are only processed if generated by a system program (program number less then 15 or FSTSLV). In addition, only the following GEFSYE's will be

processed: 2, 3, 4, 5, 8, 9, 10, 11, 18, 19, 20, 22, 23, 27, 28, and 29. The format of the GMC record generated is as follows:

Word	Bits	Information
1	0-17	Size of record (variable)
	18-26	Not used
	27-35	Octal 15 (trace number)
2	0-35	-1 if program generating this record is not PALC. Otherwise this word contains a file code in bits 18-35.
3	0-35	-1 if program generating this record is not PALC. Otherwise it contains the SNUMB of the job for which PALC is working.
4	0-17	GEFSYE type
	18-20	Processor #
	21-26	Program #
	27-35	Activity #
5-n	0-35	Catalog file string name - not to exceed 40 words.

In order to monitor the catalog file string names of user files that are being processed by the File Transfer System (FTS), it is necessary to capture the occurrence of a MME GEMORE when issued by FTS. All MME GEFSYEs issued by FTS are ignored so as not to cause a conflict with the MME GEMOREs. An FTS GEMORE is processed only if it is a GEMORE for a permanent file. The record generated is identical to the MME GEFSYE record, except that word 2 is a minus one (-1) and word 3 is the file code being used by FTS.

5.4.3.5 FMS Cache Record. During the execution of MSM or CM a special record is written at preselected times during the monitoring session. These records are used to analyze FMS catalog cache (when configured) and also the effectiveness of the incore space tables for disk devices. This record is not generated on a WW6.4 system. The format of this GMC record is shown below.

Word	Bits	Information
1	0-17	Size of record (=15)
	18-26	Not used
	27-35	Octal 77 (special flag)
2	0-35	Number of cache hits (word -12 from entry point of .MFSIO)
3	0-35	Number of writes (word -11 from entry point of .MFSIO)
4	0-35	Number of reads (word-10 from

		entry point of .MFSIO)
5	0-35	Number of reads not in CC (octal 1511 from entry point of .MFSIO)
6	0-35	Number of non-320 word reads (octal 1512 from entry point of .MFSIO)
7	0-35	Number of skips caused by .SSTAK (octal 1513 from entry point of .MFSIO)
8	0-35	Number of cache clears (octal 1514 from entry point of .MFSIO)
9	0-35	Number of no hits (octal 1520 from entry point of .MFSIO)
10	0-35	Number of hits (octal 1521 from entry point of .MFSIO)
11	0-35	.CRSHR
12	0-35	Number of times buffer allocation attempted (word -6 from entry point of .MASO4)
13	0-35	Number of times buffer busy (word -5 from entry point of .MASO4)
14	0-35	Number of times available space table was already in memory (word -4 from entry point of .MASO4)
15	0-35	Number of times available space table was in memory but was busy (word -3 from entry point of .MASO4)

5.4.4 CPUM. The CPU Monitor processes the GMC generated event trace type 70.

5.4.4.1 Trace Type 70 - Standard. This GMC record allows six processors to be monitored and allows differentiation between TS1 executive processor time and TS1 subdispatch processor time. An activity is recognized as a copy of TSS if bit 13 in its .STAT1 word is set and if its SNUMB is TS2, TS3, or TS4 (TS1 always has program number 5). The check on the .STAT1 word eliminates possible confusion between legitimate copies of TSS and GEIN execution of spawn files or termination of DRL TASK jobs by the same names. If a system program has a SNUMB of $PACT, $MOLT, $POLT, $COLT, $SOLT, or $SALT its processor time is accumulated, along with that for program number six (test and diagnostics). The format for this GMC trace type record is shown below.

Word	Bits	Information
1	0-17	Size of record (79 or 173)
	18-26	Not used
	27-35	Trace number (octal 70)
2	0-35	Not used
3	0-35	Processor time (clock pulses) for program 1 ($CALC, Core Allocator)

4	0-35	Processor time (clock pulses) for program 2 ($PALC, Peripheral Allocator)
5	0-35	Processor time (clock pulses) for program 3 ($SYOT, SYSOUT writer)
6	0-35	Processor time (clock pulses) for program 4 ($RTLN, scheduler)
7	0-35	Processor time (clock pulses) for program 5 (TS1, TSS Executive)
8	0-35	Processor time (clock pulses) for program 6 ($TOLT, T&D executive; also includes time for special T&D SNUMBs)
9-16	0-35	Processor time (clock pulses) for programs 7-14 (decimal) (Transaction Processor, Log-on, TTLSYS protection, WIN, DMTEX). In commercial releases, several of these words will contain no data since many of these programs are strictly WWMCCS related.
17	0-35	Processor time (clock pulses) for GMC
18	0-35	Processor time (clock pulses) for user programs
19	0-35	Subdispatch time (clock pulses) for program 5 (TS1)
20	0-35	Processor time (clock pulses) for TS2 executive
21	0-35	Processor time (clock pulses) for TS3 executive
22	0-35	Processor time (clock pulses) for TS4 executive
23	0-35	Subdispatch time (clock pulses) for TS2
24	0-35	Subdispatch time (clock pulses) for TS3
25	0-35	Subdispatch time (clock pulses) for TS4
26	0-35	Miscellaneous subdispatch time (clock pulses) (expansion capability for TDS, TPE II)
27-51	0-35	Number of dispatches to those programs associated with words 2-26
52	0-35	RSCR time
53-58	0-35	Idle time for processors 0-5 (.CRIDT)
59-64	0-35	Overhead time for processors 0-5 (.CROVH)
65-69	0-35	Processor time (clock pulses) for 5 specially requested SNUMBs
70-74	0-35	Name of special SNUMB in BCD
75-80	0-35	Gate loop time for each processor. Module .MFALT deducts gate loop time from the processor time charged to jobs and adds it to overhead time reported in .CROVH.

(If the user disables the monitoring of dispatcher queuing, this is the end of the CPU record. If dispatcher queuing is left enabled, then the record will contain the following additional data:)

81-144	0-35	The dispatcher queue table
145-169	0-35	Dispatcher queue time (clock pulses) for those programs associated with words 2-26. It should be noted that words 162, 166-169 will contain no data since subdispatch queuing is currently not monitored.
170-174	0-35	Dispatcher queue time (clock pulses) for 5 specially requested SNUMEs

5.4.4.2 Trace Type 63 - Initial. The CPU Monitor generates an initial record which describes the dispatcher options that are active.

Word	Bits	Information
1	0-17	Record size (8)
	18-26	Not used
	27-35	Octal 63 (trace number)
2-5	0-35	First four words of .MDISP
6	0-35	.CRPJT+2
7-9	0-35	Not used

5.4.4.3 Trace Type 63 - Activity Termination. Whenever an activity terminates, a record is generated describing the dispatcher queue time accumulated by the activity. This record is only generated when the dispatcher queuing option is left enabled.

Word	Bits	Information
1	0-17	Record size (8)
	18-26	Not used
	27-35	Octal 63 (trace number)
2	0-35	SNUMB
3	0-35	Queue time in clock pulses
4	0-35	CPU time in clock pulses
5	0-17	Program number
	18-35	Activity number
6	0-35	RSCR stop time of job
7	0-35	Start time of job taken from .CRTOD
8	0-35	Accumulated swap time of job in clock pulses
9	0-35	Stop time of job taken from .CRTOD

5.4.4.4 Trace Type 63 - Termination Record. When the SMF is terminated and the dispatcher queuing option is enabled, the following termination record is generated.

Word	Bits	Information
1	0-17	Record size (386)
	18-26	Not used
	27-35	Octal 63 (trace number)
2-65	0-35	Queue time for all currently active programs in clock pulses
66-129	0-35	CPU time for all currently active programs in clock pulses
130-193	0-35	SNUMB for all currently active programs
194-257	0-35	Start time for all currently active programs taken from .CRTOD
258-321	0-35	Activity number for all currently active programs
322-385	0-35	Cumulative swap time for all currently active programs in clock pulses
386	0-35	RSCR clock time
387	0-35	.CRTOD clock time

5.4.5 TM. The Tape Monitor processes three GCOS system traces: 50, 51, and 52 and creates its own data collection records to evaluate the effect of these events.

5.4.5.1 Trace Type 50. This GCOS system trace is generated whenever an activity goes to the core allocator and will result in the generation of a GMC trace type 50 record. The format for this GMC trace type record is shown below.

Word	Bits	Information
1	0-17	Size of record (variable)
	18-26	Not used
	27-35	Octal 50 (trace number)
2	0-29	SNUMB
	30-35	Not used
3	0-35	Time stamp
4	0-11	Activity number
	12-17	Program number
	18-29	Urgency
	30-35	Octal 50
5-N	0-35	Words 1 and 2 of SCT entry for each tape used by this program

5.4.5.2 Special Trace Type 50. The first GMC trace type 50 record is a special trace 50 to indicate the status of all tape drives when the monitor first begins. Its structure is shown below.

Word	Bits	Information
1	0-17	Size of record (variable)
	18-26	Not used
	27-35	Octal 50 (trace number)

4	0-35	Processor time (clock pulses) for program 2 ($PALC, Peripheral Allocator)
5	0-35	Processor time (clock pulses) for program 3 ($SYOT, SYSOUT writer)
6	0-35	Processor time (clock pulses) for program 4 ($RTIN, scheduler)
7	0-35	Processor time (clock pulses) for program 5 (TS1, TSS Executive)
8	0-35	Processor time (clock pulses) for program 6 ($TOLT, T&D executive; also includes time for special T&D SNUMBs)
9-16	0-35	Processor time (clock pulses) for programs 7-14 (decimal) (Transaction Processor, Log-on, FILSYS protection, WIN, DMTEX). In commercial releases, several of these words will contain no data since many of these programs are strictly WWMCCS related.
17	0-35	Processor time (clock pulses) for GMC
18	0-35	Processor time (clock pulses) for user programs
19	0-35	Subdispatch time (clock pulses) for program 5 (TS1)
20	0-35	Processor time (clock pulses) for TS2 executive
21	0-35	Processor time (clock pulses) for TS3 executive
22	0-35	Processor time (clock pulses) for TS4 executive
23	0-35	Subdispatch time (clock pulses) for TS2
24	0-35	Subdispatch time (clock pulses) for TS3
25	0-35	Subdispatch time (clock pulses) for TS4
26	0-35	Miscellaneous subdispatch time (clock pulses) (expansion capability for TDS, TPE II)
27-51	0-35	Number of dispatches to those programs associated with words 2-26
52	0-35	RSCR time
53-58	0-35	Idle time for processors 0-5 (.CRIDT)
59-64	0-35	Overhead time for processors 0-5 (.CROVH)
65-69	0-35	Processor time (clock pulses) for 5 specially requested SNUMBs
70-74	0-35	Name of special SNUMB in BCD
75-80	0-35	Gate loop time for each processor. Module .MFALT deducts gate loop time from the processor time charged to jobs and adds it to overhead time reported in .CROVH.

(If the user disables the monitoring of dispatcher queuing, this is the end of the CPU record. If dispatcher queuing is left enabled, then the record will contain the following additional data:)

81-144	0-35	The dispatcher queue table
145-169	0-35	Dispatcher queue time (microseconds) for those programs associated with words 2-26. It should be noted that words 162, 166-169 will contain no data since subdispatch queuing is currently not monitored.
170-174	0-35	Dispatcher queue time (microseconds) for 5 specially requested SNUMBs

5.4.4.2 Trace Type 63 - Initial. The CPU Monitor generates an initial record which describes the dispatcher options that are active.

Word	Bits	Information
1	0-17	Record size (8)
	18-26	Not used
	27-35	Octal 63 (trace number)
2-5	0-35	First four words of .MDISP
6	0-35	.CRPJT+2
7-9	0-35	Not used

5.4.4.3 Trace Type 63 - Activity Termination. Whenever an activity terminates, a record is generated describing the dispatcher queue time accumulated by the activity. This record is only generated when the dispatcher queuing option is left enabled.

Word	Bits	Information
1	0-17	Record size (9)
	18-26	Not used
	27-35	Octal 63 (trace number)
2	0-35	SNUMB
3	0-35	Queue time in microseconds
4	0-35	CPU time in clock pulses as determined by system
5	0-17	Program number
	18-35	Activity number
6	0-35	RSCR stop time of job
7	0-35	Start time of job taken from .CRTOD
8	0-35	Accumulated swap time of job in clock pulses
9	0-35	Stop time of job taken from .CRTOD
10	0-35	CPU time in microseconds as computed from traces

5.4.4.4 Trace Type 63 - Termination Record. When the GMF is terminated and the dispatcher queuing option is enabled, the following termination record is generated.

Word	Bits	Information
1	0-17	Record size (450)
	18-26	Not used
	27-35	Octal 63 (trace number)
2-65	0-35	Queue time for all currently active programs in microseconds
66-129	0-35	CPU time for all currently active programs in clock pulses as determined by system
130-193	0-35	SNUMB for all currently active programs
194-257	0-35	Start time for all currently active programs taken from .CRTOD
258-321	0-35	Activity number for all currently active programs
322-385	0-35	Cumulative swap time for all currently active programs in clock pulses
386-449	0-35	CPU time for all currently active programs in microseconds as determined from traces
386	0-35	RSCR clock time
387	0-35	.CRTOD clock time

5.4.5 TM. The Tape Monitor processes three GCOS system traces: 50, 51, and 52 and creates its own data collection records to evaluate the effect of these events.

5.4.5.1 Trace Type 50. This GCOS system trace is generated whenever an activity goes to the core allocator and will result in the generation of a GMC trace type 50 record. The format for this GMC trace type record is shown below.

Word	Bits	Information
1	0-17	Size of record (variable)
	18-26	Not used
	27-35	Octal 50 (trace number)
2	0-29	SNUMB
	30-35	Not used
3	0-35	Time stamp
4	0-11	Activity number
	12-17	Program number
	18-29	Urgency
	30-35	Octal 50
5-N	0-35	Words 1 and 2 of SCT entry for each tape used by this program

5.4.5.2 Special Trace Type 50. The first GMC trace type 50 record is a special trace 50 to indicate the status of all tape drives when the monitor first begins. Its structure is shown below.

Word	Bits	Information
1	0-17	Size of record (variable)
	18-26	Not used
	27-35	Octal 50 (trace number)

2	0-35	Not used
3	0-35	Time stamp
4	0-35	Flag (-2)
5	0-35	SNUMB
6	0-35	Word 0 of SCT table
7	0-35	.CRSN1 for SNUMB
8	0-35	Word 3 of SCT table
9-N	0-35	Repeat of words 5-8 for as many tape drives as are currently configured

5.4.5.3 Trace Type 51. This GCOS system trace is generated every time an activity terminates and will result in the generation of a GMC trace type 51 record. This record structure is identical to the type 50 record structure except for word 4. The format for this word is shown below.

Word	Bits	Information
4	0-11	Abort code
	12-17	Program number
	18-26	Urgency
	27-29	Type of abort
	30-35	Octal 51

5.4.5.4 Trace Type 52. This GCOS system trace is generated every time a program is denied a tape due to unavailability and will result in the generation of a GMC trace type 52 record. The format for this GMC trace type 52 record is shown below.

Word	Bits	Information
1	0-17	Record size (variable)
	18-26	Not used
	27-35	Octal 52 (trace number)
2	0-35	Not used
3	0-35	Time stamp
4	0-35	Flag (-2)
5	0-35	Register A of standard type 52 trace
6	0-35	Register Q of standard type trace
7-N	0-35	Same as special trace 50, words 5-n

5.4.6 CM. The Channel Monitor processes four different GCOS system trace types: 4, 7 and 22. These events cause the generation of the following GMC trace records.

5.4.6.1 Trace Type 4. This GCOS system trace is generated at the termination of an I/O and will cause the generation of a trace type 4 GMC record. The format of the GMC trace type 4 record is shown below.

Word	Bits	Information
1	0-17	Size of record (=3)
	18-26	Not used
	27-35	Octal 4 (Trace number)
2	0-35	Time stamp
3	0-29	A-register of trace
	30-35	IOM number
4	0-17	MBA of job
	18-35	I/O queue address

5.4.6.2 Trace Type 7. This GMC trace record is the same as the MSM trace type 7 record described in subsection 5.4.4.1 and 5.4.4.3. Only one trace type 7 is generated if both the CM and MSM are active.

5.4.6.3 Trace Type 22. This GCOS system trace is generated every time an I/O is requested and will result in the generation of a GMC trace type 22 record. The format of this GMC trace type 22 record is shown below.

Word	Bits	Information
1	0-17	Size of record (=3)
	18	Device busy flag (set if device is busy)
	19	Abort flag (set if bit 5 or bit 6 in .STATE is set)
	20	Swap/move flag (set if bit 7 or bit 8 in .STATE is set)
	21	Flag (=1 - I/O status is stopped)
	22-26	I/O status (bits 31-35 of I/O queue word 0)
	27-35	Octal 22 (trace number)
2	0-35	Time stamp
3	0-35	Device type
	6-11	Device number
	12-17	Channel number
	18-23	Program number
	24-33	Logical primary channel index
	34-35	IOM number
4	0-17	MBA of job
	18-35	I/O queue address

5.4.7 CAM. The Communications Analysis Monitor processes the GMC generated event trace type 14.

5.4.7.1 Trace Type 14. This GMC trace is generated whenever the H6000-355 mailbox in DNWW/MDNET is changed and will result in the

writing of a GMC trace type 14 record. The format for this GMC trace type 14 record is shown below.

Word	Bits	Information
1	0-17	Record size (=variable)
	18-26	Not used
	27-35	Octal 14 (trace number)
2	0-35	Time stamp
3	0-2	355 number
	3-17	Logical line number
	18-35	Terminal ID
4	0-8	Terminal type
	9-17	ICM count
	18-26	OP code
	27-35	Command
5-7	0-23	ICM Data
8	0-23	Not used
	24-35	Data tally
9	0-11	Ignore
	12-17	Status
	18-29	Input data tally
	30-35	Not used
10	0-17	Slave LAL
	18-35	Checksum
11-12	0-17	Number datanets
	18-35	Not used
13-n	0-35	A series of values will follow depending upon the number of datanets configured. The values collected are the number of special interrupts processed, number of special interrupts unanswered, number of requests waiting mailbox. For each type of entry, a single value will appear for each datanet configured. This set of numbers will then be followed by the second table type and finally the third table type. This series of tables is next followed by a fourth table containing the number of lines waiting mailbox over each datanet.
n+1	0-35	Userid or -1 if not logged on to TSS
n+2	0-35	Userid (present only if user logged onto TSS)
n+3	0-17	Not used (present only if user logged onto TSS)
	18-35	TSS flags (.LFLAG) (present only if user logged onto TSS)
n+4	0-17	Not used (present only if user logged onto TSS)
	18-35	Subsystem size (words) (.LSIZE) (present only if user logged onto TSS)
n+5-529	0-35	Communications traffic (only if specific CAM option specified)

5.4.7.2 Special Trace Type 14. This trace record is written during CAM initialization to specify the start time. Its structure is shown below.

Word	Bits	Information
1	0-35	Octal 7000014
2	0-35	Time from MME GETIME
3-8	0-35	Not used

5.4.7.3 Special TSS Trace Type 14. This trace record is written a single time and provides a description of the Timesharing Environment as it exists during this data collection period.

Word	Bits	Information
1	0-17	Record size (=30)
	18-26	Not used
	27-35	Octal 14 (trace number)
2	0-35	special flag (= -2)
3	0-35	.TFMAX - maximum number terminals
4	0-35	.TAMRI - time interval for memory size reduction
5	0-35	.TATMC - maximum time allowed for size changes
6	0-35	.TAGMI - minimum time between GEMORE requests
7	0-35	.TAMMS - initial maximum TSS size
8	0-35	.TASMS - minimum TSS size
9	0-35	.TAMII - memory size growth factor
10	0-35	.TASRI - memory size reduction factor
11	0-35	.TSFS - minimum size reduction factor
12	0-35	.TSGRW - minimum swap file size
13	0-35	.TSSF - number of swap files
14	0-35	.TCSF - swap file #S
15	0-35	.TCSF+1 - swap file #T
16	0-35	.TCSF+2 - swap file #U
17	0-35	.TCSF+3 - swap file #V
18	0-35	.TIMER - reconnect timer
19	0-35	.TAMIS - large subsystem fence size
20	0-35	.TALPP - large subsystem wait time
21	0-17	.TAGPP - #32 ms time quantums
	18-35	.TAGPP - frequency of Priority B dispatching with a 1 meaning every other dispatch and a 2 meaning every third dispatch
22	0-35	.TTASK - maximum number of concurrent derail tasks
23-31	0-35	.TSFDV-.TSFDV+8 - device allocation for TSS files #D, #P, #Q, .D, SS, #S, #T, #U and #V

5.4.7.4 Special GEROUT Trace Type 14. This trace record is written whenever a user transfers from one major subsystem to another (Example: TSI to FTS).

Word	Bits	Information
1	0-17	Record size (=7)
	18-26	Not used
	27-35	Octal 14 (trace number)
2	0-35	-3 (flag)
3	0-17	0
	18-23	GEROUT type
	24-35	Terminal ID
4	0-35	New DAC name
5	0-35	Time
6-7	0-35	Not used

5.4.8 GRTM. The GRTS Monitor processes one GMC trace record, type 62.

5.4.8.1 Trace Type 62. This GMC trace is generated whenever the DATANET-355 GRTS Monitor data record is transmitted to the GMC. The format for this GMC trace type 62 record is shown below.

Word	Bits	Information
1	0-17	Record size (variable)
	18-20	DATANET number
	21-26	Not used
	27-35	Octal 62 (trace number)

Word	Bits	Information
2	0-35	Time stamp - H6000
3	0-17	DAC character count
	18-35	Time stamp - DN-355
4	0-17	010101
	18-35	Buffer denials (cumulative)
5	0-17	010201
	18-35	Buffer availability (current)
6	0-17	010301
	18-35	Number of users (current)
7	0-17	010401
	18-35	Number of transactions sent to host (cumulative)
8	0-17	010501
	18-35	Number of transactions received from host (cumulative)
9	0-17	010601
	18-35	Number of 36-bit words sent to host (cumulative)
10	0-17	010701
	18-35	Number of 36-bit words received from the host (cumulative)
11	0-17	011001
	18-35	Number of host RSVPs received (cumulative)
12	0-17	011101
	18-35	Amount of time in milliseconds spent in idle loop since the last buffer was sent
13	0-17	011201
	18-35	Number of calls to the buffer allocation routine (cumulative)
14	0-17	030105
	18-26	HSLA
	27-35	HSLA subchannel
15	0-17	Number of transmits on S/C (cumulative)
	18-35	Number of receives on S/C (cumulative
16-N		Additional entries depending on the number of subchannels specified on the data card.
N+1		Response Time Buffer. This portion of data is variable depending upon the activity occurring on the DN-355. The various types of data that can be collected in this buffer are illustrated below.

Every line displayed require 18 bits of buffer space.

0000100000001000010 Terminal Type / Line ID Time Stamp Direct Access Name (1/2)	Connect-To-Slave Record
0000100000010000010 Terminal Type / Line ID Time Stamp	Send Output Record
0000100000011000010 Terminal Type / Line ID Time Stamp	Accept Direct Output Record
0000100000100000010 Terminal Type / Line ID Time Stamp	Accept Direct Outpu' nd Then Input Record
0000100000101000010 Terminal Type / Line ID Time Stamp	Accept Direct Input Record
0000100000110000010 Terminal Type / Line ID Time Stamp	Input Accepted Record
0000100000111000010 Terminal Type / Line ID Time Stamp	Disconnect Record
0000100001000000010 Terminal Type / Line ID Time Stamp	Output Received Record
0000100001001000010 Terminal Type / Line ID Time Stamp	Output Started Record
0000100001010000010 Terminal Type / Line ID Time Stamp	Output Complete Record
0000100001011000010 Lost Message Count	Lost Message Record (Always first data in type 2)

The following definitions apply to the above record types.

- o Connect-To-Slave A terminal has logged on to an H6000 DAC program.
- o Send Output DATANET-355 requesting more output from the H6000.
- o Accept Direct Output The H6000 has told the DATANET-355 that it has data to send.
- o Accept Direct Output and Then Input The H6000 has told the DATANET-355 that it has data to send and expects input to be returned.
- o Accept Direct Input The DATANET-355 asking the H6000 to accept data.
- o Input Accepted Input received by the H6000.
- o Disconnect A terminal has disconnected from a H6000 DAC program.
- o Output Received DATANET-355 has received data from the H6000.
- o Output Started DATANET-355 has started to transmit data to a terminal.
- o Output Complete DATANET-355 has finished sending data to the terminal.

5.4.9 IDLEM. The Idle Monitor processes a GCOS system trace type 21 by creating its own data collection record to describe the effect of this event.

5.4.9.1 Trace Type 21. The format of the GMC trace type 21 record is shown below. The data in this record is obtained from the occurrence of many different system trace events (see table 5-1).

Word	Bits	Information
1	0-17	Size of record (=3)
	18-20	Processor number
	21-26	Not used
	27-35	Octal 21 (trace number)
2	0-35	Time idle was broken
3	0-35	Time idle started
4	0-35	Not used

5.4.10 TPS. The TPSM processes GCOS system trace types 0, 1, 2, 4, 5, 6, 13, 42, 51, 65 and 74 by creating its own data collection records to describe the effect of these system events.

5.4.10.1 Trace Types 0, 1, 2, 4, 5, 6 and 65. These GCOS system traces are generated during job I/O and execution and result in the generation of a GMC trace type 74 record. The format for this GMC trace type record is shown below.

Word	Bits	Information
1	0-17	Number of words following (record control word)
	18-26	Not used
	27-35	Octal 74 (trace number)
2	0-35	Register A of standard GCOS trace
3	0-35	Register Q of standard GCOS trace
4	0-35	Time Stamp

5.4.10.2 Trace Types 13, 42 and 51. These GCOS system traces are generated for each job start, process, and termination. The TPSM processes these traces for each TPAP and generates the following GMC trace type 74.

Word	Bits	Information
1	0-17	Number of words following (record control word)
	18-26	Not used
	27-35	Octal 74 (trace number)
2	0-35	Register A of standard GCOS trace
3	0-35	TPAP SNUMB
4	0-35	Time Stamp

5.4.10.3 Trace Type 74. This internally generated TPE trace has a variable format depending upon where within TPE the trace was generated.

5.4.11 TSS Monitor. The TSS Monitor (TSSM) collector generates traces by using the GCOS trace 74 mechanism. Since both TPSM and TSSM use the same GCOS trace (74), they are incompatible.

5.4.11.1 TSSM Trace Type 74. Formats of the TSSM trace 74 vary according to the subtype; there are 105 subtypes.

Except where noted, all traces written by TSSM are four words long. The first two words are always in the following format:

Word 1	Bits	0-17	Number of words following (normally 3)
		18-20	Processor number
		21	Flag (=1, TSS in courtesy call)
		22-35	GMC record type (=74 octal)
2		0-35	RSCR time

TSSM normally stores the A-register of GCOS trace 74 into word 3 and the Q-register of GCOS trace 74 into word 4. There are three exceptions to this rule. For the I/O request trace (24), TSSM adjusts the I/O queue address passed by TSS from a negative offset to a positive offset using the Lower Address Limit. The second exception is that the TSSM detected delay trace (105) is created during parallel CPU monitor processing so that if TSS receives no dispatches, a trace can be made anyway to give SSA data about TSS. The last exception is for the initialization trace (90). For this trace, receipt of GCOS trace 74 tells the TSSM capture routine to produce a series of special records. The only other trace subtype having a length other than four words is subtype 32. In that subtype, two words are appended onto the normal four-word format. Word 3 takes the following format in most traces:

Word 3	Bits	0-17	UST address
		18-20	Flags or program stack (usually zero)
		21-29	Subtype number
		30-35	GCOS trace number (=74 octal)

The above format will be called "Common format" in following trace descriptions when the flag field is zero.

Subtype 1, Log-on complete

Word 3	Common format		
4		0-23	Reserved (DRUN ID in trace 32)
		24-35	Line ID or blanks (BCD)

Subtype 2, Swap space denied

Word 3	Common format
4	Not used

Program stack - described by following scheme:

```
(UST+.LFILE) contains ZERO    TALLY,AFTPTR
                      TALLY   TALLY   *+n,6-n
                              DUP     1,5
                              ZERO    ASCII name, address of
                                      primitive
                      n=depth of stack (0-5, 0=no entries)
                      primitive=3 if build mode
```

Flag for build mode - set nonzero if a primitive 3 is found

Subsystem BAR and LAL - found at offset .LSIZE from current UST address (bits 0-17)

Time type - found at offset .LTCW from current UST address (bits 27-35)

Log-on time - found at offset .LTALC from current UST address

Extra buffer memory address - found at offset .LKMS from current UST address (bits 0-17)

BTOS flags - found at offset .LPQF from current UST address (bits 18-35). If bit 18 nonzero, store EBM address in bits 0-17; otherwise ignore .LKMS

Memory demand - found at offset .LSIZE from current UST address (bits 18-35)

Subsystem size for penalty - found at location .TAMIS in TSSA (bits 0-17)

Penalty factor - found at location .TALPP in TSSA (bits 27-35)

5.4.12 Special Records. During the execution of the GMC, it sometimes is necessary to generate special records that describe the occurrence of a special event. Following is a description of these special records.

5.4.12.1 Lost Data Record. If the rate of data collection does not allow GMC to dump its internal buffers to tape or if the system develops a tape malfunction, it is possible for GMC to generate a lost data condition. When this condition occurs, a special trace is

generated in the last good record recorded on tape. The next good trace recorded will be found at the beginning of the next physical record.

Word	Bits	Information
1	0-35	000777200100

5.4.12.2 Termination Record. Upon termination, GMC writes out a special termination record. The format is described in subsection 5.6.1.

5.4.12.3 End-of-Reel Flag. When the multireel option is enabled, GMC writes the following special record header at the end of each tape.

Word	Bits	Information
1	0-35	007773400100
2	0-35	Continuation reel number (BCD)

5.4.12.4 MUM Lost Data. If GMC generates a lost data condition while executing the MUM, the MUM generates a special flag in the header. The format of this flag is described below.

Word	Bits	Information
1	0-17	Record size
	18-35	Octal 200110 (lost data flag)

5.4.12.5 Reconfiguraton Record. The format of this record is described in subsection 5.4.1 under the reconfiguration discussion.

5.5 GMF User Input Parameter Options

The user must start here to define his intended use of the GMC so that he can then determine his JCL and system structure.

User control over GMC is total as a complete parameter capability is provided. Sample data parameter cards are shown in figure 5-2.

The GMC functional options are:

(1) Capability to turn off any particular monitor or combination of monitors.

(2) Specifying that collection is to halt after filling 1-9 data tapes. The default is to collect an unlimited number of tapes.

```
CC
1
M0 M3                    Turn off Monitor 0 and 3
M1                       Turn off Monitor 1
M1 M9                    Turn off Monitor 1 and collect only a single
                         reel
M1 *12.36,05.00          Turn off Monitor 1, start collecting data at
                         12.36, and collect for 5 hours
*,03.00                  All Monitors are present on the R* file and are
                         active, collection is to start at once and
                         continue for 3 hours
+CK                      All Monitors are present on the R* file, and
                         communication traffic is to be monitored for
                         terminal CK
M1 M4 M93                Turn off Monitors 1 and 4, and collect maximum
                         of three reels of data
M*                       Suppress abort if GMC cannot move
#VIDEO,HEALS             All Monitors are present on the R* file, and
                         accumulate processor time in the CPU Monitor
                         for these SNUMBs.
M0 M5 M8 ?1              Turn off monitors 0, 5, and 8.  Collect
                         only tape connects with the MSM and CM.
M0 M2 M5 X               Turn off monitors 0,2,5, read second data
MS2755T,RTOS             card, turn on MSM/CM special SNUMB option to
                         include TSS, FTS, $PALC, 2755T and RTOS.
M0 M4 MS                 Turn off monitors 0,4, and collect MSM/CM traces
                         for only TSS, FTS, $PALC.
```

Figure 5-2. Data Card Examples

(3) Requesting complete communication data for 1 or 2 terminal IDs.

(4) Suppressing a GMC abort if it cannot move to an acceptable location.

(5) Specifying up to three SNUMBS to be processed by the CPU Monitor.

(6) Requesting that only tape connects or mass storage connects be collected, but not both. The default is to collect both types.

(7) Declaring the start and stop times of monitoring.

(8) Special Debug Request.

(9) Specifying that the Mass Store Monitor and/or Channel Monitor are to collect data only for certain jobs.

(10) Specifying that the data card options are continuing on a new card. Without this option, only a single data card will be processed.

(11) Turn off the dispatcher queuing option for the CPU Monitor.

(12) Specifying the monitoring requirements for the GRTM.

All options are listed on a single data card, unless a dual data card is stated as being used. Each option should be separated from a previous option by the presence of a single blank column.

5.5.1 On/Off Option. This option allows the user to turn off all monitors not required for his purposes. Since the GMC default is to have all monitors turned on, unless specifically turned off, and since the TSS and TPE Monitors are incompatible, the user must have a data card and at least one of these two monitors must be turned off. The code format to turn off a given monitor is:

M0 = Memory Utilization
M1 = Mass Store Monitor
M2 = CPU Monitor
M3 = Tape Monitor
M4 = Channel Monitor
M5 = Communications Monitor
M6 = GRTS Monitor
M7 = TPE Monitor
M8 = Idle Monitor
MA = TSS Monitor
MB-MF = User Developed Monitors

(See section 13 for a discussion of user developed monitors.)

Multiple monitors are turned off by listing a series of the above monitor codes, with each code separated by a single blank column.

CAUTION: The TPE Monitor and the TSS Monitor are incompatible and cannot be active at the same time.

While it is optional to turn off a monitor, a user must turn off, on the parameter card, any monitor that is not loaded in the compiled R*. Failure to do so will result in an M0-M8,MA abort. The hexadecimal digit following the M represents the monitor that is not present on the R* file, but yet was not turned off on the parameter card. Details for creation of the R* file are given in subsection 5.6.

5.5.2 Tape Selection Option. This option allows the user to specify the number of data collector tapes to be accumulated in a job run.

M9 = One reel of tape
M91 = One reel of tape
M92 = Two reels of tape
M93 = Three reels of tape
M94 = Four reels of tape
.
.
.
M99 = Nine reels of tape

If more than nine reels of tape are desired, the user must collect an unlimited number of tapes until the GMC is aborted by either a time request or an operator console abort.

When this option is used, GMC will terminate normally with an SF abort.

5.5.3 Terminal Specification Option. This option allows the user to specify the one or two terminal ID's for which it is desired to collect total terminal data. This option may only be selected if the CAM is active and causes all data seen by the H6000 for the selected terminal to be written to the GMC tape. The following format should be used for specifying the desired terminals:

+ID = Collect data for a single terminal (replace ID with the actual terminal ID)
+ID,ID = Collect data for two terminals

CAUTION: This option can possibly collect passwords and therefore the data tape will be classified to the highest level on the system. Without this option all tapes are unclassified.

5.5.4 Move Option. This option allows the user to suppress an abort if GMC cannot relocate out of the growth range of TSS. See subsection 5.3 for an explanation of the GMC relocation procedure. The proper code for this option is shown below.

M* = suppress abort

This option should not be used if the user wants to ensure that TS1 performance will not be affected by the presence of the GMC. The present version of GMC will not move out of the way of any other copies of TSS that might be active.

5.5.5 CPU SNUMB Option. This option allows the user to specify SNUMBS for the CPU Monitor special option. The CPU Monitor will separately accumulate processor times in its data records for up to five SNUMBS. Multiple SNUMBs must be separated by a comma with no intervening blanks. The last SNUMB must be followed by a blank before any new option is requested.

#SNUMB1 = Accumulate processor time for a single job (replace SNUMB1 with actual SNUMB of a job)

#SNUMB1,SNUMB2 = Accumulate processor time for two jobs

#SNUMB1,SNUMB2,SNUMB3,... = Accumulate processor time for three, four, or five jobs

5.5.6 Connect Option. This option allows the user to suppress collection of tape connects or mass store connects within the CM and/or MSM.

?1 = Only tape connects wanted
?2 = Only mass store connects wanted

The default for this option is that all tape and mass storage connects are collected. The use of this option can significantly reduce the amount of data collected by the CM and MSM.

5.5.7 Time Option. This option allows the user to specify run time parameters. The time capability is to pre-set a time to begin data collection with a time length to run after collection starts.

*07.00,04.00 = start at 7:00 a.m., run for four hours and stop at 11:00 a.m.
*16.00 = start at 4:00 p.m., no stop specified
*,02.50 = start immediately, stop after 2 1/2 hours

Rules for this option are:

a. The time option must be the last parameter on the data parameter card as the card is read left to right and time is the last entry processed.

b. Asterisk signals GMC to process the time input option.

c. Use four characters for all times in each time entry field. Time is expressed as a 24-hour clock. All zero's must be present on the parameter card.

d. If the time option specifies a start at 0900 for a 4 hour run to 1300, and GMC is not spawned until 1000, the run will still terminate at 1300. In this case, only 3 hours of data will be collected, even though 4 hours of collection was specified.

e. The user can request the following: *22.00, 04.00. This means data collection should begin at 22:00 and continue until 02:00. The GMC will handle the problem of a clock rollover.

f. GMC allocates a tape drive as soon as it initially goes into execution. It keeps this tape drive even when it goes to sleep until told to start up. Therefore, if GMC is spawned at 0700 and told to collect data starting at 1100, the tape drive will be allocated from 0700.

g. If no time option is used, the GMC will start collecting data upon entry into the system and terminate upon a console request or tape limit request. When a time option is used, the GMC will terminate normally with a TS abort.

5.5.8 Special Debug. This option should not be used except under special request by the developers of GMC or by someone extremely familiar with the GMC design. It has been observed that for some reason (yet to be determined), the GMC will begin to record data to tape in a nonsequential fashion. This happens very infrequently and will probably never be observed by a site. When it does occur, the data reduction programs will print out warning messages that bad trace records are being processed, or that record sequence numbers are out of order. In most cases, the data reduction programs will recover and process normally.

This input option will cause the GMC to verify every record as it is written to tape. At any time that the collector determines an improper sequence is being written, it will abort with one of the following aborts:

R1 - Bad record written at WRBUF1

R2 - Bad record written at WRBUF2

R3 - First Record on a new tape has a bad record number at TAPD2

R4 - Tape write from disk has a bad record number at TAPWRT

R5 - Disk write has a bad record number at WRITIT

R6 - The disk write block number is bad at status check MLCC

R7 - The disk read block number is bad at status check DIRDCC

To use this option, simply type the letter K on the GMC data card.

5.5.9 Limited Mass Store Monitor/Channel Monitor Collection. In order to limit the amount of data being collected by the Mass Store and/or Channel Monitors, the user can request that only data being generated from certain jobs should be collected. To use this option, the characters MS must appear on the data card. If the "S" character is immediately followed by a blank, then only data for FTS, TS1 and $PALC will be captured. If the user desires, he may request that data

THIS PAGE LEFT INTENTIONALLY BLANK

from five additional SNUMBs also be captured. To use this option, the SNUMBs must appear on the data card immediately after the "S" character. No blank should be present between the "S" of the MS option and the first letter of the first SNUMB. The SNUMBs must be separated by commas (no intervening blanks) and the last SNUMB must be followed by at least one blank column before a new input option is requested.

5.5.10 Request the Next Data Card. If the user is unable to fit all the aforementioned parameters on a single data card, an additional data card can be used. The user must inform the GMC that a second data card will be present by placing an "X" on the first data card. The "X" must be the last entry on the first data card and must not be placed in the middle of an input option. A given input option should be completely described before the "X" option is used. No more than two cards can be used to describe all of the standard GMF options.

5.5.11 Disable the Dispatcher Queuing Option. In order to turn off the dispatcher queuing option of the CPU Monitor, the user should type the letter "Q" on the GMF data card. See subsection 5.2.3 for an explanation of dispatcher queuing monitoring.

5.5.12 Specifying Monitoring Requirements for the GRTM. In order to collect GRTM data, an M6 must not appear on the first data card. If the M6 is omitted from the first card, then the GRTM is active, in which case additional data cards are required. The datanet specifications must be placed on a separate data card and are not directly related with the previously described options. An "X" option must not be used to indicate that the datanet option card is present. If an M6 does not appear on the standard data cards then the GMC will expect an additional data card to follow any of the standard data cards that are present. The additional data cards are free format and indicate the datanets to be monitored, the HSLA subchannels to be monitored, and finally whether response time monitoring is to be performed. By default, only CPU resource monitoring will be conducted. As stated earlier, if the entire monitoring function is to be selected, the GRTS monitor will require approximately 2K of DATANET memory. On the other hand, if only the default option is selected, the monitor will require only 1K of DATANET memory. The required parameter categories are as follows:

Dn = FEP number (n=0 to 7)

HSLAn = High Speed Line Adapter (HSLA) number (n=1 to 3 per FEP)

SCHn = Subchannel numbers associated with each HSLA entry (n=0 to 31)

Rn = Performance response time monitoring on FEP number (n=0 to 7)

Semicolons delimit each of the categories. They also indicate that more GRTS data follows. However, the last data card should not have semicolon at the end. Commas delimit subchannel sets. A "-" specifies a range of values for subchannels.

Example:

D1;HSLA1;SCH0-10,14,18-30;HSLA2;SCH3-15,20-30;D0;R0

In this example, we will monitor DATANET #1 for CPU resources (by default), for subchannel usage on HSLA 1 subchannels 0-10, 14, 18-30, and HSLA2 subchannels 3-15 and 20-30. No response time monitoring will be performed. For DATANET 0, we will monitor CPU resources (by default), no subchannels will be monitored but response time will be monitored. In order to determine the total memory used by the monitor, the following formula should be used:

Memory Used = 1K (default for CPU resource monitoring)
+ 32 words * (number of HSLAs)
+ 8 words * (number of subchannels requested) +
1K (if response monitoring selected)

5.5.13 General Rules of the GMC Data Parameter Card. The following are general rules to be followed in defining the data card:

a. The time option, if selected, must be the last option on the data card.

b. All input elements of the eight options should be separated by a blank character.

5.6 JCL for Creation of an Object File

5.6.1 Introduction to JCL. After the user has completed a study of the options to be specified in the parameter cards described in subsection 5.5 above, the user then must build the JCL that will create the user version of a GMC object file.

The user has optional control over the creation of a GMC object file that will serve his purposes based on functions specified on the parameter cards. This optional structure minimizes the size of the GMC monitor as only necessary code is used and provides, in addition, for easier extension to the capabilities of GMC.

The four functions to be built are initialization, system patch, remove the patch, and primary monitor collection subroutines.

5.6.2 Creation of an Object File. The GMC executive routine has been subdivided into discrete sections of code based on function. In order to generate a usable GMC object file, the individual sections of code must be merged to create a single routine. This procedure has been utilized for two reasons. First, this structure permits the easy addition of new programs, i.e., monitors. Secondly, this structure allows the simple generation of a GMC containing only that code required to capture the necessary data. If a user wants to create a GMC containing only the code necessary for the Memory Utilization Monitor and Mass Store Monitor, he can easily do so. If the user then decides he does not want to run the Mass Store Monitor, he can either recreate a GMC object file or he can turn off the Mass Store Monitor via a data card. Although this procedure sounds complicated, it minimizes the size of the GMC.

When GMF is restored to the user's system, the GMC data collector program is in the form of a catalog structure. This structure is shown in figure 5-3. The GMC catalog structure as it relates to the creation of GMC R* file is shown in figure 5-4. For a description of each file function, refer to table 5-3.

As illustrated in figure 5-1, the GMC consists of an Executive Routine which initializes all the required programs, installs any required system patches, records the system patches, and removes all patches from the system when it is finished. Table 5-3 gives a breakdown of all required GMC files and any optional files. The user selects the programs he wants to monitor in the system, assembles those programs in numerical order (as listed in table 5-3), and runs a file-edit job to create a GMC object file. Figure 5-5 is an example of a GMC containing the Memory Utilization, Idle, and CPU Monitors. Figure 5-6 is an example of a GMC containing the Mass Store Monitor, Channel Monitor, and Idle Monitor. A $ GMAP card is required before the SELECTA for /GMF.TOP, and before every SELECTA after /GMF.BTM. Under normal operating conditions, the FILEDIT activity will contain multiple "Inconsistent Deck Name" error messages, which can be ignored. If a user should improperly create an R* file by omitting a required file, the GMC execution will abort with an S1-S9 abort, or an SA, SC, or SD abort, depending upon which routine is missing. The user should refer to table 5-2 for an explanation of these aborts.

Within the GMC release, precanned JCL has already been provided for the generation of 18 different monitor combinations. See table 5-3 for a list of the files containing this set of JCL.

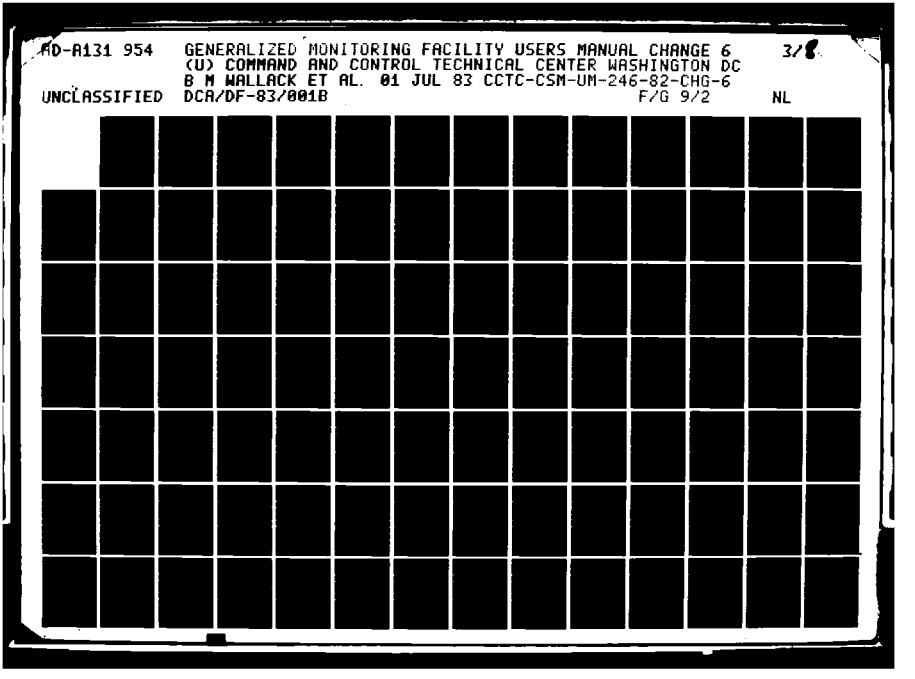
The GMF is designed to be run on a HIS 6000 computer system, running with WWMCCS GCOS release 6.4 or 7.2. These releases are equivalent to the HIS commercial 2H or 4JS1 GCOS releases. The CPU, TSS and TPE Monitors must be run under a WW7.2 release.

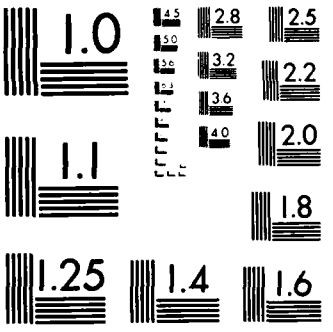
B29IDPX0

GMFCOL

MAKE	GMF.OBJ	PATLOOK	GMF	MUM	MSM	CM	CPU	TM
ALL			GMF.TOP	MUM.INIT	MSM.INIT	CM.INIT	CPU.INIT	TM.INIT
MUM			GMF.BTM	MUM.T10	MSM.PAT	CM.T04A	CPU.PAT	TM.T50
CPU			GMF.MID	MUM.T46	MSM.REMO	CM.T22A	CPU.T70	
TM			GMF.MON		MSMDOIT	CM.T07A	CPU.REMO	
MSM							CPUDOIT	
CAM								
CM								
GRT								
TPE								
TSS								
MO2								
MO2S8								
M148								
M48								
M14								
M56								
MO[illegible]								
MO1345								

Figure 5-3. GMC Catalog File Structure (Part 1 of 2)


AD-A131 954
GENERALIZED MONITORING FACILITY USERS MANUAL CHANGE 6
(U) COMMAND AND CONTROL TECHNICAL CENTER WASHINGTON DC
B M WALLACK ET AL. 01 JUL 83 CCTC-CSM-UM-246-82-CHG-6
UNCLASSIFIED
DCA/DF-83/001B
F/G 9/2
NL



MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS 1963 A

B29IDPX0

GMFCOL (Continued)

CAM	GRT	IDLE	TPE	TSS	RUN.GMF
CAM.INIT	GRT.INIT	IDL.INIT	TP.INIT	TSS.PAT	
CAM.REMO	GRT.T62	IDL.TRCS	TPE200	TSS.COL	
CAM.PAT	GRT.COL	IDL.T21		TSSMON	
CAM.T14				TSS.INIT	
CAMDOIT				TSGMF	

Figure 5-3. (Part 2 of 2)

$:SELECTA:B29IDPXO/GMFCOL/(see below for FILE requirements and options)

REQUIRED	OPTIONAL, WHEN A MONITOR IS USED ALL PROGRAM ELEMENTS OF THE MONITOR MUST BE USED									
GMF.	MUM.	MSM.	CM.	CPU.	TM.	CAM.	GRT.	IDLE	TP	TSS
GMF.TOP	MUM.INIT	MSM.INT	CM.INIT	CPU.INIT	TM.INIT	CAM.INIT	GRT.INIT	IDL.INIT	TP.INIT	TSS.INIT
GMF.MID		MSM.PAT		CPU.PAT		CAM.PAT				
GMF.PATLOOK*1		MSMDOIT		CPUDOIT		CAMDOIT				
GMF.MON		MSM.REMO		CPU.REMO		CAM.REMO				
GMF.BTM	MUM.T10		CM.T04A	CPU.T70	TM.T50	CAM.T14	GRT.T62	IDL.TRCS	TPE200	TSS.COL
	MUM.T46		CM.T22A				GRT.COL	IDL.T21		
		CM.T07A*2	CM.T07A							

NOTE-1 GMF.PATLOOK is required only with MSM, CPU and CAM.
NOTE-2 CM.T07A is required with MSM. However, if using both MSM. and CM., then only use only one copy of CM.T07A.

Figure 5-4. GMC JCL Structure

```
$       FILEDIT SOURCE,OBJECT, INITIALIZE
$       PRMFL   R*,W,S,B29IDPXO/GMFCOL/GMF.OBJ
$       FILE    K*,NULL
$       DATA    *C,COPY
$       LOWLOAD
$       OPTION  ERCNT/500/
$       GMAP    NDECK
$       SELECTA B29DIPXO/GMFCOL/GMF/GMF.TOP
$       SELECTA B29IDPXO/GMFCOL/MUM/MUM.INIT
$       SELECTA B29IDPXO/GMFCOL/CPU/CPU.INIT
$       SELECTA B29IDPXO/GMFCOL/IDLE/IDL.INIT
$       SELECTA B29IDPXO/GMFCOL/GMF/GMF.MID
$       SELECTA B29IDPXO/GMFCOL/CPU/CPU.PAT
$       SELECTA B29IDPXO/GMFCOL/PATLOOK
$       SELECTA B29IDPXO/GMFCOL/CPU/CPUDOIT
$       SELECTA B29IDPXO/GMFCOL/GMF/GMF.MON
$       SELECTA B29IDPXO/GMFCOL/CPU/CPU.REMO
$       SELECTA B29IDPXO/GMFCOL/GMF/GMF.BTM
$       GMAP    NDECK
$       SELECTA B29IDPXO/GMFCOL/IDLE/IDL.TRCS
$       GMAP    NDECK
$       SELECTA B29IDPXO//GMFCOL/IDLE/IDL.T21
$       GMAP    NDECK
$       SELECTA B29IDPXO/GMFCOL/MUM/MUM.T10
$       GMAP    NDECK
$       SELECTA B29IDPXO/GMFCOL/MUM/MUM.T46
$       GMAP    NDECK
$       SELECTA B29IDPXO/GMFCOL/CPU/CPU.T70
$       EXECUTE
$       ENDEDIT
$       ENDCOPY
```

Figure 5-5. Creation of R* File for Memory, Idle, and CPU Monitor

```
$       FILEDIT     SOURCE,OBJECT,INITIALIZE
$       PRMFL       R*,W,S,B29IDPX0/GMFCOL/GMF.OBJ
$       FILE        K*,NULL
$       DATA        *C,COPY
$       LOWLOAD
$       OPTION      ERCNT/500/
$       GMAP        NDECK
$       SELECTA     B29IDPX0/GMFCOL/GMF/GMF.TOP
$       SELECTA     B29IDPX0/GMFCOL/MSM/MSM.INIT
$       SELECTA     B29IDPX0/GMFCOL/CM/CM.INIT
$       SELECTA     B29IDPX0/GMFCOL/IDLE/IDL.INIT
$       SELECTA     B29IDPX0/GMFCOL/GMF/GMF.MID
$       SELECTA     B29IDPX0/GMFCOL/MSM/MSM.PAT
$       SELECTA     B29IDPX0/GMFCOL/PATLOOK
$       SELECTA     B29IDPX0/GMFCOL/MSM/MSMDOIT
$       SELECTA     B29IDPX0/GMFCOL/GMF/GMF.MON
$       SELECTA     B29IDPX0/GMFCOL/MSM/MSM.REMO
$       SELECTA     B29IDPX0/GMFCOL/GMF/GMF.BTM
$       GMAP        NDECK
$       SELECTA     B29IDPX0/GMFCOL/CM/CM.T04A
$       GMAP        NDECK
$       SELECTA     B29IDPX0/GMFCOL/CM/CM.T22A
$       GMAP        NDECK
$       SELECTA     B29IDPX0/GMFCOL/CM/CM.T07A
$       GMAP        NDECK
$       SELECTA     B29IDPX0/GMFCOL/IDLE/IDL.TRCS
$       GMAP        NDECK
$       SELECTA     B29IDPX0/GMFCOL/IDLE/IDL.T21
$       EXECUTE
$       ENDEDIT
$       ENDCOPY
```

Figure 5-6. Creation of R* File for Mass Store Monitor, Idle and Channel Monitor

When GMF is used on WWMCCS release 7.2, or commercial release 4JS1-4JS3, the user must ensure that the value for variable "SYS64" is set to 0. This variable is defined in an "EQU" statement located in the following files: B29IDPXO/GMFCOL/GMF/GMF.TOP, B29IDPXO/GMFCOL/CM/CM.TO7A, and B29IDPXO/GMFCOL/MUM/MUM.T10. See subsection 2.6.2 for a discussion of other GMC modifications that are required under different GCOS software releases.

Having created a GMF object file, no other system modifications are required unless the GRT Monitor, the TPE Monitor or the TSS Monitor are desired. Each of these monitors require system modifications to be made prior to their use. The system modifications required by the GRTS Monitor are described in subsection 5.2.7 of this section. The system modifications required by the TPE Monitor are described in subsection 5.2.9 of this section. The system modifications required by the TSS Monitor are described in subsection 5.2.10 of this section.

5.7 JCL for Executing the GMC

The JCL needed to execute the GMC is shown in figure 5-7. The size to be placed on the $ LIMITS card depends on the number of monitors present in the R* file. The size should range from 15K to 24K, depending on the number of monitors. The load map produced on the compilation listing from the General Loader will specify the actual memory size required to load GMC. In figure 5-7, parameter card following $ DATA I* demonstrates a turnoff for the MUM, CPUM, TM, GRTM, IDLE, TPEM monitors, and a collection of an unlimited number of tapes. If the GRTS monitor is active, it may dynamically grow GMC during its initialization procedure. Because of this feature, the GRTS monitor requires a $ LIMITS card with a large SYSOUT request to force the system to obtain an extra 1K memory. Figure 5-8 shows the JCL needed to execute a GMC session which includes the GRTM. Since the GMC runs in master mode, it requires a PRIVITY card, and the operator must grant it permission to run. The DK file is optional but must be present if more than one reel will be collected. The size of the DK file is approximately 300 links but may require more than 300 links on a very busy system. This file is used to collect data during rewind of a completed data tape. When GMF loads, many load error messages will be produced. These error messages may all be ignored.

```
$       OPTION  ERCNT/999/
$       LOWLOAD
$       EXECUTE DUMP
$       PRMFL   R*,R,S,B29IDPX0/GMFC0L/GMF.0BJ
$       LIMITS  99,16K
$       PRIVITY
$       TAPE    0T,X2D,,99999
$       FILE    DK,X1R,300R
$       DATA    I*
M0 M2 M3 M6 M7 M8 (optional)
$       ENDJOB
```

*NOTE - The execution report from GMC will contain multiple "nonfatal error" messages which can be ignored.

**NOTE - If high density tape collection is desired, the tape card should be modified with the addition of four commas after the tape number followed by the words DEN16:

```
$     TAPE     0T,X2D,,12345,,,,DEN16
```

and also the option M; should be on the data card.

Figure 5-7. JCL for Executing the GMC

```
$       LOWLOAD
$       OPTION          ERCNT/500/
$       EXECUTE         DUMP
$       PRMFL           R*,R,S,B29IDPX0/GMFC0L/GMF.CBJ
$       LIMITS          99,15K,,99999
$       PRIVITY
$       TAPE            0T,X2D,,,,GRTSII-DATA
$       FILE            DK,X1D,300R
$       DATA            I*
M0 M1 M2 M3 M4 M5 M7 M8 M9 MA
 D0;HSLA1;SCH0-8,14,18,21-23;D1;R0
$       ENDJOB
```

Figure 5-8. GRTM JCL

SECTION 6. MEMORY UTILIZATION MONITOR DATA REDUCTION PROGRAM

The Memory Utilization Data Reduction Program (MUDRP) is a FORTRAN program which uses the GMC output tape to produce a series of reports presenting various aspects of memory utilization. Data collected by the Idle Monitor will also be processed by MUDRP and reported as part of the MUM reports.

6.1 Inputs

The Memory Utilization Data Reduction Program has two inputs. The first is the data collection tape produced by the General Monitor Collector (see section 5). The second input is a set of control cards specifying the output report options.

6.1.1 Report Options. The MUDRP reports consist of histograms, serial plots, and general report outputs. These reports have a number of parameter selections which can be modified by card input at run time. Any report may be turned on or off, and a set of time spans may be specified for measurement.

6.1.2 Default Options. Table 6-1 lists all reports produced by default. All reports except the Memory Map and Out of Core Report are produced unless specifically turned off, by user request.

6.1.3 Histogram Options. There are five characteristics for each histogram directly available to the user. These include the ranges and resolution of display as well as the descriptive titles. The format of a histogram report is completely described later in this chapter (section 6.3.2).

The first of these characteristics is the lowest value, LOWVAL. Any input received by the histogram that is less than this value will be entered into the first histogram bucket, along with all entries that are less than the starting value of the next bucket.

The next entry range, the interval size of each histogram entry, is determined by the INTVSZ parameter. It specifies the size of each entry range or "bucket."

The size of the histogram is determined from the parameter TABSIZ which specifies how many "buckets" or entry ranges are contained in this histogram. TABSIZ also specifies the upper value which could fall within the display range of the histogram. If an entry is made above this range, it is placed in an out-of-range "bucket" of the histogram. The number of the out-of-range occurrences are output with the histogram, along with their average value, so that the histogram can be altered to handle the variability of the range better. The user should note that a compile time parameter of MXTBSZ determines how large any one histogram can become. If an increase beyond this is desired, MXTBSZ is increased along with the memory size of the data reduction program. There are two ways to modify the histograms so as to include the out-of-range values. The first method would involve increasing the internal size of a histogram bucket by altering the value of INTVSZ. In this way, the histogram would

Table 6-1. Default Reports (Part 1 of 4)

ID Number	Histogram Title
1	Memory Demand Size of New User Activities in 1K Word Blocks
2	The Original Allocation Time for User Memory in 1/10 Second
3	The Total Elapsed Time a User Activity was Swapped
4	The Number of Times a User Activity was Swapped
5	The Request Size of GEMORES
6	The Percent of Size-Time Product Used by a User Activity
7	The GEMORE Service or Denial Time - 1/10 Second
8	The Elapsed Time a User Activity was in Memory
9	The Elapsed Duration of a User Activity in Tenths of a Second
10	The Ratio of User Activity Duration Versus Memory Use Time
11	The Memory Demand Size of All Demand Types
12	Elapsed Time Between Allocator Calls in 1/100 of a Second
13	Number of Extra Activities that Could Fit in Memory With Compaction
14	Number of Activities Waiting for Memory in the Allocator Queue *
15	Number of Activities Residing in Memory *
16	Elapsed Time of a Busy State (All Processors)
17	The Total Amount of Available Memory *
18	The Potential Time-Weighted Memory Demand
19	Demand Outstanding When a Processor Went Idle *
20	Number of Activities in Memory When a Processor Went Idle *

Table 6-1. (Part 2 of 4)

ID Number	Histogram Title
21	Number of Activities Waiting Memory When a Processor Went Idle *
22	Memory Available When a Processor Went Idle *
23	Delay Time in the System Scheduler
24	Delay Time Until Core Allocation
25	Percent of Assigned Memory Used (Time-Corrected) *
29	Number of User Activities Waiting Memory in the Allocator Queue *
30	Number of User Activites in Memory *
31	Elapsed Time of a Busy State Processor 0
32	Elapsed Time of a Busy State Processor 1
33	Elapsed Time of a Busy State Processor 2
34	Elapsed Time of a Busy State Processor 3
35	CP Time Per User Activity
36	I/O Time Per User Activity
40	Number of Activities Waiting Memory (Time-Corrected) *
41	Number of Activities in Memory (Time-Corrected) *
42	Memory Available (Time-Corrected) *
43	Number of User Activities Waiting Memory (Time-Corrected) *
44	Number of User Activities in Memory (Time-Corrected) *
45	Total Demand Outstanding (Time-Corrected) *

* Jobs with 0 urgency are not included.

Table 6-1. (Part 3 of 4)

ID Number	Histogram Title
46	Number of Extra Activites That Could Fit in Memory Without Compaction
48	Length of an Idle State (All Processors)
49	Length of an Idle State Processor 0
50	Length of an Idle State Processor 1
51	Length of an Idle State Processor 2
52	Length of an Idle State Processor 3
53	Number of Times System Activity Swapped
54	Elapsed Time a System Activity was Swapped
55	Elapsed Time of a Busy State Processor 4
56	Elapsed Time of a Busy State Processor 5
57	Length of an Idle State Processor 4
58	Length of an Idle State Processor 5

ID Number/Name	Plot Title
26/PLOT1	Availability of Memory vs. Outstanding Demand In Core Allocator Queue vs. Outstanding Demand in Peripheral Allocator Queue Plus Outstanding Demand in Core Allocator Queue
27/PLOT2	Memory Shortfall in Core Allocator Queue vs. Memory Shortfall in Core Allocator Queue Plus Memory Shortfall in Peripheral Allocator Queue
28/PLOT3	Number of Activities in Core Allocator Queue vs. Number of Activities in Peripheral Allocator Queue
59/PLOT4	Memory Demand Outstanding vs. Potential Demand vs. Zero Urgency Demand

Table 6-1. (Part 4 of 4)

ID Number/Name	Other Reports
37/PALC	Peripheral Allocator Report
38/ACTIVE	Activity Report/Excessive Resource Report/Abort Report/IDENT Report
39/MAP	Memory Map
47/OUT	Out of Core Report
---	Special Job Memory Reports
---	System Program Usage Report
---	Memory Statistics Report
---	Distribution of Urgency Over Time Report
---	Zero Urgency Job Report

cover a larger range of values. This change could be made via data cards and would not increase the size of the program.

The second method would involve increasing the size of the histogram by altering the value of TABSIZ. As long as the size requested does not exceed 50, this change can also be done via a data card. However, if an individual histogram needs to be larger than 50 buckets, the user will need to change the value of MXTBSZ. This change will require a change to source code, a recompile, and probably an increase in program size. All references to MXTBSZ must be altered. This would need to be done in the EDIT subsystem of Timesharing.

The remaining items that can be modified are the title and the vertical axis headers. The method for altering the histogram parameters is detailed in subsection 6.1.6. Table 6-2 shows the default values for all histograms.

6.1.4 Plot Options. There are three characteristics directly available to the user for each individual plot axis used.

The first characteristic, MAXNUM, is the maximum number of entries to be plotted on each vertical plot axis.

The second characteristic, YMAX, defines the upper limit of the horizontal display axis.

The third characteristic, YMIN, defines the lower limit of the horizontal display axis. The method for altering these values is explained in subsection 6.1.7. Table 6-3 shows the default values for all plots.

6.1.5 Default Option Alteration. The general format for an option request is as follows: the first card contains an action code describing the action to be taken. Subsequent cards modify report parameters for some of the action codes. All input cards are free format with the only requirements being that at least one blank space separates multiple input parameters. The very last input card must have the word "END" typed on it. This card must be present whether or not any other input options are selected. Available actions with their (default) implications are shown in table 6-4. There is no order required for the options. In reading the following sections it should be remembered that the first card for any input option must be the action code specification with no other data present on the card.

The user should take special note that if this software is executed under a VW6.4/2H GCOS release, an additional data card is required. The data card should contain the letters RN.

Table 6-2. Default Values for Histograms (Part 1 of 2)

ID #	Low Value	Interval Size	Number of Buckets
1	4	4	50
2	0	50	50
3	0	250	50
4	0	1	50
5	4	4	50
6	0	1	50
7	0	5	50
8	0	200	50
9	0	200	50
10	.95	.1	50
11	4	4	50
12	0	10	50
13	0	1	50
14	0	1	50
15	0	1	50
16	0.0	5.0	50
17	0	10	50
18	4	10	50
19	4	20	50
20	0	1	50
21	0	1	50
22	4	8	50
23	0	25	50
24	0	25	50
25	50	2	50
29	0	1	50
30	0	1	50
31	0.0	5.0	50
32	0.0	5.0	50
33	0.0	5.0	50
34	0.0	5.0	50
35	5	5	50
36	5	5	50
40	0	1	50
41	0	1	50
42	0	10	50
43	0	1	50
44	0	1	50
45	0	10	50
46	0	1	50
48	0.0	5.0	50
49	0.0	5.0	50
50	0.0	5.0	50

Table 6-2. Default Values for Histograms (Part 2 of 2)

ID #	Low Value	Interval Size	Number of Buckets
51	0.0	5.0	50
52	0.0	5.0	50
53	0	1	50
54	0	250	50
55	0.0	5.0	50
56	0.0	5.0	50
57	0.0	5.0	50
58	0.0	5.0	50

Table 6-3. Default Values for Plots

ID #	Max Size of Plot	Lower Plot Limit	Upper Plot Limit
26	Unlimited	0.	456.
27	Unlimited	0.	456.
28	Unlimited	0.	114.
59	Unlimited	0.	228.

Table 6-4. Available Report Actions and Their (Default) Values (Part 1 of 2)

HISTG - Modify a histogram (see table 6-2)

PLOT - Modify a Plot (see table 6-3)

ON - Turn a specific report on - (all reports on except Memory Map and Out of Core Report)

OFF - Turn a specific report off - (all reports on except Memory Map and Out of Core Report)

TIME - Set a timespan(s) for reporting - (total time reported)

ALLOFF - Turn all reports off except those specified - (all reports on except Memory Map and Out of Core Report)

ALLON - Turn all reports on except those specified - (all reports on except Memory Map and Out of Core Report)

ERROR - Do not stop on an option request error - (stop on an input error)

DEBUG - Program debug requested - (no debug)

ALLOC - Stop program after a specified number of memory allocations have been requested - (entire tape processed)

NREC - Stop program after a specified number of tape records have been processed - (entire tape processed)

NOUSER - Do not print USERID on any report - (USERID printed on certain reports)

IDLE - Turn off all Idle Monitor reports - (all IDLE reports on)

WASTED,CORE,IO,CPU,RATIO,URG - Changes parameters used in the Excessive Resource Usage Report - (20K,50K,30MIN, 30MIN,5,40)

ABORT - SNUMBs not to report in the ABORT Report - (all SNUMBs that abort are reported)

PLTINT - Change Interval at which plots are printed - (10 MIN)

FSTSLV - Change the lowest allowable user program number - (14 decimal)

Table 6-4. (Part 2 of 2)

MASTER - Define SNUMBs that are considered to be SYSTEM programs - (all programs with a program number less than FSTSLV). In addition, the programs $TRAX, $HEALS and VIDEO will be considered system programs.

PALC - Change the print control for the PALC report (600 secs)

END - Required as last card of input. It must be present.

SPECL - Produce the Special Job Memory Reports

RN - Processing on a WW6.4 system

MAPART - Produce a memory map only when the number of jobs waiting memory surpasses a predetermined value.

ZERO - Include zero urgency jobs in all memory calculations (zero urgency jobs are not included).

6.1.6 Histogram Alterations (Action Code HISTG). A complete description of histogram default values and their meanings is provided in section 6.1.3. In order to change any histogram parameter, the user is required to supply a series of four data cards. The first card contains the action code HISTG. The second card specifies which histogram the user wants to alter. This specification is made by inputting the histogram ID number as obtained from table 6-1. The third data card describes the parameter to be changed and the fourth card provides the new value for the parameter. The following options are available:

Card #3	Card #4
LOWVAL	A new low value
SIZE	A new maximum histogram size
BUCKET	A new bucket size
HEADER	A two-word header separated by at least one blank. Each header word must not exceed six characters in length. If one of the headers is to be blank, the word BLANK must be typed on the data card.

Two additional points must be stressed.

- If the user wants to change multiple parameters for a single histogram, then cards 3 and 4 should be repeated as often as required. When alterations for a given histogram are completed, the user must supply a new action request. If the user desires to change another histogram, then the entire sequence of four data cards must be repeated.

- When inputting the new parameter values, the user must consult table 6-2 in order to determine whether the parameters must be inputted as integer (no decimal point), or as real numbers (decimal point must appear on data card).

Figure 6-1 shows a standard histogram format. Column five of this histogram has the words "NUMBER" and "WAIT." These are called the header labels and are used to describe the function being reported by this histogram. It is these two words that the user may modify with the HEADER parameter card. Figure 6-2 shows the input option formats for this action code.

6.1.7 Plot Alterations (Action Code PLOT). Modifications to a plot allow the user to specify a new plot size, a new maximum horizontal axis limit and a new minimum horizontal axis limit. The default values for the existing plots are described in section 6.1.4. As with the histogram option, the user is required to supply a series of four data cards for each parameter change desired. The first data card contains the action

DISTRIBUTION COLLECTED ON SYSTEM NMCC2 AT 12:39:46 ON 80-12-15

THE # OF USER ACTIVITIES WAITING MEMORY IN ALLOCATORS QUEUE

INDIV. NUMBER	CUMUL. NUMBER	CUMUL. PRC	INDIV. PRC	NUMBER WAIT	PERCENT OF OCCURRENCE 00 05 10 15 20 25 30 35 40 45 50	REPORT 10
					I....I....I....I....I....I....I....I....I....I....I	
1456	1456	6.220	6.220	0- 0	IXXXXXX	
2257	3713	15.861	9.641	1- 1	IXXXXXXXXXX	
3949	7662	32.730	16.869	2- 2	IXXXXXXXXXXXXXXXXX	
5456	13118	56.036	23.306	3- 3	IXXXXXXXXXXXXXXXXXXXXXXX	
4053	17171	73.349	17.313	4- 4	IXXXXXXXXXXXXXXXX	
2924	20095	85.839	12.490	5- 5	IXXXXXXXXXXX	
1836	21931	93.682	7.843	6- 6	IXXXXXXXX	
878	22809	97.433	3.751	7- 7	IXXXX	
317	23126	98.787	1.354	8- 8	IX	
114	23240	99.274	0.487	9- 9	I	
133	23373	99.842	0.568	10- 10	IX	
34	23407	99.987	0.145	11- 11	I	
3	23410	100.000	0.013	12- 12	I	

23410 ENTRIES TOTAL AVERAGE = 3.40961 VARIANCE = 3.821 STANDARD DEVIATION = 1.955

Figure 6-1. Standard Histogram

Card		
	1	A
	2	B
	3	C
	4	D

where

A = The word HISTG
B = Histogram ID number (table 6-1)
C = Parameter control word
D = New parameter value. If parameter control word was HEADER, then this card must contain two words separated by at least one blank. Each word cannot exceed six characters in length. If the user desires one of the words to contain blanks, he must type the word BLANK on the card.

* Cards 3 and 4 are repeated for as many parameter changes as desired for the single histogram defined by card 2.

Figure 6-2. HISTG Action Code Format

code PLOT. The second card specifies which plot the user wants to alter. This specification is made by inputting the plot ID number from table 6-1. The third data card describes the parameter to be changed and the fourth data card provides the new value for the parameter. The following options are available:

Card #3	Card #4
SIZE	A new maximum plot size
LOWVAL	A new low value
HIVAL	A new high value

Several points must be stressed.

- When inputting parameter values for LOWVAL and HIVAL, these values must be inputted as real numbers (decimal point must appear on data card). If a new LOWVAL or HIVAL value is selected, the user should try to ensure that (HIVAL-LOWVAL) is divisible by 114.

- When inputting parameter values for SIZE, the value must be inputted as an integer. If the size is to be unlimited, then a -1 must be inputted.

Figure 6-3 shows a standard plot format. The maximum and minimum values of the plot are used to determine the value of each dash across the horizontal axis. In this figure, each dash has a delta value of 1.0. Figure 6-4 shows the format for this input option.

6.1.8 Turn a Report On (Action Code ON). This card allows a user to turn on a single report that is off by default. Card 1 contains the word ON and card 2 contains the NAME of the report to be turned on (table 6-1). No change to the default parameters, of the report, will be made. This option may be used only for Plots and Other Reports and cannot be used to control individual histograms. Note that only those reports and/or plots that have a specific NAME can be controlled with this option.

6.1.9 Turn a Report Off (Action Code OFF). This card allows a user to turn off a single report that is on by default. Card 1 contains the word OFF and card 2 contains the NAME of the report to be turned off (table 6-1). No change to the default parameters, of the report, will be made. This option may be used only for Plots and Other Reports and cannot be used to control individual histograms. Note that only those reports and/or plots that have a specific NAME can be controlled with this option.

6.1.10 Set a Timespan of Measurement (Action Code TIME). The timespan of data collection can cover many hours of which only a few may be of interest. This option allows a user to specify the timespan (or spans) to be displayed in all reports. For example, the user may specify that he wants to collect data from 0500 to 2200 and wants to display data only

from 0900 to 1700 in all reports. As another option, the user may request to see the memory map from 0900 to 1000, plot #1 from 1200 to 1500 and all other reports from 0800 to 1700.

The user must specify the report NAME to be affected by the time request (table 6-1). If the entire reduction (all reports) are to be time controlled, a report NAME of "TOTAL" must be used. Histogram reports cannot be individually time spanned. Note that only those reports and/or plots that have a specific NAME can be controlled with this option. All time spans of plots or other reports will be bound by the total report timespan, if one is to be used. Up to five timespans for each report or plot may be specified, and they must be serially ordered. All times are expressed as SIMI where SI is the hour and MI is the minute. All time

DISTRIBUTION COLLECTED ON SYSTEM NMCC2 AT 12:39:46 ON 80-12-5

OF ACTIVITIES IN CALC QUEUE VS # IN PALC QUEUE

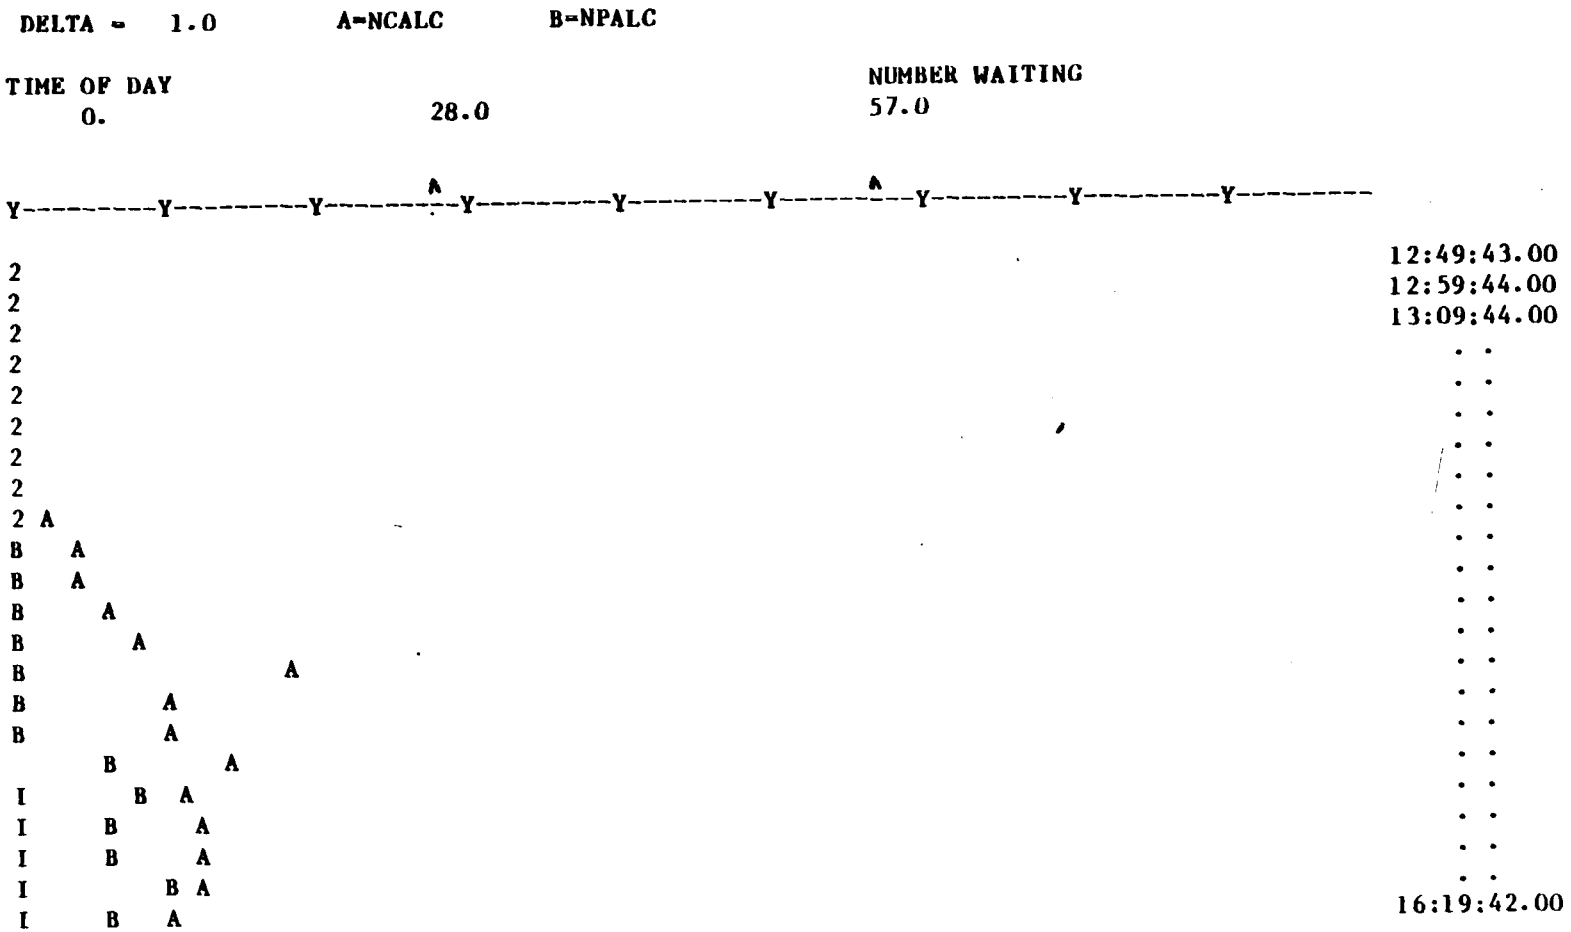


Figure 6-3. Standard Plot

Card		
	1	A
	2	B
	3	C
	4	D

A = The word PLOT
B = Plot ID number (table 6-1)
C = Parameter control word
D = New parameter value

* Cards 3 and 4 are repeated for as many parameter changes as desired for the single plot defined by card 2.

Figure 6-4. PLOT Action Code Format

must be expressed as four character fields with no intervening blanks. Time is based on a 24-hour clock. If a user wants to request the time 4:07, he must input 0407. All times must include four characters.

If a start time, but no stop time, is desired, no characters should be entered after the minutes of the start time. If a stop time is requested, there must be a start time corresponding to it. If the user wants to start at the beginning of data collection and stop at some specified time, but is not sure of the start time, a start time of 0001 should be used. Figure 6-5 shows the format for this option.

6.1.11 Turn All Reports Off Except Those Specified (Action Code ALLOFF). All reports except those explicitly identified here are to be turned off. The inputs consist of

A B C . . . Y (max of 25)

where A through Y are the report ID numbers (table 6-1) to be turned on. The format is shown in figure 6-6. This option will control the printing of all reports, including histograms, if they contain a specific ID number.

6.1.12 Turn All Reports On Except Those Specified (Action Code ALLON). All reports except those explicitly identified here are to be turned on. The input consists of

A B C . . . Y (max of 25)

A through Y are the report ID numbers (table 6-1) to be turned off. The format is the same as action code ALLOFF (see figure 6-6). This option will control the printing of all reports, including histograms, if they contain a specific ID number.

6.1.13 Continue Data Reduction After an Input Option Error (Action Code ERROR). This code allows data reduction to continue when an error has been detected and reported in an input option request. The default value will abort data reduction and report the error. Only the Action Code card is required.

6.1.14 Debug For a Given Program Number (Action Code DEBUG). This is a debug option which supplies large amounts of output for a given program number. It should be used only in cases of data reduction problems. Card 1 contains the word DEBUG and card 2 contains a program number. A program number of -150 will provide detailed debug on system scheduler activities.

6.1.15 Stop After a Specified Number of Tape Records Processed (Action Code NREC). This option is useful when a tape problem occurs and the entire tape cannot be processed. When this occurs, the program will usually abort with an I/O error and some reports might be lost. If a tape error does occur during data reduction, the operator should type a "U" in

```
Card  1   A
      2   N M
      3   B C D E ...
```

where
A = The word TIME
N = Report NAME to be time spanned (table 6-1)
M = Number of different times appearing on Card 3
B,C,D,E = Start and stop times used to define the time spans.
Times must be separated by one or more blanks.

Figure 6-5. TIME Action Code Format

```
Card  1 = A
      2 = N
      3 = B C D E ...

 where
A = The word ALLOFF or ALLON
N = The number of report ID's appearing on card 3    cannot exceed 25
B,C,D,E = ID numbers of those reports not to be turned OFF/ON
```

Figure 6-6. ALLOFF/ALLON Action Code Format

in order to stop data reduction processing prior to the tape error. The first card contains the word NREC and the second card contains the number of tape records to be processed.

6.1.16 Suppress USERID (Action Code NOUSER). This action code is used to suppress the printing of USERIDs on those reports where the USERID normally appears. Only the Action Card is required.

6.1.17 Turn Idle Reports Off (Action Code IDLE). This option will turn off histograms dealing with idle CPU information (i.e., report IDs 16, 19, 20, 21, 22, 31-34, 43-54, 55-58). The user should realize that these reports are useful in determining the I/O boundness of the system. However, on most systems, the idle trace is 70 percent of the entire tape, so that, by turning off this processing, processing time can be reduced by over 50 percent. Only the action card is required for this option.

6.1.18 Change Excessive Resource Limits Used in Excessive Resource Report (Action Code WASTED, CORE, IO, CPU, RATIO and URG). This report lists all jobs which are above a preset threshold for any of the following resources:

Wasted Memory
Excessive Memory
Excessive CPU time
Excessive I/O time
Excessive Ratio
Excessive Urgency

These limits are currently set to the values specified in table 6-4 and may be changed by using this option. The format for this option consists of Card 1 specifying the action code and Card 2 specifying the new threshold limit. This report is explained later in this chapter.

6.1.19 Eliminate SNUMBs From Abort Report (Action Code ABORT). This report lists all activities that fail to go to EOJ (i.e., Abort). The details of the report are given in section 6.3 of this chapter. At times, jobs are designed in such a way that they can be terminated only via a MME GEBORT or operator command. While these jobs do not go to EOJ, they have processed correctly and have not resulted in wasted computer resources. This option allows the user to request that these jobs not be included in the Abort Report. The first card contains the action code ABORT. The second card contains the number of jobs that will be deleted from the Abort Report. This number may not exceed 10. If more than 10 jobs are listed, only the first 10 will be deleted. The third card contains the SNUMB of each job to be deleted. Each SNUMB must be followed by at least one blank column.

6.1.20 Change the Plot Interval (Action Code PLTINT). Currently, all plots are outputted at 10-minute intervals. The plot interval controls the output of all plots; i.e., one plot cannot have a different time interval than another plot. The first card of this option contains the action code PLTINT. The second card contains the new plot interval inputted in minutes.

6.1.21 Change the Program Number for the First Slave Job (Action Code FSTSLV). In the GCOS system, certain program numbers are assigned to system jobs. For example $CALC is program number 1, $PALC is program number 2, $SYOT is program number 3, etc. In the WWMCCS system, all programs with a program number less than 14 (decimal) are considered system programs. This option allows the user to alter this program number from its default value of 14. The first card contains the word FSTSLV and the second card contains the new program number. For non-WWMCCS systems, FSTSLV should normally be set to 8.

6.1.22 Request that Certain Jobs be Considered System Jobs (Action Code MASTER). There are certain jobs executed during the course of a day which have program numbers that would designate these jobs as user jobs (see subsection 6.1.21). However, in actuality they are system jobs and should be considered as system overhead. Examples of such jobs are VIDEO, HEALS, the GMF MONITOR, etc. This option allows the user to define up to ten jobs that should be considered as system jobs. The first card contains the Action Code MASTER. The second card contains the number of jobs to be defined as system jobs. The third card contains the SNUMB of each job to be considered as a system program. Each SNUMB must be followed by at least one blank column. It should be noted that VIDEO, $HEALS, the GMF Monitor and $TRAX will automatically be reported as system jobs and do not need to be requested via this option.

6.1.23 Allocation Status Report Print Control (Action Code PALC). Due to the excessive amount of output possible from the Allocation Status report, a time control can be set to print only those activities that are in any allocation state greater than the time limit. This time limit defaults to 600 seconds (10 minutes). The first card contains the word PALC and the second card contains the new time limit, in seconds.

6.1.24 Request the Special Job Memory Reports (Action Code SPECL). If the analyst desires to track the memory demands for a specified number of jobs (not to exceed ten), this input option should be invoked. This option will cause two reports to be produced. One report will indicate every time the requested job(s) was swapped or issued a MME GEMORE/GEMREL for memory, how long it was swapped, or how long the GEMORE was outstanding, and how much memory the job(s) required. In addition, every time the special job issues a MME GEMORE, a line from the Memory Map Report will be generated. This line is generated by default and is not dependent upon whether or not the Memory Map Report is enabled. When the analyst wants to match the Memory Map output to the Special Job output, he must do so based on the time value. For example, if the Special Job Report indicates that FTS issued a MME GEMORE at 16.81057, the user would then examine the Memory Map for a line of output with a time smaller than 16.81057, but where the time on the next line is greater than or equal to 16.81057. For example, the Memory Map might have a line of output with a time indication of 16.81052 where the next line of output was 16.81065. In this case, the line of output at 16.81052 shows what memory looked like at the instant in time that FTS issued the MME GEMORE. If the Special Job

Report indicates that FTS was swapped after issuing the MME GEMORE, the analyst could examine the Memory Map in order to determine why FTS was forced to swap.

A line of the Memory Map is also generated every time the GEMORE for the special job was denied or the special job was forced to swap in order for the GEMORE to be satisfied. This line of the Memory Map would show what memory looked like when the special job was denied the memory request or was swapped from memory. A final line of the Memory Map is produced whenever the special job's memory demand was met. For the swap/denied case and the memory-met case, the Special Job Report and Memory Map are matched by locating identical time values on each report. By generating the Memory Map, the analyst can determine if there are certain jobs that are preventing other jobs from acquiring required memory resources. In this case, the Special Job Report and Memory Map Report can be correlated by matching up the time values from both reports with the identical time values. This is especially useful in an analysis of the Timesharing Subsystem or the File Transfer System.

A second report will also be produced which indicates the average memory size of the job(s) during the course of its execution. This average is taken over increments of time where the time increment used, is the same increment that is used to produce the series of plots. The option consists of three cards where the first card contains the word SPECL, the second contains the number of jobs to be analyzed, and the third card contains the list of SNUMBs separated by at least one blank column.

6.1.25 Process Data on a WW6.4 System (Action Code RN). If the data reduction program is to be run on a WW6.4 system, the user must use this input option. It consists of the letters RN typed on a data card.

6.1.26 Produce a Memory Map Only Under Certain Memory Demand Conditions (Action Code MAPART). Due to the enormous amount of output generated by the Memory Map and Out of Core Reports, it is suggested that a site not produce these reports as a standard procedure. However, these reports are very useful in that they do provide a complete picture of memory as well as a total list of all jobs waiting for memory. In order to provide an analyst with the capability of obtaining these reports, without being buried in computer output, this option has been designed. When used, this option states that a line of the Memory Map and Out of Core Reports should be generated only when the number of activities waiting for memory surpasses a certain limit. To invoke this option, a two-card format is required. Card 1 contains the word MAPART and card 2 contains the number of activities that must be waiting memory before a line of output will be generated for the Memory Map and Out of Core Reports.

6.1.27 Include Zero Urgency Jobs in all Memory Calculations (Action Code ZERO). In several of the reports produced by the MUM, jobs with zero urgencies are not included in the statistical calculations. With the use of this option, zero urgency jobs will be included in all calculations.

6.1.21 Change the Program Number for the First Slave Job (Action Code FSTSLV). In the GCOS system, certain program numbers are assigned to system jobs. For example $CALC is program number 1, $PALC is program number 2, $SYOT is program number 3, etc. In the WWMCCS system, all programs with a program number less than 14 (decimal) are considered system programs. This option allows the user to alter this program number from its default value of 14. The first card contains the word FSTSLV and the second card contains the new program number. For non-WWMCCS systems, FSTSLV should normally be set to 10.

6.1.22 Request that Certain Jobs be Considered System Jobs (Action Code MASTER). There are certain jobs executed during the course of a day which have program numbers that would designate these jobs as user jobs. However, in actuality they are system jobs and should be considered as system overhead. Examples of such jobs are VIDEO, HEALS, the GMF MONITOR, etc. This option allows the user to define up to ten jobs that should be considered as system jobs. The first card contains the Action Code MASTER. The second card contains the number of jobs to be defined as system jobs. The third card contains the SNUMB of each job to be considered as a system program. Each SNUMB must be followed by at least one blank column.

6.1.23 PALC Report Print Control (Action Code PALC). Due to the excessive amount of output possible from the PALC report, a time control can be set to print only those activities that are in any PALC state greater than the time limit. This time limit defaults to 600 seconds (10 minutes). The first card contains the word PALC and the second card contains the new time limit, in seconds.

6.1.24 Request the Special Job Memory Reports (Action Code SPECL). If the analyst desires to track the memory demands for a specified number of jobs (not to exceed ten), this input option should be invoked. This option will cause two reports to be produced. One report will indicate every time the requested job(s) was swapped or issued a MME GEMORE for memory, how long it was swapped, or how long the GEMORE was outstanding, and how much memory the job(s) required. A second report will also be produced which indicates the average memory size of the job(s) during the course of its execution. This average is taken over increments of time where the time increment used, is the same increment that is used to produce the series of plots. The option consists of three cards where the first card contains the word SPECL, the second contains the number of jobs to be analyzed, and the third card contains the list of SNUMBs separated by at least one blank column.

6.1.25 Process Data on a WW6.4 System (Action Code RN). If the data reduction program is to be run on a WW6.4 system, the user must use this input option. It consists of the letters RN typed on a data card.

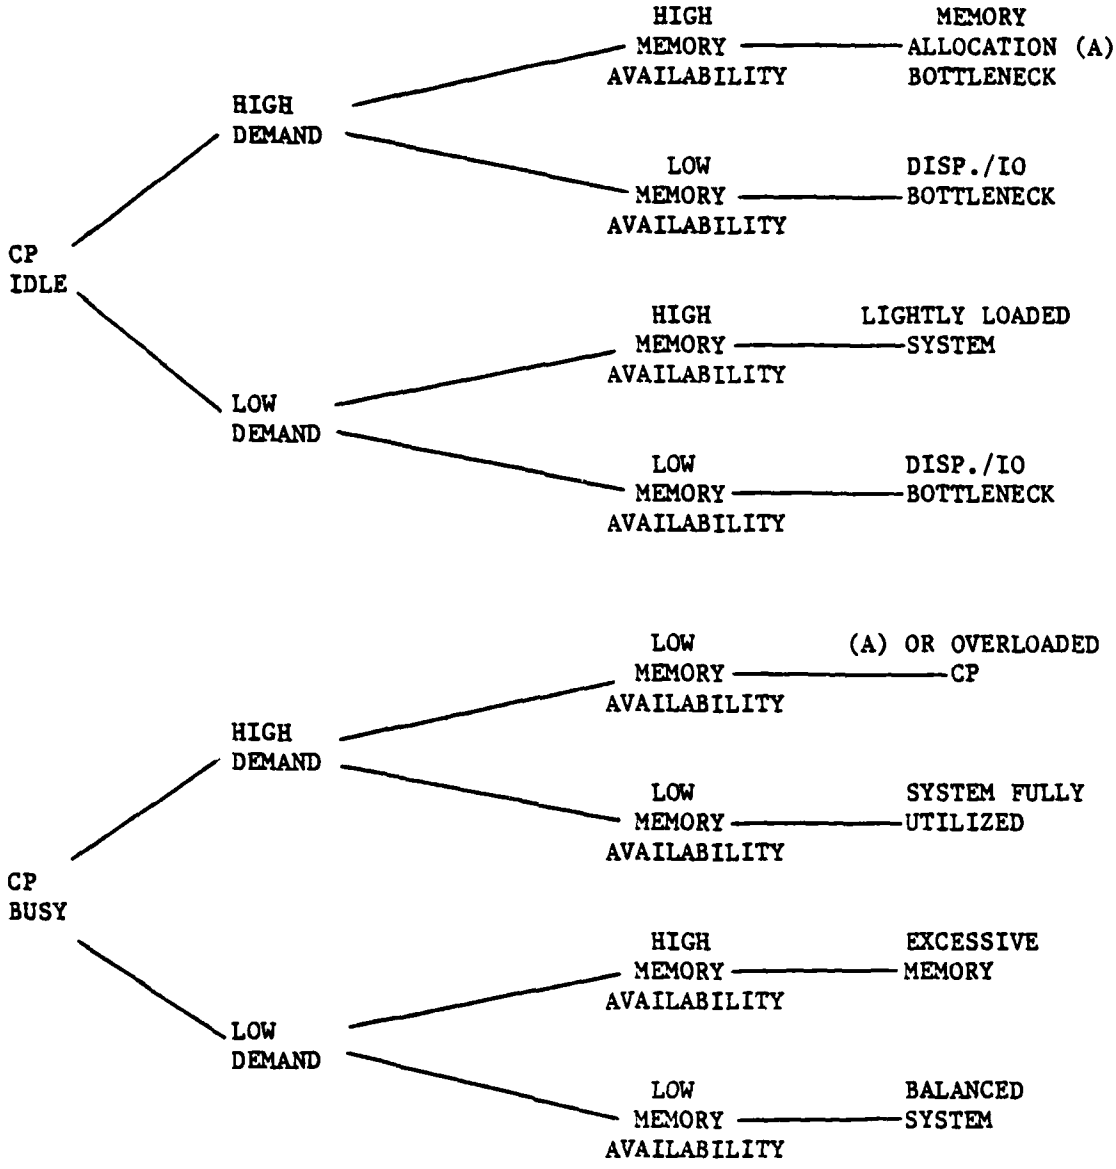


Figure 6-7. System Bottleneck Chart

6.2 Processing

6.2.1 General. The reports of the MUDRP are intended to aid in the following:

- o System sizing - both memory sizing and processor utilization.
- o Job flow analysis - determining if and where a bottleneck exists and the user memory loading and the daily load distributions.
- o System perturbation measures - allows the user to evaluate how a new procedure or new load may alter the utilization of the system as well as determine the total utilization/capacity of the system.
- o Large user jobs - aid in identifying specific jobs which are misusing or "hogging" system resources.

Figure 6-7 illustrates how the monitor will pinpoint these various areas. For example, if the monitor indicates a large percentage of processor idleness with high memory demand and low memory availability, a dispatching or I/O bottleneck would be indicated. This would be caused by the I/O not completing its services in a sufficiently timely manner to allow full use of the processors. If processor use was very high and memory demand and availability were high, a memory allocation bottleneck or an overloaded processor would be indicated.

6.2.2 JCL. Figure 6-8 presents the JCL needed to run a total MUM reduction. The following points describe key features of the required JCL.

- o 62K required for memory
- o 15K sysout requirement would vary depending on amount of data collected. This figure would be significantly higher if the Memory Map or Out of Core Report were produced.
- o The DATA I* card is used to indicate the presence of data cards. All data cards must immediately follow this card. At least one data must be present. That card will contain the word END and is used to signify the end of input data. The END card must be present even if no other data cards are desired. The data cards shown in the example are those recommended for most analyses.
- o An additional 12K will be required to load the MUM reduction program, but this 12K will be released immediately upon loading.

```
$   IDENT          1820251/30/3044,C702
$   SELECT         B29IDPXC/OBJECT/MUM
$   TAPE           01,X1D,,18897
$   LIMITS         999,62K,-4K,15K
$   DATA           I*
    DATA cards
    ALLON
    14
    5 6 7 12 13 14 15 17 18 29 30 46 39 47
    SPECL
    1
    TS1
    END
```

Figure 6-8. JCL to RUN MUDRP

Table 6-5 shows all the MUDRP file codes and their corresponding reports.

6.3 Outputs

In this section, a simple explanation of how each report was derived from the data is given. Subsection 6.1 discussed how the ranges and other options of each report may be modified to fit an individual installation. While this section will provide some insight as to how an analyst should proceed to review all the reports produced by the MUDRP, section 14 provides a step-by-step approach as to how a memory analysis might be conducted.

Immediately prior to the output of the histograms, the user will find a printout containing processing information. Included in this information is the following:

- Printout of all input options selected by user
- Indication of multireel tapes that are being requested and have been mounted
- Indication of the monitors that were active during data collection
- Error messages - all error messages are either self-explanatory or else followed by the words "For Information Only." The latter messages are used by CCTC for future enhancements and as such can be ignored by the user.
- If the time frame option was used, and indication of when the various time frames were reached.

6.3.1 MUM Title Page. The Memory Utilization Monitor (MUM) title page contains a summary of the systems configuration and activity over the measurement period (see figure 6-9). It displays the time the monitor was initiated and terminated, as well as identifying the system which was monitored and the tape number(s) containing the data. The configuration information is augmented by the amount of memory dedicated to the operating system itself, including that used by the memory allocation program. These figures will give the user a good idea of how much hard core space remains and could be used for SSA module hard core loading. If SSA cache is also configured, the amount of memory being used for this feature is also listed. The version number should be 01-82.

Immediately following is a summary of the work processed over the measurement period. The first set of lines provides information concerning the overhead generated by the actual data collection. The monitor name is given, its CPU time in seconds, an its overhead as a function of total processor power. The GMF executive overhead is separated from the actual monitors and is listed as "EXEC". The monitor "NAME" is an area of code within the Mass Store Monitor and even though

Table 6-5. File Code for MUM Reports
(Part 1 of 2)

20	Activity Resource Report, Special Job Reports
21	IDENT Report
22	Special Job Report (temporary file)
23	Special Job Report (temporary file)
24	Urgency Over Time Report (temporary file)
26	Zero Urgency Job Report (temporary file)
27	Activity Abort Report
31	Plot 1 - (see table 6-1 for Plot Definition) (temporary file)
32	Plot 2 - (see table 6-1 for Plot Definition) (temporary file)
33	Plot 3 - (see table 6-1 for Plot Definition) (temporary file)
34	Excessive Resource Report
35	Plot 4 - (see table 6-1 for Plot Definition) (temporary file)
36	Used for outputting all plots
37	Used for outputting Out of Core Report, Memory Map, and Peripheral Allocator Report
42	Histograms, System Program Usage Report, Memory Statistics Report, Distribution of Urgency Over Time Report, Zero Urgency Job Report
45	Out of Core Report (temporary file)
51	Memory Map Report with one file required for each 128K Memory configured (temporary file)
52	Memory Map Report with one file required for each 128K Memory configured. (temporary file)
53	Memory Map Report with one file required for each 128K Memory configured (temporary file)
54	Memory Map Report with one file required for each 128K Memory configured (temporary file)

Table 6-5. (Part 2 of 2)

55	Memory Map Report with one file required for each 128K Memory configured (temporary file)
56	Memory Map Report with one file required for each 128K Memory configured (temporary file)
57	Memory Map Report with one file required for each 128K Memory configured (temporary file)
58	Memory Map Report with one file required for each 128K Memory configured (temporary file)
59	Demand List Report (temporary file)

THIS PAGE LEFT INTENTIONALLY BLANK

****** THE MEMORY UTILIZATION MONITOR ******

VERSION 01-82

MONITORING ON 79-05-04 STARTED AT 11:39:50 AND COMPLETED AT 13:19:56 FOR A TOTAL TIME OF 1.67 HOURS
ON SYSTEM OSCC2 RUNNING 6.4.1D OF TAPE D0002

THE SYSTEMS CONFIGURATION CONSISTED OF:

2 - 6680 CENTRAL PROCESSORS
2 - INPUT/OUTPUT MULTIPLEXORS
WITH 24 I/O CHANNELS
512 - 1024 WORD BLOCKS OF MEMORY
52 OF WHICH WERE USED BY THE HARD CORE SYSTEM ITSELF
3 OF WHICH WAS USED BY CALC
5K USED BY SSA CACHE (NOT HARD CORE)

THE SYSTEM PROCESSED THE FOLLOWING OVER THE MEASUREMENT SESSION:

MONITOR	TIME(SECS)	% OVERHEAD
EXEC	100	1.7
MUM	60	1.0
CPU	40	.7
IDLE	40	.7
TOTAL		4.1

175 ACTIVITIES WERE PROCESSED AT A RATE/HOUR OF 104.89
OF THESE 50 WERE SYSTEM SCHEDULAR ACTS
125 ACTUAL ACTIVITIES WERE PROCESSED AT A RATE/HOUR OF 74.85

128 MOVES WERE PERFORMED AT A RATE/HOUR OF 76.72
266 SWAPS WERE PERFORMED AT A RATE/HOUR OF 159.44
201466 TIMES THE PROCESSORS WENT IDLE YIELDING A
20 % IDLENESS OF THE PROCESSORS
THE FIRST PROCESSOR WENT IDLE 122479 TIMES, YIELDING 23 % IDLENESS
THE SECOND PROCESSOR WENT IDLE 78987 TIMES, YIELDING 17 % IDLENESS

THE MEMORY ALLOCATOR WAS CALLED 1907 TIMES - 0 OF WHICH RESULTED IN NO STATE CHANGE

THE TOTAL CPU TIME IN SECS WAS 6991 THE TOTAL IO TIME IN SECS WAS 15662 CPU/IO RATIO IS 0.446373
WEIGHTED MEMORY SURPLUS IN K WORDS WAS 46
WEIGHTED MEMORY SHORT-FALL IN K WORDS WAS 50 INCLUDES CALC AND PALC QUEUES

Figure 6-9. MUM Title Page Report - Idle Monitor Active

listed separately it is also included under the monitor "MSM". The Monitor "FMS" is also an area of code within the Mass Store Monitor, but in this case it has not been included under the monitor "MSM". These two special areas of code, within subroutine T7 (connect trace processing), are considered to be high usage areas and as such consume significant processing resources. In order to determine the true overhead of these areas, so that future code optimization can be considered, these areas are being reported separately.

Monitor "CM" in this report describes the processor overhead of subroutine T4 (terminate processing) and subroutine T22 (start I/O processing). Monitor "MSM" in this report describes the processor overhead of subroutine T7 (connect processing). Therefore, if the Channel Monitor was active, but the Mass Store Monitor was not, this report will still list both "CM" and "MSM" as contributing to the processor overhead. The total Channel Monitor overhead will be found by adding the overhead of the "CM" monitor, to the overhead of the "MSM" monitor, to the overhead of the "FMS" monitor.

If both the Channel Monitor and Mass Store Monitor were active, then the combined overhead of both monitors can be found as the sum of "MSM" + "CM" + "FMS".

For purposes of this report, % overhead is computed as

$$\frac{\text{(CPUTIME used by monitor)}}{\text{(Total Elapsed Time) x (number of Processors)}}$$

Following this are several lines describing the work performed during the monitoring session. These lines are self-explanatory.

If a termination record is not processed, either because the monitor aborted before a termination record could be written or else time frames were used, the lines describing GMC overhead will not be printed.

The number of times a processor went idle is derived from the idle processor traces captured by the IDLEM, with the percentage of processor idle also being gathered by the collection of idle state information. This is shown system-wide (i.e., for all the central processors and then individually for each processor). This information will not be present if the IDLEM was not active or if its output reports have been disabled by a data card option (see figure 6-10).

The number of memory allocator calls, as counted by the monitor, is shown. This much less than the number of calls to the multitude off SSA modules used by the Core Allocator and consists only of those that may have altered the memory state of the system. The second figure shows how many times a memory state change might have taken place and did not. This could be caused by no allocation being possible or by a call to the allocator pertaining to a matter other than allocation (i.e., a console message).

****** T H E M E M O R Y U T I L I Z A T I O N M O N I T O R ******
VERSION 01-82

MONITORING ON 80-12-15 STARTED AT 12:39:46 AND COMPLETED AT 19:29:06 FOR A TOTAL TIME OF 6.82 HOURS
ON SYSTEM NMCC2 RUNNING W64000 OF TAPE D0001

THE SYSTEMS CONFIGURATION CONSISTED OF:

2 - 6680 CENTRAL PROCESSORS
2 - INPUT/OUTPUT MULTIPLEXORS
WITH 16 I/O CHANNELS
512 - 1024 WORD BLOCKS OF MEMORY
52 OF WHICH WERE USED BY THE HARD CORE SYSTEM ITSELF
3 OF WHICH WERE USED BY CALC
5K USED BY SSA CACHE (NOT HARD CORE)

THE SYSTEM PROCESSED THE FOLLOWING OVER THE MEASUREMENT SESSION:

MONITOR	TIME(SEC)	% OVERHEAD
EXEC	300	1.2
MUM	200	.8
TOTAL		2.0

1033 ACTIVITIES WERE PROCESSED AT A RATE/HOUR OF 151.41
OF THESE 33 WERE SYSTEM SCHEDULAR ACTS
1000 ACTUAL ACTIVITIES WERE PROCESSED AT A RATE/HOUR OF 146.63

1660 MOVES WERE PERFORMED AT A RATE/HOUR OF 243.31
6363 SWAPS WERE PERFORMED AT A RATE/HOUR OF 932.66
IDLE MONITOR WAS NOT ACTIVE OR WAS TURNED OFF

THE MEMORY ALLOCATOR WAS CALLED 23409 TIMES - 14 OF WHICH RESULTED IN NO STATE CHANGE

THE TOTAL CPU TIME IN SECS WAS 31022 THE TOTAL IO TIME IN SECS WAS 61341 CPU/IO RATIO IS 0.505735
WEIGHTED MEMORY SHORT-FALL IN K WORDS WAS 101 INCLUDES ONLY CALC QUEUE
WEIGHTED MEMORY SHORT-FALL IN K WORDS WAS 150 INCLUDES CALC AND PALC QUEUES

Figure 6-10. MUM Title Page Report - Idle Monitor Off

The next line printed out is the Total CPU and I/O times in seconds and the ratio of CPU to I/O time. This figure gives the user an idea of whether the workload processed by the system is I/O or CPU dominant. It should be noted that these numbers are the amount of CPU and I/O time generated during the measurement period.

The next two lines give an indication of whether the system has a surplus or shortfall of memory. The weighted figure is calculated by using the following formula:

$$W = \frac{\sum_{i=1}^{N} \left(\begin{matrix} \text{memory} \\ \text{available} \end{matrix} - \begin{matrix} \text{demand for} \\ \text{memory} \end{matrix} \right) * \left(T_{i+1} - T_i \right)}{\text{TOTAL TIME}}$$

Where i = calls to the core allocator

T T - length of time over which memory availability was in this state.

If W comes out positive, there is a core surplus and if W comes out negative, there is a core shortfall. In the first line, the demand for memory is taken only from the Core Allocator's queue. In the second line the demand for memory is taken from the demand in both the Core Allocator and Pripheral Allocator queues. The Peripheral Allocator's queue consists of the memory demand that is currently being processed by the Peripheral Allocator and has not yet reached the Core Allocator. The Peripheral Alloctor will stop transferring jobs to the Core Allocator when the Core Allocator's queue reaches a predefined length. This second figure presents a truer picture of memory availability. Jobs from the Peripheral Allocator are only included if they have been completely processed by the Peripheral Allocator. These figures present a good first indication of whether or not availability of memory is a system constraint. In calculating demand, a job is only included if it does not have a zero urgency. Any activity with a zero urgency will not be considered to have a core demand unless the activity is in a loading (activity 0) or terminating status.

6.3.2 System Program Usage. The report immediately following the title page provides an overview of the system program load on the memory subsystem. The data presented consists of the following:

- o $\frac{\text{Total Memory Time for This System Program} * 100}{\text{Memory Time for all Programs}}$

This figure would indicate what percentage of the total memory time was used by this program.

- o Percentage of the Elapsed Time in Memory
- o $\frac{\text{Total Size Time Product for This System Program} * 100}{\text{Total Size Time Product for all Programs}}$

This figure would indicate what percentage of the total size time produce was used by this program. The size-time product of a job is an attempt to determine the memory effect of a job based not only on its size, but on the length of time that it runs. A 20K jobs that runs for three hours might be more detrimental to a system than a 60K jobs that runs ten minutes.

$$\text{o} \quad \frac{\text{Total Size Time Product for This System Program} * 100}{\text{Total Size Time Produce Available to System}}$$

Where Total Size Time Produce Available to System = (The Elapsed Run Time) * (Total Allocatable Memory)

The next two figures are weighted memory sizes for this program. The first figure is the weighted memory size of this program while it is in memory. Therefore, if TSS was in memory during three different time periods for 1/2 hour, 3/4 hour, and 1 hour, and during these periods its memory size was 40K, 100K, 180K respectively, its weighted in memory size would be calculated as follows:

$$\text{Weighted (IN)} = \frac{(40)*(.5)+(100)*(.75)+(180)*(1)}{2.25} = \frac{275}{2.25} = 122K$$

Had the calculation not been weighted by time, the average size of TSS would have been:

$$\frac{(40)+(100)+(180)}{3} = 73K$$

In the above calculation, the report would be stating that the amount of memory being taken away from the system, by TS, was 122K. However, what if TSS was swapped for 50 percent of the total elapsed time? Then TSS really did not take 122K from the system, but rather only 61K. The second weighted figure takes into accounte the total time the program was actually in memory.

The final figure is the number of times this program was swapped.

In addition to the standard system programs, any jobs requested by the user, to be considered as system jobs, also appear in this report. In figure 6-11 we see 6 user-requested jobs appearing on the report. The user had actually requested nine jobs to be considered as system jobs, but three of those jobs never appeared. In a system using multicopies of TSS only TS1 (prog #5) will appear in this report. Other copies of TSS must be requested by user input option "MASTER". In a WWMCCS system, a program is considered to be a system program if it has a program number less than decimal 14. Commercial users should use the FSTSLV option to change the

THE SYSTEM PROGRAM USAGE OF MEMORY WAS:

	PROGRAM	% USED OF MEMORY USE TIME	% OF TIME IN MEMORY	% USED OF SIZE-TIME USED OF MEMORY	% USED OF SIZE-TIME AVAILABLE	WEIGHTED IN/TTM	# SWAPS
(001)	GMF MONITOR	7%	100%	5%	5%	23/ 24	0
(002)	PERIPHERAL ALLOCATOR	5%	71%	2%	2%	17/ 12	988
(003)	GENERAL SYSTEM OUTPUT	6%	93%	6%	5%	26/ 25	58
(004)	REMOTE INPUT	2%	35%	1%	1%	16/ 5	415
(005)	TIME-SHARING EXECUTIVE	7%	100%	31%	28%	129/129	0
(006)	TEST AND DIAGNOSTICS	0%	1%	0%	0%	6/ 0	0
(010)	LOG ON - SECURITY	2%	38%	1%	0%	11/ 4	292
(011)	FILE MANAGEMENT SYSTEM	0%	12%	0%	0%	9/ 1	0
(012)	WW NETWORK CONTROL (NCP)	6%	88%	5%	5%	27/ 24	79
(013)	WW TELNET	4%	65%	2%	2%	16/ 10	150
(014)	WW T2EXEC	3%	47%	3%	3%	28/ 13	233
(015)	WW TLCF	2%	38%	1%	1%	16/ 6	54
(016)	DMS PROCEDURES (DMTEX)	0%	0%	0%	0%	4/ 0	1
SPEC 1	HEALS	1%	21%	0%	0%	7/ 1	150
SPEC 4	NACE	3%	50%	2%	2%	18/ 9	1652
SPEC 6	VIDEO	7%	99%	1%	1%	6/ 5	2
SPEC 7	LCEN	7%	99%	4%	3%	17/ 16	4
SPEC 8	NEWRM	0%	8%	0%	0%	3/ 0	0
SPEC 9	DMSTA	5%	79%	1%	1%	8/ 6	717
					TOTAL	387/290	

Figure 6-11. System Program Load

value of 14 to an 8. In addition, the GMC Monitor, $HEALS, VIDEO and $TRAX are all considered to be system programs.

6.3.3 MUM Reports. The following paragraphs describe the reports output by MUM.

Report numbers 1-50 are all presented in histogram format (see figure 6-12). At the top of the report, the system name, as well as the time and date of data collection, are given. This is followed by the title line of the histogram. Column number 1 indicates the number of occurrences of a given event, with column number 5 describing the event. In figure 6-12, we find that 229 times there were 0 user activities in memory, while 2899 times there were 5 activities in memory. Column number 2 is simply a running total of column number 1. Therefore, the second line in column number 2 (2008) is merely a running total of the first two lines of column number 1 (229 + 1779). The fourth column is the percentage of all activities which will fall into that line of the report showed two user activities in memory. For example, 4063 entries out of a total 23410 entries. This results in a 17.35 percentage figure. This means that 17.35% of all measurements (4063/23410) showed 2 activities being in memory. Column number 3 is simply a running total of column number 4. It presents the percentage of measurements which will fall into a given line, or earlier line. For example, 25.93% of all measurement showed the number of activities in memory to be 2 or less. There is a graphic display of these measurements presented to the right of the fifth column. At the bottom of the report, summary information is provided and is calculated in the standard statistical manner.

In figure 6-13, we see a similar histogram report. As displayed by column 5, we find that each line of the histogram represents a range of values, with an interval size of 200. This interval size can be modified by the user. The lowest value in this histogram is 0 (modifiable by the user) and the size of the histogram defaults to 45 lines (also modifiable by the user). Actually, for this run, the lowest value recorded was 42. Since we can output only 45 lines and each line represents a range of values of 200, the largest value that could be reported would be 9,000 (200 x 45). If a measurement falls outside this maximum value, it is reported as an out-of-range value. In figure 6-13, we find the 21 measurements exceeded 9,000. The average of these 21 measurements was 20188.48. The first line of the summary includes all measurements that were taken. Therefore, 21 out of 79 entries (26 percent) of all measurements were out of range. the average of all measurements taken was 5953.62, while the average of the in-range masurements, (all out of range values are eliminated) was 799.62.

6.3.3.1 Report 1 - Memory Demand Sizes of New Activities in 1K Word Blocks. This report shows the demand size, in 1K-word blocks, of each individual user activity as it was first seen by the memory allocator. The demand sizes are presented to the allocator by the Peripheral Allocator. This is a good measure of the memory demand load of a systems

DISTRIBUTION COLLECTED ON SYSTEM NMCC2 AT 12:39:46 ON 80-12-15

THE # OF USER ACTIVITIES IN MEMORY

REPORT 15

INDIV. NUMBER	CUMUL. NUMBER	CUMUL. PRC	INDIV. PRC	NUMBER WAIT	PERCENT OF OCCURRENCE 00 05 10 15 20 25 30 35 40 45 50
					I....I....I....I....I....I....I....I....I....I....I
229	229	0.978	0.978	0- 0	IX
1779	2008	8.578	7.599	1- 1	IXXXXXXXX
4063	6071	25.933	17.356	2- 2	IXXXXXXXXXXXXXXXXX
6597	12668	54.114	28.180	3- 3	IXXXXXXXXXXXXXXXXXXXXXXXXXXXX
7062	19730	84.280	30.167	4- 4	IXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
2899	22629	96.664	12.384	5- 5	IXXXXXXXXXXXX
633	23262	99.368	2.704	6- 6	IXXX
134	23396	99.940	0.572	7- 7	IX
14	23410	100.000	0.060	8- 8	I

23410 ENTRIES TOTAL AVERAGE = 3.30145 VARIANCE - 1.622 STANDARD DEVIATION - 1.273

Figure 6-12. Standard Histogram Report

DISTRIBUTION COLLECTED ON SYSTEM NMCC2 AT 18:50:03 ON 81-03-13

THE ELAPSED DURATION OF USER ACT IN TENTHS OF A SECOND

INDIV. NUMBER	CUMUL. NUMBER	CUMUL. PRC	INDIV. PRC	TENTHS SECOND	PERCENT OF OCCURRENCE 00 05 10 15 20 25 30 35 40 45 50	REPORT 20
					I....I....I....I....I....I....I....I....I....I....I	
22	22	27.848	27.848	42- 199	IXXXXXXXXXXXXXXXXXXXXXXXXXX	
14	36	45.570	17.722	200- 399	IXXXXXXXXXXXXXXXXX	
5	41	51.899	6.329	400- 599	IXXXXXX	
2	43	54.430	2.532	600- 799	IXXX	
1	44	55.696	1.266	800- 999	IX	
0	44	55.696	0.	1000-1199	I	
1	45	56.962	1.266	1200-1399	IX	
2	47	59.494	2.532	1400-1599	IXXX	
1	48	60.759	1.266	1600-1799	IX	
0	48	60.759	0.	1800-1999	I	
1	49	62.025	1.266	2000-2199	IX	
3	52	65.823	3.797	2200-2399	IXXXX	
0	52	65.823	0.	2400-2599	I	
0	52	65.823	0.	2600-2799	I	
1	53	67.089	1.266	2800-2999	IX	
1	54	68.354	1.266	3000-3199	IX	
1	55	69.620	1.266	3200-3399	IX	
0	55	69.620	0.	3400-3599	I	
2	57	72.152	2.532	3600-3799	IXXX	
0	57	72.152	0.	3800-3999	I	
0	57	72.152	0.	4000-4199	I	
1	58	73.418	1.266	4200-4399	IX	

79 ENTRIES TOTAL AVERAGE = 5953.62024 VARIANCE = 04446990.000 STANDARD DEVIATION = 10219.931
21(26%) OUT OF RANGE AVERAGE FOR THESE = 20188.48 IN RANGE AVERAGE = 799.62

Figure 6-13. Out-of-Range Histogram

operation and can be used to set System Scheduler classes to correctly balance the load cross varying memory size jobs. In this type of report, the distribution shows the percentage of activities which had a particular memory size. For this report, an entry is made for each new user activity demand at each allocator call. See Report 10 for an explanation of user vs. system activity.

6.3.3.2 Report 2 - The Memory Demand Size of All Demand Types. This report contains the information in Report 1, with the addition of all other individual demand types. These include activities that are swapped or involved in a memory compaction procedure. This report should be similar to Report 1, unless a great amount of GETMORE, GETLEC, or swap operations are performed by the users load. This would alter the memory size demands from that seen by the allocator at the initial request. For this report, an entry is made for each activity with an outstanding demand for each allocator call. Activities with an urgency of 0 are not counted.

6.3.3.3 Report 3 - The Potential Time-Weighted Memory Demand. This report shows the total memory required if all jobs currently in the system were to have their memory demand satisfied at one time. The data shown in this report is the sum of all jobs currently in memory, plus all jobs currently swapped, plus all jobs waiting for their original allocation of memory, plus all jobs currently in the system with a zero urgency. This report is an indication of the worst case demand that may be potentially placed upon the system. An entry is made to this report at each allocator call. See report 16 for an explanation of time correction.

6.3.3.4 Report 4 - The Demand That Was Outstanding When a Processor Went Idle. This report shows the sum of demand for all activities in the system including outstanding GETMOREs. It is a distribution of memory demand that is not satisfied, across the measurement session. It should be remembered that all data is collected at the Core Allocator and does not represent the full system load. Portions of the load may be held in the System Scheduler and the Peripheral Allocator. Activities with an urgency of 0 are not counted. An entry is made only if a processor has gone idle since the last allocator call. If a large demand should be outstanding during processor idleness, a system bottleneck may be present. In this case, memory is probably fully utilized (i.e., demand cannot be satisfied), but the activities that are occupying memory are not using the processor, (i.e., a processor has gone idle). This is a good sign of an I/O backlog. IDLEM data is used to produce this report. If the Idle Monitor was not active, this report will not be produced.

6.3.3.5 Report 5 - The Total Amount of Available Memory. The total amount of available memory is a key indicator of the system memory utilization. If this amount is continually low, the memory is being fully utilized and possibly in need of expansion. A continually high amount may indicate another system bottleneck or an excess of memory. This report, when used in conjunction with Reports 3, 4, and 6 should give a good first-level indication of system memory utilization. It should be noted that the availability shown here exists in all quadrants. The availability is the sum of any and all "holes" in the system and does not mean that this memory is contiguously available.

Any activity with an urgency of 0 that is currently in memory will have its memory size included in this availability figure. The reason for this is that if memory becomes a constraint, these activities can be swapped and their memory will become available for use.

For this report, an entry is made for each allocator call. For most analyses, this report will not be used since report 8 provides a more statistically accurate representation of this data.

6.3.3.6 Report 6 - The Memory Available When a Processor Went Idle. The previous report is repeated with the additional restraint that a processor has gone idle since the last allocator call. This aids in identifying either a bottleneck or a lightly loaded system.

For this report, an entry is made at each allocator call that had a processor go idle since the last allocator call. IDLEM data is used to produce this report. This report will not be produced if IDLEM was not active or the IDLEM Reports have been disabled via user input command.

6.3.3.7 Report 7 - The Time-Corrected Total Demand Outstanding. See report 16 for an explanation of time correction. The time-corrected total demand is the sum of all requests for memory known to the allocator as indicated in report 4. Activities with urgency 0 are not counted.

6.3.3.8 Report 8 - The Time-Corrected Memory Available. See report 16 for an explanation of time correction. This report reflects the time-corrected amount of total memory available as indicated in report 5.

6.3.3.9 Report 9 - The Number of Activities Waiting for Memory in Allocator Queue. This report identifies the depth of the allocator demand queue and includes all activites that are waiting for memory allocation. Activities with a 0 urgency are not considered as waiting for memory. This report aids in determining if too many or too few activities are getting to the Core Allocator from the Peripheral Allocator. For this report, an entry is made at each allocator call. For most analyses, this report will not be used since report 11 provides a more statistically accurate representation of this data.

6.3.3.10 Report 10 - The Number of User Activities Waiting Memory in Allocator Queue. This report is the same as report 9 except that it only counts those activites of a slave job as identified by their program number (program number 14 or greater). In order to change this program number test, the user should see Input Action FSTSLV. In addition, the user may specify up to ten additional programs that he wants considered as system programs, even though their program number exceeds 14. The user

should see Input Action MASTER in order to select this option. This report indicates the "user" work waiting allocation. For this report, an entry is made on each allocator call. For most analyses, this report will not be used since report 12 provides a more statistically accurate representation of this data.

6.3.3.11 Report 11 - The Time-Corrected Number of Activities Waiting Memory. See report 16 for an explanation of time correction. This report indicates the time-corrected number of activites waiting memory as in report 9.

6.3.3.12 Report 12 - The Time-Corrected Number of User Activities Waiting Memory. See report 16 for an explanation of time correction. This report indicates the time-corrected number of user jobs waiting memory in the allocators queue as in report 10. See report 10 for additional user options.

6.3.3.13 Report 13 - The Number of Activities Waiting Memory When a Processor Went Idle. Report 9 is the basis for this report, with the additional criteria that a processor must have gone idle since the last allocator call. An entry is made for each allocation where a processor has gone idle since the last call. IDLEM data is used to produce this report. This report will not be produced if IDLEM is not active or the IDLEM reports were disabled via user input commands.

6.3.3.14 Report 14 - The Number of Activites Residing in Memory. This report represents the number of activities allocated memory. It indicates the multiprogramming depth the system is obtaining. It is probably an upper level since an activity is allocated memory prior to and past actual usage. Any activity in memory, with a 0 urgency, is not considered as residing in memory. For this report, an entry is made for each allocator call. For most analyses, this report will not be used since report 16 provides a more statistically accurate representation of this data.

6.3.3.15 Report 15 - The Number of User Activities in Memory. The activites shown in this report are those that are in memory and have a program number greater than or equal to 14. These are user programs. For this report, an entry is made at each allocator call. As explained in report 14, any activity which has an urgency of zero will not be counted as being in memory. See report 10 for additional user options in defining system jobs and user jobs. For most analyses, this report will not be used since report 17 provides a more statistically accurate representation of this data.

6.3.3.16 Report 16 - The Time-Corrected Number of Activities in Memory. This report presents the same information as in report 14. The number of entries at each allocator call is determined by the time since the last allocator call. The result is a simulation of a uniform sample rate of allocator calls. Therefore, the noncorrected reports display the distributions as seen by the allocator itself. The time-corrected reports present the time weighted distributions. As an example assume that three

measurements are taken. It is found that 6 activities are in memory for 2 minutes, 20 activites for 5 minutes, and 8 activities for 1 minute. The average number of activites in memory is (6+20+8)/3=11. If we correct for time however, we get ((6)*(2)+(8)+(20)*(5))/8=100/8=12.5 activites in memory. The division of 8 was the total time (5+2+1) spent collecting data. All of the time-corrected reports are of the same nature.

6.3.3.17 Report 17 - The Time-Corrected Number of User Activities in Memory. This report indicates the time-corrected number of user jobs with allocated memory as in report 15. See report 10 for additional user options in defining system jobs and user jobs. See report 16 for a definition of Time Correction.

6.3.3.18 Report 18 - The Number of Activities in Memory When a Processor Went Idle. This report indicates the total number of activities with allocated memory when a processor went idle. This report can show an I/O bottleneck if the multiprogramming depth is high but there is no work for a procesor to perform. For this report, an entry is made on each allocator call for which a processor went idle since the last call. IDLEM data is used to produce this report. This report will be not be produced if IDLEM is not active or if IDLEM reports have been disabled by user input options.

6.3.3.19 Report 19 - The Ratio of User Activity Duration Versus Its Memory Use Time. This report indicates the ratio of total elapsed time (report 20) over the total allocated memory time (report 21). This shows how activity run time is stretched due to memory resource contention. If there was no memory contention then the total elapsed time of an activity would be very close to total memory time of an activity. As memory contention increases (i.e., swapping occurs), the total elapsed time will begin to increase in comparison to total memory time. It must be realized that total memory time can increase if there is CPU or I/O contention. As a job sits in memory, waiting on the CPU or I/O subsystems, its memory time will increase.

For this report, an entry is made for each user activity that terminates. See report 10 for an explanation of user vs. system activities.

6.3.3.20 Report 20 - The Elapsed Duration of User Activity in 10ths of a Second. This report presents the clock time that the allocator knew of a user activity's existence, measured from its first memory demand to its termination. This includes all time spent in a GEWAKE, in memory, and swapped.

For this report, an entry is made for each user activity that terminates. See report 10 for an explanation of user vs. system activities.

6.3.3.21 Report 21 - The Total Elapsed Time a User Activity Was in Memory. This report shows the duration of elapsed clock time each user

activity had memory allocated to it. It helps describe the system workload requirements.

For this report, an entry is made for each user activity that terminates. See report 10 for an explanation of user vs. system activities.

6.3.3.22 Report 22 - The GEMORE Service or Denial Time - 1/10 Second, Elapsed. The time from a GEMORE request until the activity is allocated the extra memory, swapped to achieve the aditional memory, or denied the memory is displayed in this report.

For this report, an entry is made for each activity whose GEMORE request is no longer present. This report is not used in most analyses.

6.3.3.23 Report 23 - The Request Size of GEMOREs. All GEMORE requests are shown in this report with the displayed size in 1K blocks.

For this report, an entry is made for each GEMORE request. This report is not used in most analyses.

6.3.3.24 Report 24 - Delay Time in the System Scheduler. The amount of time a job spends in one of the scheduler queues is displayed in this report. The Allocation Status Report can be used to display particular jobs that are delayed for excessive amounts of time.

6.3.3.25 Report 25 - Delay Time Until Core Allocation. This report displays the total amount of time activities spent in the various allocation phases prior to core allocation. The Allocation Status Report can be used to display particular activities that are delayed for excessive periods of time.

6.3.3.26 Reports 26 through 31 - The Elapsed Time of a Busy State of the Processors. These reports present the elapsed clock time between the idle states of each individual processor. The reports supply an indication of how each processor is utilized versus the others in the system.

For these reports, an entry is made at each idle state of a processor. IDLEM data is used to produce these reports. These reports will not be produced if the IDLEM was not active or if the IDLEM reports have been disabled by user input option.

6.3.3.27 Report 32 - The Elapsed Time of a Busy State of Processors. The elapsed clock time between idle states of all processors is presented in this report.

For this report, an entry is made for each processor idle state. IDLEM data is used to produce this report.

6.3.3.28 Report 33 - Elapsed Time Between Allocator Calls in 1/100 of a Second. This report shows the elapsed clock time between calls to the allocator and shows if the allocator is receiving sufficient service.

For this report, an entry is made for each user activity that terminates. This report will not be used for most analyses.

6.3.3.29 Report 34 - The I/O Time Charged per User Activity in Seconds. This report indicates the I/O time charged to each user activity.

For this report, an entry is made for each user activity that terminates.

6.3.3.30 Report 35 - The CP Time Charged per User Activity in Seconds. This report presents the CP time charged to each user activity. For this report, an entry is made for each user activity that terminates.

Reports 34 and 35 report the total CPU and I/O times used by a user activity while the monitor was active. These histograms are not generated for programs with program numbers less than 14 (i.e., system programs). See report 10 for additional user options in defining system activities and user activities.

6.3.3.31 Report 36 - The Number of Times a User Activity was Swapped. This report shows the swap count per user activity. The total number of swaps a user activity incurs is the user argument, as counted by the monitor. See report 10 for additional user options in defining system activities and user activities.

For this report, an entry is made for each user activity that terminates.

6.3.3.32 Report 37 - The Total Elapsed Time a User Activity was Swapped. This report indicates the total time a user activity was inactive due to a swap. After each swap is completed, an accumulator is updated, and if an activity is terminated, an entry is made to this report. See report 10 for additional user options in defining system activities and user activities.

6.3.3.33 Report 38 - The Number of Times a System Activity was Swapped. This report is the same as report 34 except for system activities. See report 10 for additional user options in defining system activities and user activities.

6.3.3.34 Report 39 - The Total Elapsed Time a System Activity was Swapped. This report is the same as report 35 except for system activites. See report 10 for additional user options in defining system activities and user activities.

6.3.3.35 Report 40 - Number of Extra Activities That Might Fit in Memory Using Compaction. This report shows how memory might have been used more optimally. It takes the total amount of available memory (displayed in

report 5) and attempts to fit in those activities waiting memory. If an activity fits, the memory available is decreased, and the next activity is tried. If an activity does not fully fit, the next activity is tried. This continues until all available memory is used or until all the activities waiting have been tried. The search starts at the first waiting program an progresses serially down the program numbers of those waiting. This search ignores the actual size of "holes" or quadrant-crossing and is not necessarily obtainable or optimal. For this report, an entry is made at each allocator call.

6.3.3.36 Report 41 - Number of Extra Activities that Might Fit Memory Without Compaction. This report is the same type as report 40. In this case, activities are fit into existing holes and are ordered by urgency. The search progresses down the activities serially, beginning at the highest urgency activity. This histogram presents a good picture of how well the core allocator is performing its function.

For this report, an entry is made at each allocator call.

6.3.3.37 Report 42 - The Percent of Size-Time Product Used by a User Activity. This report shows the percentage of users each activities' size-time product over its run-time duration. An entry is made for each user activity that terminates. See report 10 for an explanation of user and system activities.

6.3.3.38 Report 43 through 49 - The Length of Idle State in the Processors. The elapsed clock time of an idle state is given in these reports for each individual processor and also as an average for all processors. They supply an indication of how each processor was utilized versus the others in the system. They also provide information on how busy the processors are. These reports should be used in conjunction with reports 26 through 32.

IDLEM data is used to produce these reports. These reports will not be produced if IDLEM was not active or if the IDLEM reports have been disabled by user input command.

6.3.3.39 Report 50 - Original Allocation Time for User Memory in I/O Second. This report gives the time each user activty waited for its original allocation of memory. See report 10 for an explanation of user and system activities.

6.3.3.40 Report 51 - The Time-Corrected Percent of Assigned Memory Used. This report gives the time-corrected percentage of slave memory used over the monitoring period. Any memory being utilized by jobs with zero urgency will not be included in the memory-used figure for this report. See report 16 for a definition of Time Correction.

6.3.4 Activity Resource Usage Report. For each activity known to the monitor, a detailed Resource Usage Report is made upon termination of the activtyl. The report is ordered by termination time sequence, and the resource usage is that known to the system at the last allocator call (refer to figure 6-15).

Each activity is displayed via the SNUMB and activity number followed by the CP an I/O charge times expressed in milliseconds. This is the CP and I/O tmes generated during the monitoring session. The size-time product is the total K words times the microseconds of allocation time, which gives a better expression for the memory used by the job than the size of the job. The minimum and maximum core requirements of the job are then shown, including the activity Slave Service Areas (SSAs) as well as slave size.

The elapsed time, in hours, an activity was known to the allocator is followed by the number of timesthe job size changed for any reason. The wasted core column is calculated from the job Slave Prefix Area (SPA) word 37 octal. This is filled by the System Loader and may not be valid for all job types (i.e., an H* file is not loaded in the normal system load manner). This clumn is shown in order to help locate users that do not have the $LIMITS card set correctly for the memory being used. If the user appears to be requesting excessive core on his $LIMITS card, he may be using this extra space as a spare buffer area. If this figure shows an excessive misuse of the $LIMITS card, the user should be contacted and questioned.

The next two columns provide a count of the total number of swaps and moves incurred by the activity. The final columns of the entry gives memory allocation time, wait time, swap time, memory time, and GEWAKE time, all in tenths of a second for each activity. An entry will be made in this report for every activity of a job, when the activity completes. Upon termination of the monitor, the resource usage of all actiites known to the allocator will be reported, including system jobs. This output follows a full line of asterisks to denote that no termination records were found for these activities.

COLLECTED ON SYSTEM NMCC2 ON 80-12-15 AT TIME 12:39

ACTIVITY RESOURCE USAGE REPORT - REPORTED PER ACT SIZE								ELAPSED	SIZE WASTED				TIME(.1 SEC) SPENT IN			
SNUMB-ACT		CPU & IO TIME	(MS)	SIZE-TIME PROD		MIN	MAX	TIME	CHANGE	CORE	SWAPS	MOVES	ALLOC	SWAP	MEMORY	GEWAKE
7332T-	0	44	108	1.5829E 08		31	31	0.003	0	0	0	0	51	0	55	0
7329T-	0	12077	29043	2.6766E 09		32	32	0.023	0	0	0	0	0	0	837	0
$GENB-	0	275	1003	1.7747E 07		11	11	0.001	0	0	0	0	2	0	16	0
7339T-	0	1300	2483	1.5904E 08		40	40	0.008	0	0	0	0	234	0	41	0
7348T-	0	1421	2646	1.3063E 08		31	31	0.042	0	0	0	0	1463	0	43	0
7338T-	1	991	2761	8.1595E 07		12	12	0.016	0	2	0	0	512	0	67	0
7338T-	2	2267	4284	2.6228E 08		25	36	0.003	2	0	0	2	15	0	83	0
7338T-	3	2761	4297	3.7449E 08		45	56	0.006	2	0	0	2	131	0	83	0
7354T-	0	1360	2909	2.1781E 08		40	40	0.002	0	0	0	0	5	0	55	0
7338T-	4	1766	3848	3.5940E 08		42	53	0.005	1	6	0	1	110	0	76	0
XXXXX-	0	2388	13991	4.8413E 08		11	11	0.016	0	0	0	0	146	0	440	0

FOLLOWING INFO MAY BE INCOMPLETE DUE TO LOST, DATA, NO EOF, OR ACTIVITY WAS ACTIVE WHEN MONITOR ENDED.

SYSTEM SCHEDULER CPU TIME, IO TIME AND SIZE-TIME PRODUCT ARE 43122911 129261145 0.19D 12

$PALC-	1	615091	2394913	2.9917E 11	17	17	6.822	0	0	988	108	0	17546	175982	52077
SYSOT-	0	626346	21551415	6.2101E 11	27	27	6.822	0	0	58	119	0	1877	230004	13728
$RGIN-	0	52768	167826	1.3920E 11	16	16	6.822	0	0	415	28	0	34039	87000	124567
-TSS—	1	12163737	13709653	3.1906E 12	77	187	6.822	58	0	0	0	0	0	245610	0
LOGON-	0	91433	59898	1.0459E 11	11	11	6.822	0	0	292	23	0	10302	95078	140228
FSYS -	0	0	90	2.7697E 10	9	9	6.822	0	0	0	0	0	0	30774	214835
NCP -	1	387705	72254	5.9289E 11	26	28	6.822	2	0	79	43	0	3452	216492	25664
TELNE-	1	231955	27456	2.5972E 11	16	17	6.822	2	0	150	22	0	5644	161220	78745
FTS -	1	216277	488796	1.7443E 11	24	51	4.732	176	0	191	111	9	3053	55649	111637
TLCF -	1	35090	32496	1.5089E 11	16	16	6.822	0	0	54	44	0	1839	94307	149462
DMTEX-	1	0	60	7.6714E 06	4	4	6.822	0	0	1	0	0	13	19	245577
DMSTA-	1	450328	1113536	1.5680E 11	8	8	6.822	0	0	717	76	0	13430	195994	36182
VIDEO-	1	116630	496	1.4735E 11	6	6	6.822	0	0	2	0	0	21	245588	0
7871T-	2	505393	384600	6.4366E 10	38	38	0.645	0	0	45	9	26	6263	16938	0
WAITL-	999	0	0	0.	58	58	0.175	0	0	0	0	6291	0	0	0
$GENB-	0	673795	2019705	1.9096E 11	** MEMORY USE TIME * 8.6242E 09											

Figure 6-15. Activity Resource Usage Report

In both this report and the following report, two special SNUMBs may appear. These SNUMBs are SWAKE an WAITL. These SNUMBs will appear for any activity that was in SWAKE or waiting original core allocation during the entire monitoring sesion. Due to the manner in which the data collector determines the SNUMB of a job, the real SNUMBs are not known for activities in either of the above categories.

6.3.5 SNUMB-IDENT Report. The SNUMB-IDENT Report is used to correlate the SNUMB, IDENT, and USERID of an activity. This allows operations personnel to identify each job reported in the Memory Map and Activity Usage Reports with a particular user (see figure 6-16).

The report displays the SNUMB, activity number, as well as the $IDENT and $USERID cards supplied at run time. The report also supplies the start and stop times of the activity. Following the stop time are four columns of data describing the urgency demands made by the activity. The four columns are: IU (Initial Urgency), AU (Average Urgency), HU (Highest Urgency) and MU (Minimum Urgency). At the far end of an entry, the type function being performed by an activity (i.e., GELOAD, FILSYS, COBOL, FORTY, etc...) is also presented.

The average urgency of an activity is merely the summation of urgencies recorded for this activity in each memory trace divided by the total number of times the urgency was recorded. It should be noted that GCOS may alter a job's urgency to as high as 63 for very short periods of time, and this will be recorded under the HU column. However, the value recorded under the AU column will be a much more accurate representation of the activity's true urgency during processing.

As each activity terminates, its entry is made to this report. Upon termination of the monitor, a summary of each activity still being processed at monitor termination will be given below a line of asterisks. These activities will include the system jobs and will provide an indication of system costs.

6.3.6 Memory Map Report. The Memory Map Report supplies a complete mapping of memory allocation. Memory is broken down into 128K half-quadrant sections and is displayed as such. This report can produce a tremendous quantity of data. Users should consider using time intervals any time the Memory map is to be produced.

Total memory can be pictured by laying each half quadrant side by side, with a time correlation being made by using the page and line numbers supplied on each output. The output is lined up by matching page number and quadrant numbers. The absolute clock time for each quadrant is found on the map of the first half quadrant of the system (0 to 128K). (Refer to figure 6-17).

The first half quadrant shows the time the state was present, the time since the last state change, and the memory used by the Hard Core Modules (HCM). The HCM usage is shown via a ****HCM*** character string.

The remaining space and all other half quadrants display the allocation of memory per job activity. All memory allocated is shown and the format is as follows:

----JOBOL-XXXUYY---

with the left and right asterisks representing the upper and lower addresses of the job whose SNUMB is JOBOL, with an activity number of XXX and an urgency of YY. Each character displayed for a job (from and including the asterisks) represents 1K of memory. As a job size decreases, the format changes as follows:

THIS PAGE LEFT INTENTIONALLY BLANK

```
                SNUMB-IDENT REPORT COLLECTED ON SYSTEM NMCC2  ON 80-12-15 AT TIME 12:39
SNUMB           START    STOP    IU  AU  HU  MU   USERID          IDENT
8437T-   2      18:21    18:27    5   5   5   5   DJ3JC32406      18200/30/4641,TEST ESET                                  GELOAD
8002T-  35      18:21    18:28    5   5   5   5   DJ8XC342LE      1820214/30/2353,TSTCRE                                   JOVIAL
28023-   1      18:24    18:28    5   5   5   5   DJ4XC31611      2806138/10/5093,NAV-LOOK2,C316,LINCOLN                   GMAP
29827-   3      18:21    18:29    5   5   5   5   DJ8XC32301      1820011/10/2389,STATREP      SLOAN  323              GELOAD
8496T-   0      18:29    18:29    5   5   5   5   DJ3JC32406      1820020/30/4641,TEST ESET                                GELOAD
28923-   7      18:28    18:29    5   5   5   5   DJ4XI31622      2806177/30/5168,AEAHAM                                   FORTY
8501T-   1      18:30    18:30    5   5   5   5   ABC005EEEEEE    ::000000000000000000000000000000000000000000000000000   OR
8454T-   2      18:24    18:32    5   5   5   5   DJ3NI33202      2800165/30/4683,RP                                       COBOL
7776T-   8      18:26    18:32    5   5   5   5   DJ8XC341HK      1820223/30/5202,GIPSY-SRC                                GELOAD

FOLLOWING INFO MAY BE INCOMPLETE DUE TO LOST, DATA, NO EOF, OR ACTIVITY WAS ACTIVE WHEN MONITOR ENDED.
$PALC-   1      12:39    19:29   55  55  55  55
SYSOT-   0      12:39    19:29   55  55  55  55
$RGIN-   0      12:39    19:29   55  55  55  55
-TSS--   1      12:39    19:29   55  55  55  55
LOGON-   0      12:39    19:29   55  55  55  55
FSYS -   0      12:39    19:29   55  55  55  55
NCP  -   1      12:39    19:29   55  55  55  55
TELNE-   1      12:39    19:29   55  55  55  55
FTS  -   1      14:45    19:29   55  55  55  55   NET-EXEC        000000000000000000000000000000000000000000000000000000000.MF
TLCF -   1      12:39    19:29   55  55  55  55
DMTEX-   1      12:39    19:29   55  55  55  55
DMSTA-   1      12:39    19:29   55  55  55  55   OPNSUTIL        ZZZ,OPNSTIL                                              DMSTA
VIDEO-   1      12:39    19:29   55  55  55  55   OPNSUTIL        ZZZ,OPNSTIL                                              VIDEO
7871T-   2      18:50    19:29   55  55  55  55   DJ3JC32402      1820020/30/4641,AIR PATCH TEST 5027      MEACHAM (CT20AIR
```

Figure 6-16. SNUMB-IDENT Report

```
                                 30        40        50        60        70        80        90        100       110       120
LINE    TOD        DELTA         5678901234567890123456789012345678901234567890123456789012345678901234567890123456789012345678
  1- 16.57942 0.00163889 ********************AHCM********** *----------------------FTS -001062----------------------**--------JDSIP-
  2- 16.58025 0.00083351 ********************AHCM********** *----------------------FTS -001062----------------------**--------JDSIP-
  3- 16.58050 0.00024986 ********************AHCM********** *----------------------FTS -001062----------------------**--------JDSIP-
  4- 16.58056 0.00005555 ********************AHCM********** *----------------------FTS -001062----------------------**--------JDSIP-
  5- 16.58058 0.00002789 ********************AHCM********** *----------------------FTS -001062----------------------**--------JDSIP-
  6- 16.58097 0.00038886 ********************AHCM********** *----------------------FTS -001062----------------------**--------JDSIP-
  7- 16.58097 0.         ********************AHCM********** *----------------------FTS -001062----------------------**--------JDSIP-
  8- 16.58169 0.00072217 ********************AHCM********** *----------------------FTS -001062----------------------**--------JDSIP-
  9- 16.58175 0.00005555 ********************AHCM********** *----------------------FTS -001062----------------------**--------JDSIP-
 10- 16.58258 0.00083327 ********************AHCM********** *----------------------FTS -001062----------------------**--------JDSIP-
 11- 16.58308 0.00049996 ********************AHCM********** *----------------------FTS -001062----------------------**--------JDSIP-
 12- 16.58314 0.00005579 ********************AHCM********** *----------------------FTS -001062----------------------**--------JDSIP-
 13- 16.58314 0.         ********************AHCM********** *----------------------FTS -001062----------------------**--------JDSIP-
 14- 16.58325 0.00011110 ********************AHCM********** *----------------------FTS -001062----------------------**--------JDSIP-
 15- 16.58400 0.00074983 ********************AHCM********** *----------------------FTS -001062----------------------**--------JDSIP-
 16- 16.58403 0.00002789 ********************AHCM********** *----------------------FTS -001062----------------------**--------JDSIP-
 17- 16.58408 0.00005555 ********************AHCM********** *----------------------FTS -001062----------------------**--------JDSIP-
 18- 16.58411 0.00002766 ********************AHCM********** *----------------------FTS -001062----------------------**--------JDSIP-
 19- 16.58478 0.00066662 ********************AHCM********** *----------------------FTS -001062----------------------**--------JDSIP-
 20- 16.58483 0.00005555 ********************AHCM********** *----------------------FTS -001062----------------------**--------JDSIP-
 21- 16.58483 0.         ********************AHCM********** *----------------------FTS -001062----------------------**--------JDSIP-
 22- 16.58528 0.00044465 ********************AHCM********** *----------------------FTS -001062----------------------**--------JDSIP-
 23- 16.58542 0.00013876 ********************AHCM********** *----------------------FTS -001062----------------------**--------JDSIP-
 24- 16.58547 0.00005555 ********************AHCM********** *----------------------FTS -001062----------------------**--------JDSIP-
 25- 16.58550 0.00002789 ********************AHCM********** *----------------------FTS -001062----------------------**--------JDSIP-
 26- 16.58594 0.00044441 ********************AHCM********** *----------------------FTS -001062----------------------**--------JDSIP-
 27- 16.58678 0.00083327 ********************AHCM********** *----------------------FTS -001062----------------------**--------JDSIP-
 28- 16.58683 0.00005555 ********************AHCM********** *----------------------FTS -001062----------------------**--------JDSIP-
 29- 16.58683 0.         ********************AHCM********** *----------------------FTS -001062----------------------**--------JDSIP-
 30- 16.58767 0.00083327 ********************AHCM********** *----------------------FTS -001062----------------------**--------JDSIP-
 31- 16.58772 0.00005555 ********************AHCM********** *----------------------FTS -001062----------------------**--------JDSIP-
 32- 16.58772 0.         ********************AHCM********** *----------------------FTS -001062----------------------**--------JDSIP-
 33- 16.58775 0.00002789 ********************AHCM********** *----------------------FTS -001062----------------------**--------JDSIP-
 34- 16.58783 0.00008321 ********************AHCM********** *----------------------FTS -001062----------------------**--------JDSIP-
 35- 16.58786 0.00002789 ********************AHCM********** *----------------------FTS -001062----------------------**--------JDSIP-
 36- 16.58869 0.00083327 ********************AHCM********** *----------------------FTS -001062----------------------**--------JDSIP-
 37- 16.58953 0.00083327 ********************AHCM********** *----------------------FTS -001062----------------------**--------JDSIP-
 38- 16.58958 0.00005555 ********************AHCM********** *----------------------FTS -001062----------------------**--------JDSIP-
 39- 16.58958 0.         ********************AHCM********** *----------------------FTS -001062----------------------**--------JDSIP-
 40- 16.58958 0.         ********************AHCM********** *----------------------FTS -001062----------------------**--------JDSIP-
 41- 16.59017 0.00058341 ********************AHCM********** *----------------------FTS -001062----------------------**--------JDSIP-
 42- 16.59100 0.00083327 ********************AHCM********** *----------------------FTS -001062----------------------**--------JDSIP-
 43- 16.59144 0.00044441 ********************AHCM********** *----------------------FTS -001062----------------------**--------JDSIP-
 44- 16.59153 0.00008345 ********************AHCM********** *----------------------FTS -001062----------------------**--------JDSIP-
 45- 16.59153 0.         ********************AHCM********** *----------------------FTS -001062----------------------**--------JDSIP-
 46- 16.59161 0.00008345 ********************AHCM********** *----------------------FTS -001062----------------------**--------JDSIP-
 47- 16.59164 0.00002766 ********************AHCM********** *----------------------FTS -001062----------------------**--------JDSIP-
 48- 16.59164 0.         ********************AHCM********** *----------------------FTS -001062----------------------**--------JDSIP-
 49- 16.59247 0.00083327 ********************AHCM********** *----------------------FTS -001062----------------------**--------JDSIP-
 50- 16.59331 0.00083351 ********************AHCM********** *----------------------FTS -001062----------------------**--------JDSIP-
```

Figure 6-17. Memory Map Report
(Part 1 of 3)

```
          140       150       160       170       180       190       200       210       220       230       240       250
LINE 2345678901234567890123456789012345678901234567890123456789012345678901234567890123456789012345678901234567890123456
  1- U33-------**----TLNT -C01U61----* *----------RUNPD-001U05----------**------MUM -001U62------**--------NCP -001U61--------*
  2- U33-------**----TLNT -001U61----* *----------RUNPD-001U05----------**------MUM -001U62------**--------NCP -001U61--------*
  3- U33-------**----TLNT -001U61----* *----------RUNPD-001U05----------**------MUM -001U62------**--------NCP -001U61--------*
  4- U33-------**----TLNT -001U61----* *----------RUNPD-001U05----------**------MUM -001U62------**--------NCP -001U61--------*
  5- U33-------**----TLNT -CC1U61----* *----------RUNPD-001U05----------**------MUM -001U62------**--------NCP -001U61--------*
  6- U33-------**----TLNT -C01U61----*                                   *------MUM -001U62------**--------NCP -001U61--------*
  7- U33-------**----TLNT -001U61----*                   *--NACE -001U00--**------MUM -001U62------**--------NCP -001U61--------*
  8- U33-------**----TLNT -001U61----*         JDSUP-001*--NACE -001U00--**------MUM -001U62------**--------NCP -001U61--------*
  9- U33-------**----TLNT -001U61----*         JDSUP-001*--NACE -001U00--**------MUM -001U62------**--------NCP -001U61--------*
 10- U33-------**----TLNT -001U61----*         JDSUP-001*--NACE -001U00--**------MUM -001U62------**--------NCP -001U61--------*
 11- U33-------**----TLNT -C01U61----*         JDSUP-001*--NACE -001U00--**------MUM -001U62------**--------NCP -001U61--------*
 12- U33-------**----TLNT -C01U61----*         JDSUP-001*--NACE -001U00--**------MUM -001U62------**--------NCP -001U61--------*
 13- U33-------**----TLNT -C01U61----*         JDSUP-001*--NACE -001U00--**------MUM -001U62------**--------NCP -001U61--------*
 14- J33-------**----TLNT -001U61----*         JDSUP-001*--NACE -001U00--**------MUM -001U62------**--------NCP -001U61--------*
 15- U33-------**----TLNT -C01U61----*         JDSUP-001*--NACE -001U00--**------MUM -001U62------**--------NCP -001U61--------*
 16- U33-------**----TLNT -001U61----*                   *--NACE -001U00--**------MUM -001U62------**--------NCP -001U61--------*
 17- U33-------**----TLNT -C01U61----*                                   *------MUM -001U62------**--------NCP -001U61--------*
 18- U33-------**----TLNT -CC1U61----*                   *--NACE -001U63--**------MUM -001U62------**--------NCP -001U61--------*
 19- U33-------**----TLNT -001U61----*                   *--NACE -001U63--**------MUM -001U62------**--------NCP -001U61--------*
 20- J33-------**----TLNT -001U61----*                   *--NACE -001U63--**------MUM -001U62------**--------NCP -001U61--------*
 21- U33-------**----TLNT -C01U61----*                   *--NACE -001U63--**------MUM -001U62------**--------NCP -001U61--------*
 22- U33-------**----TLNT -001U61----*                   *--NACE -001U63--**------MUM -001U62------**--------NCP -001U61--------*
 23- U33-------**----TLNT -001U61----*                   *--NACE -001U63--**------MUM -001U62------**--------NCP -001U61--------*
 24- U33-------**----TLNT -C01U61----*                                   *------MUM -001U62------**--------NCP -001U61--------*
 25- U33-------**----TLNT -C01U61----*                   *--NACE -001U00--**------MUM -001U62------**--------NCP -001U61--------*
 26- U33-------**----TLNT -C01U61----*         JDSUP-001*--NACE -001U00--**------MUM -001U62------**--------NCP -001U61--------*
 27- U33-------**----TLNT -001U61----*         JDSUP-001*--NACE -001U00--**------MUM -001U62------**--------NCP -001U61--------*
 28- U33-------**----TLNT -001U61----*         JDSUP-001                 *------MUM -001U62------**--------NCP -001U61--------*
 29- U33-------**----TLNT -001U61----*         JDSUP-001*--NACE -001U00--**------MUM -001U62------**--------NCP -001U61--------*
 30- J33-------**----TLNT -C01U61----*         JDSUP-001*--NACE -001U00--**------MUM -001U62------**--------NCP -001U61--------*
 31- U33-------**----TLNT -C01U61----*         JDSUP-001                 *------MUM -001U62------**--------NCP -001U61--------*
 32- U33-------**----TLNT -001U61----*         JDSUP-001*--$HEAL-00JU32--**------MUM -001U62------**--------NCP -001U61--------*
 33- J33-------**----TLNT -001U61----*         JDSUP-001*--$HEAL-00JU32--**------MUM -001U62------**--------NCP -001U61--------*
 34- U33-------**----TLNT -001U61----*                   *--$HEAL-00JU32--**------MUM -001U62------**--------NCP -001U61--------*
 35- U33-------**----TLNT -001U61----* *--NACE -001U00--**--$HEAL-00JU32--**------MUM -001U62------**--------NCP -001U61--------*
 36- U33-------**----TLNT -C01U61----* *--NACE -001U00--**--$HEAL-00JU32--**------MUM -001U62------**--------NCP -001U61--------*
 37- U33-------**----TLNT -C01U61----* *--NACE -001U00--**--$HEAL-00JU32--**------MUM -001U62------**--------NCP -001U61--------*
 38- U33-------**----TLNT -001U61----* *--NACE -001U00--*                   *------MUM -001U62------**--------NCP -001U61--------*
 39- J33-------**----TLNT -C01U61----*                                   *------MUM -001U62------**--------NCP -001U61--------*
 40- U33-------**----TLNT -001U61----*                   *--NACE -001U00--**------MUM -001U62------**--------NCP -001U61--------*
 41- U33-------**----TLNT -001U61----*         JDSUP-001*--NACE -001U00--**------MUM -001U62------**--------NCP -001U61--------*
 42- U33-------**----TLNT -001U61----*         JDSUP-001*--NACE -001U00--**------MUM -001U62------**--------NCP -001U61--------*
 43- U33-------**----TLNT -CC1U61----*         JDSUP-001*--NACE -001U00--**------MUM -001U62------**--------NCP -001U61--------*
 44- U33-------**----TLNT -001U61----*         JDSUP-001                 *------MUM -001U62------**--------NCP -001U61--------*
 45- U33-------**----TLNT -001U61----*         JDSUP-001*--$HEAL-00JU32--**------MUM -001U62------**--------NCP -001U61--------*
 46- U33-------**----TLNT -001U61----*         JDSUP-001*--$HEAL-00JU32--**------MUM -001U62------**--------NCP -001U61--------*
 47- U33-------**----TLNT -001U61----*                   *--$HEAL-00JU32--**------MUM -001U62------**--------NCP -001U61--------*
 48- U33-------**----TLNT -001U61----* *--NACE -001U00--**--$HEAL-00JU32--**------MUM -001U62------**--------NCP -001U61--------*
 49- U33-------**----TLNT -C01U61----* *--NACE -001U00--**--$HEAL-00JU32--**------MUM -001U62------**--------NCP -001U61--------*
 50- U33-------**----TLNT -C01U61----* *--NACE -001U00--**--$HEAL-00JU32--**------MUM -001U62------**--------NCP -001U61--------*
```

Figure 6-17. (Part 2 of 3)

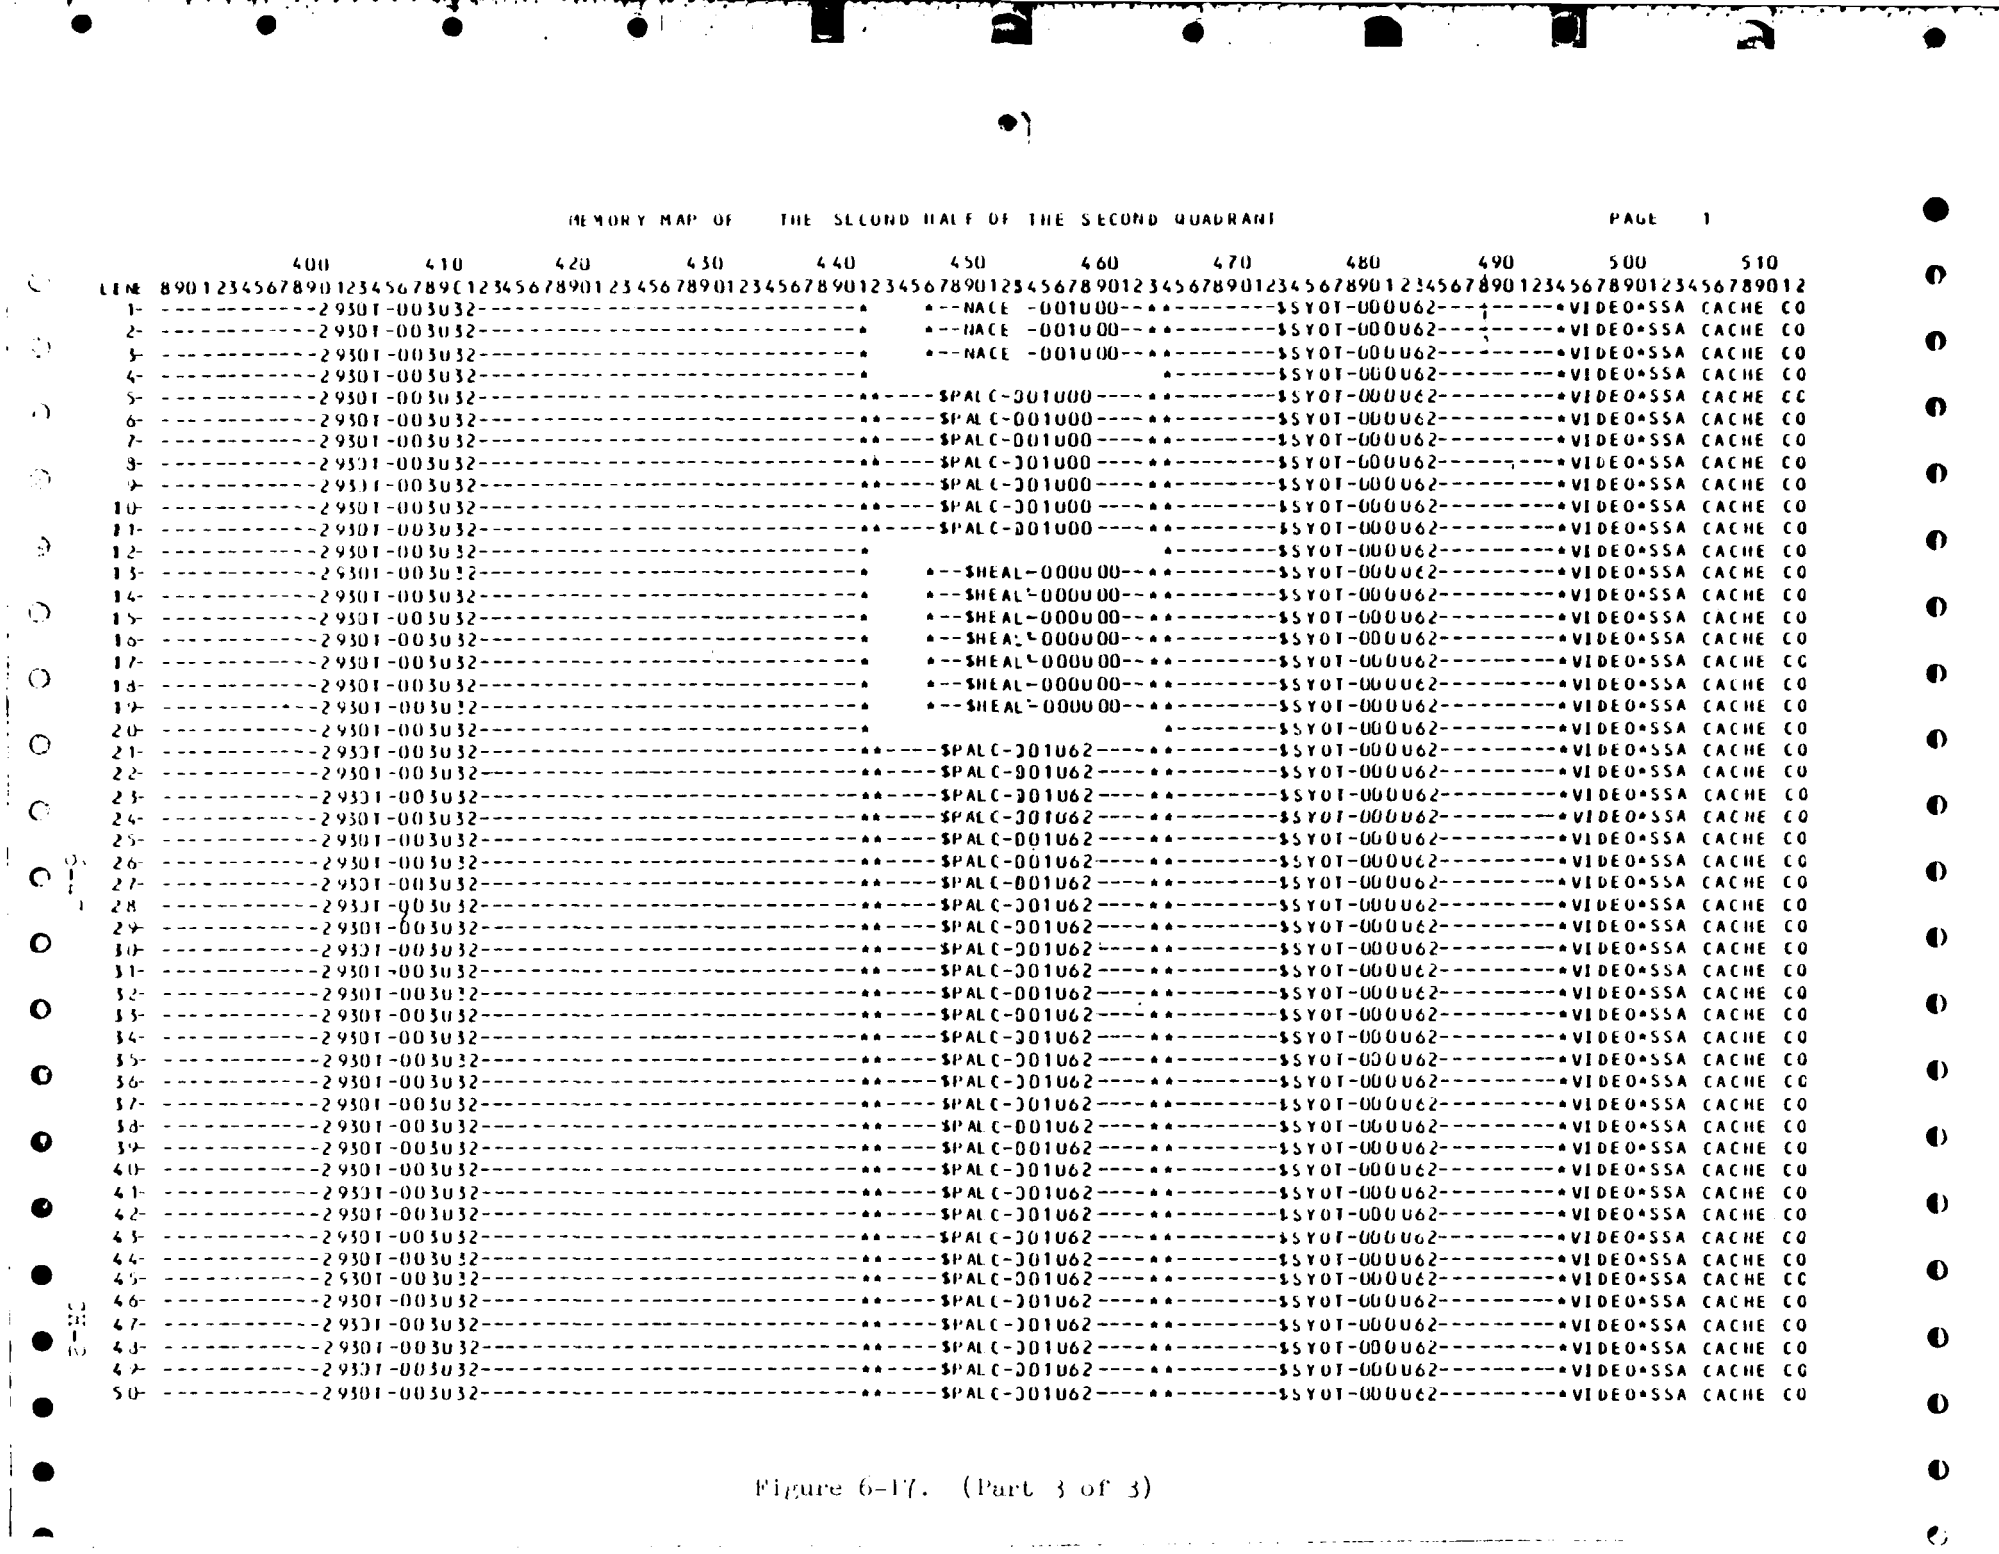


Figure 6-17. (Part 3 of 3)

-JOB01-XXXU05-	(16K) including SSAs
*JOB02-XXX--	(12K)
JOB03	(7K) - activity discarded
JOB04	(5K) - asterisks discarded
****	(4K) - no identification

As can be seen in figure 6-17, Part 1, every line of the figure has a line number ranging from 1-50. In addition, there is a page number in the upper right corner. When the user wants to match this picture of the first half of the first quadrant with its corresponding half of some other quadrant, the following steps should be followed:

1- Match page numbers (see figure 6-17, Part 2, Part 3)
2- Match the line numbers from identical page numbers.

Two special names can appear in the memory map. If SSA cache memory is configured the following letters will be found in the map, depending on the size of the SSA cache memory:

*SSA CACHE**	12K	
*SSA CACHE	10K	
*SSA CAC	8K	
SSA C	5K	(see figure 6-17, Part 3)

If memory has been released from the system, then the letters *-RELEASED* will appear in the dump. This will be repeated depending upon how much memory has been released.

6.3.7 Demand List Report. The Demand List Report shows the memory demand outstanding for each memory state displayed on the memory map. The correlation is made using the same line numbers as the half quadrants of the maps themselves (refer to figure 6-18).

The Demand List Report shows the total memory available, the number of jobs waiting memory, the demand request sizes for each job waiting memory. The memory available is the sum of all holes in memory.

6.3.8 Activity Abort Request. This report is directly related to the Activity Resource Usage Report. This report is produced whenever the Activity Report is produced. For every activity that aborts during the monitoring session, an entry is made to this report. The entry gives the SNUMB, Activity Number, Abort Code, CPU Time, Run Hours, USERID, and IDENT for the activity.

The Abort Code is either an octal number or an alphanumeric value. The meaning of these codes can be found an Appendix A of Honeywell Manual DD19 (GCOS).

DEMAND LIST IN K BLOCKS PAGE 1

LINE								
1-	14K AVAILABLE	5 WAITING-	38	38	38	34	34	
2-	41K AVAILABLE	5 WAITING-	38	38	38	34	34	
3-	41K AVAILABLE	5 WAITING-	38	38	38	34	34	
4-	41K AVAILABLE	5 WAITING-	38	38	38	34	34	
5-	52K AVAILABLE	6 WAITING-	17	38	38	38	34	34
6-	35K AVAILABLE	6 WAITING-	3	38	38	38	34	34
7-	35K AVAILABLE	6 WAITING-	3	38	38	38	34	34
8-	35K AVAILABLE	6 WAITING-	3	38	38	38	34	34
9-	46K AVAILABLE	6 WAITING-	3	38	38	38	34	34
10-	43K AVAILABLE	5 WAITING-	38	38	38	34	34	
11-	43K AVAILABLE	5 WAITING-	38	38	38	34	34	
12-	43K AVAILABLE	5 WAITING-	38	38	38	34	34	
13-	5K AVAILABLE	5 WAITING-	11	38	38	34	34	

Figure 6-18. Demand List Report

The CPU time is the amount of CPU time in milliseconds used by this activity prior to its abort. This is the total CPU time generated during the monitoring session.

The Run Hours is the amount of time this activity has been known to the monitor (see figure 6-19).

At the bottom of the report, the percent of total CPU time, total I/O time, and total Run time used by the aborted activities is given. This percentage does not include any system jobs (i.e., program number $\leq$ 13) or any selected SNUMBs processed by Action Code ABORT. The selected SNUMBs are listed at the end of the report.

6.3.9 Jobs Out of Core Report. This report gives a detailed picture of memory demand outstanding for each memory state displayed on the memory map. The correlation is again made by matching the page numbers, line numbers, and times of day (see figure 6-20). If a line number does not appear, no jobs were out of memory. If a line number appears more than once, more than two jobs were out of memory at that time. For each line, the report presents the following:

- Line number (always 50 if out of core report is run without memory map)
- SNUMB and activity number of job waiting memory
- Memory demand and urgency
- Whether the job could fit in memory if available memory were compacted (FWC)
- Whether it could fit into an existing hole of memory (FWOC)
- Whether the job is a new request (NEW), a swapped job (SWAP), or a job GEWAKE (GWKE)
- How many attempts have been made by the Core Allocator to place this job into memory

It is also possible for a page number to be repeated. This occurs because the memory map prints 50 lines per page. Since the Out of Core Report can print several lines for each memory map line, it might be necessary to use several pages for each memory map page.

This report can produce a tremendous amount of output. Users should consider using the time interval option if this report is needed.

6.3.10 Excessive Resource Use Report. This report is directly related to the Activity Resource Usage Report. For every activity that is apparently using more than a preset amount of specified resources (Wasted Memory, Memory Used, I/O Secs, CPU Secs, Ratio, Urgency), an entry is made to this report. The user must realize the Wasted Memory column is not an absolute statement that a user is wasting memory. Rather, it is a

ACTIVITY ABORT REPORT ON SYSTEM NMCC2 ON 80-12-15 AT TIME 12:39

SNUMB-ACT	ABRT CODE	CPU(MS)	RUN(HR)	USERID	IDENT	
7375T- 2	Q6	2414	0.002	DJ8XI70203	1829115/10/6544,PELLERIN	GELOAD
7413T- 3	GR	2703	0.009	DJ8XI70203	1829115/10/6544,PELLERIN	GELOAD
7431T- 6	031	10317	0.011	DJ8XI70203	1829115/10/6544,PELLERIN	GELOAD
7483T- 3	077	22691	0.048	DJ8XI70203	1829115/10/6544,PELLERIN	GELOAD
7513T- 3	SC	2090	0.004	DJ8XI70357	2802112/30/6242,TITUS	GELOAD
7545T- 3	115	2932	0.008	DJ8XI70357	2802112/30/6242,TITUS	GELOAD
7676T- 1	051	495	0.005	DJ8XI342DH	1820215/30/4938,FTS	FTSBIP
WAS01- 1	FS	584	0.040	OPNSUTIL	1727330/10/5471,WAS01	FILSYS
7713T- 1	005	108542	0.406	DJ3JC32411	1820020/30/3170,SANDY	FILEDI

THE % OF CPU TIME USED BY ABORTED JOBS WAS 5.29
THE % OF IO TIME USED BY ABORTED JOBS WAS 2.01
THE % OF RUN TIME USED BY ABORTED JOBS WAS 4.67

**** THE ABOVE PERCENTAGES DO NOT INCLUDE ANY SYSTEM JOBS OR ANY OF THE FOLLOWING JOBS ****
HEALS MUM MSM GDR NEWRM

Figure 6-19. Abort Report

JOBS OUT OF CORE REPORT FOR SYSTEM NMCC2 COLLECTED ON 81-04-09 PAGE # 1

INF	TIME	SNUMD	DMND	URG	FWC	FWOC	TYPE	TRIES	SNUMD	DMND	URG	FWC	FWOC	TYPE	TRIES
1	14.00044	FSYS	0 5	0	NO	NO	GWKE	0	NCP	1 26	0	NO	NO	GWKE	7
1	14.00044	TELNE	1 16	0	NO	NO	GWKE	3	FTS	1 24	0	NO	NO	GWKE	2
1	14.00044	TLCF	1 14	0	NO	NO	GWKE	3	DMTEX	1 4	0	NO	NO	GWKE	1
1	14.00044	HEALS	1	0	NO	NO	GWKE	3	NEWRM	1 3	0	NO	NO	GWKE	1
2	14.00083	HEALS	1 8	32	YES	YES	SWAP	1	FSYS	0 5	0	NO	NO	GWKE	0
2	14.00083	NCP	1 26	0	NO	NO	GWKE	7	TELNE	1 16	0	NO	NO	GWKE	3
2	14.00083	FTS	1 24	0	NO	NO	GWKE	2	TLCF	1 14	0	NO	NO	GWKE	3
2	14.00083	DMTEX	1 4	0	NO	NO	GWKE	1	NEWRM	1 3	0	NO	NO	GWKE	1
3	14.00158	2768T	0 46	5	NO	NO	NEW	1	FSYS	0 5	0	NO	NO	GWKE	0
3	14.00158	NCP	1 26	0	NO	NO	GWKE	7	TELNE	1 16	0	NO	NO	GWKE	3
3	14.00158	FTS	1 24	0	NO	NO	GWKE	2	TLCF	1 14	0	NO	NO	GWKE	3
3	14.00158	DMTEX	1 4	0	NO	NO	GWKE	1	NEWRM	1 3	0	NO	NO	GWKE	1
4	14.00219	2800T	0 11	24	YES	NO	NEW	1	FSYS	0 5	0	NO	NO	GWKE	0
4	14.00219	NCP	1 26	0	NO	NO	GWKE	7	TELNE	1 16	0	NO	NO	GWKE	3
4	14.00219	FTS	1 24	0	NO	NO	GWKE	2	TLCF	1 14	0	NO	NO	GWKE	3
4	14.00219	DMTEX	1 4	0	NO	NO	GWKE	1	NEWRM	1 3	0	NO	NO	GWKE	1
5	14.00292	NEWRM	1 3	60	YES	YES	SWAP	1	2791T	0 51	5	NO	NO	NEW	1
5	14.00292	FSYS	0 5	0	NO	NO	GWKE	0	NCP	1 26	0	NO	NO	GWKE	7
5	14.00292	TELNE	1 16	0	NO	NO	GWKE	3	FTS	1 24	0	NO	NO	GWKE	2
5	14.00292	TLCF	1 14	0	NO	NO	GWKE	3	DMTEX	1 4	0	NO	NO	GWKE	1
6	14.00328	2791T	0 51	5	NO	NO	NEW	2	FSYS	0 5	0	NO	NO	GWKE	0
6	14.00328	NCP	1 26	0	NO	NO	GWKE	7	TELNE	1 16	0	NO	NO	GWKE	3
6	14.00328	FTS	1 24	0	NO	NO	GWKE	2	TLCF	1 14	0	NO	NO	GWKE	3
6	14.00328	DMTEX	1 4	0	NO	NO	GWKE	1							
7	14.00342	2791T	0 51	5	NO	NO	NEW	3	FSYS	0 5	0	NO	NO	GWKE	0
7	14.00342	NCP	1 26	0	NO	NO	GWKE	7	TELNE	1 16	0	NO	NO	GWKE	3
7	14.00342	FTS	1 24	0	NO	NO	GWKE	2	TLCF	1 14	0	NO	NO	GWKE	3
7	14.00342	DMTEX	1 4	0	NO	NO	GWKE	1							
8	14.00358	2791T	0 51	5	NO	NO	NEW	4	FSYS	0 5	0	NO	NO	GWKE	0
8	14.00358	NCP	1 26	0	NO	NO	GWKE	7	TELNE	1 16	0	NO	NO	GWKE	3
8	14.00358	FTS	1 24	0	NO	NO	GWKE	2	TLCF	1 14	0	NO	NO	GWKE	3
8	14.00358	DMTEX	1 4	0	NO	NO	GWKE	1							
9	14.00372	2791T	0 51	5	NO	NO	NEW	5	FSYS	0 5	0	NO	NO	GWKE	0
9	14.00372	NCP	1 26	0	NO	NO	GWKE	7	TELNE	1 16	0	NO	NO	GWKE	3
9	14.00372	FTS	1 24	0	NO	NO	GWKE	2	TLCF	1 14	0	NO	NO	GWKE	3
9	14.00372	DMTEX	1 4	0	NO	NO	GWKE	1							
10	14.00378	2801T	0 11	24	YES	NO	NEW	1	2791T	0 51	5	NO	NO	NEW	6
10	14.00378	FSYS	0 5	0	NO	NO	GWKE	0	NCP	1 26	0	NO	NO	GWKE	7
10	14.00378	TELNE	1 16	0	NO	NO	GWKE	3	FTS	1 24	0	NO	NO	GWKE	2
10	14.00378	TLCF	1 14	0	NO	NO	GWKE	3	DMTEX	1 4	0	NO	NO	GWKE	1
11	14.00394	2791T	0 51	5	NO	NO	NEW	7	FSYS	0 5	0	NO	NO	GWKE	0
11	14.00394	NCP	1 26	0	NO	NO	GWKE	7	TELNE	1 16	0	NO	NO	GWKE	3
11	14.00394	FTS	1 24	0	NO	NO	GWKE	2	TLCF	1 14	0	NO	NO	GWKE	3
11	14.00394	DMTEX	1 4	0	NO	NO	GWKE	1							
12	14.00444	2791T	0 51	5	NO	NO	NEW	8	FSYS	0 5	0	NO	NO	GWKE	0
12	14.00444	NCP	1 26	0	NO	NO	GWKE	7	TELNE	1 16	0	NO	NO	GWKE	3
12	14.00444	FTS	1 24	0	NO	NO	GWKE	2	TLCF	1 14	0	NO	NO	GWKE	3
12	14.00444	DMTEX	1 4	0	NO	NO	GWKE	1							
13	14.00539	2791T	0 51	5	NO	NO	NEW	9	FSYS	0 5	0	NO	NO	GWKE	0

FIGURE 6-20. Jobs out of Core Report

statement that the $LIMITS card appears to be requesting more memory than is actually required by this job. The user should be questioned in order to determine if this is in fact true. In the Honeywell System, a user will receive whatever amount of memory requested on the $LIMITS card, whether or not the amount of memory is actually needed. The Ratio column shows the ratio of the total elapsed time for an activity divided by the total memory time for the activity. This value gives an indication of the activity lengthening factor; i.e., how run time is affected by resource contention. For those activities using excessive memory, the report also indicates, under the MEM MIN column, the amount of time the activity was in memory. The value being used for the urgency check is the average urgency recorded for the activity and not the maximum urgency of the activity. The default values for an entry being made to this report are listed in table 6-4. These values can be changed via a previously described input option. This report will be produced whenever the Activity Resource Report is produced and will be turned off whenever the Activity Resource Report is off (see figure 6-21).

6.3.11 Allocation Status Report. This report will track an activity as it proceeds through different phases of Allocation. The report will list the SNUMB-Activity #, amount of memory the activity will require, its current status, the time it entered that phase of allocation, the time is completed that phase of allocation, the total time spend in a given phase of allocation, the device type it is waiting for, and the number of devices the activity is waiting for. Due to the manner in which data is collected for this report, it is possible that certain phases of allocation will be missed, especially if that phase of allocation occurs within a short time span. This report will give a good indication of how long it is taking activities to pass through the various allocation phases prior to core allocation. Following is a list of the more common phases of allocation and their meanings:

New Act - Activity has just entered the Peripheral Allocator
Wait Media - Activity is waiting for a device
Wait Mnt - Activity is waiting for a patch or tape to be mounted
Core Queue Full - Activity has been completely processed and is waiting for the Peripheral Allocator to send the job to the core allocator
Alloc Done - Activity has been sent to core allocator. For this case the stop time and total time columns have no real meaning. These columns simple are reporting the amount of time it took the monitor to realize that the activity had reached the core allocator
LIMBO - Activity is in Limbo and has not even been granted permission to run
HOLD - Activity is in Hold and has not even been given permission to run
SCHED - Activity was in one of the System Scheduler queues.

Only activities found to be in a state for more than 600 seconds will be reported. This limit can be changed by using the PALC input option. See figure 6-22 for a sample of this report.

6.3.12 Plot Reports. Four different plot reports are produced by the data reduction program. All plots are produced under 10-minute intervals, where the interval can be modified by the user. At every allocator call, the various parameters to the plots are accumulated and every 10 minutes, the accumulated parameters are averaged and an average value is output

THIS PAGE LEFT INTENTIONALLY BLANK

EXCESSIVE RESOURCE USAGE REPORT ON SYSTEM DNAH66 ON 81-12-17 AT TIME 10:01

SNUMB-ACT	WASTED MEMORY	MEM USED	IO SECS	CPU SECS	RATIO	MEM MIN	MAX URG	USERID	IDENT
2027Q- 1		56				1.2		FCCCSE	1020,WA024,CSE,250,MCCANN,SECRET
2044T- 1		81				2.8		FCCCSE	1010,WA803,CSE,264,UARRETT,UNCL
2057T- 1		64				3.4	47	FCCCSD	1010,NS0033,CSD,283,PADILLA,UNCL
2083Q- 1		56				1.0		FCCCSE	1020,WA024,CSE,250,MCCANN,SECRT
2085T- 1		92				5.0	55	DBA	1020,TRAX.A,CSA,254,MATHEWS,SECRET
2106T- 1		81				2.7		FCCCSE	1010,WA803,CSE,264,BARRETT,UNCL
SR01 - 5		65				0.6		FCCCS	1030,SR01,CSC,250,MCCANN,TOPSEC

Figure 6-21. Excessive Resource Usage Report

ALLOCATION STATUS REPORT

SNUMB	MEMORY	STATUS	TIME IN	TIME OUT	TOTAL TIME(SEC)	DEVICE TYPE	NUMBER
7338T- 1	10	WAIT MNT	12:42	12:52	600.		
7338T- 4	51	CORE QUE FULL	12:46	12:57	660.		
29675- 2	36	WAIT MEDIA	12:49	13:00	660.	DSS191	2

Figure 6-22. Allocation Status Report

to the plot. Each horizontal position has a delta value, which is printed on the plot. The delta value is computed from the following formula:

$$\frac{\text{(Upper Plot Limit - Lower Plot Limit)}}{114}$$

The plot limits can be set by the user, with the default values shown in table 6-3. If the user changes the maximum limits, the new maximum limit selected should be divisible by 114. If a plotted variable is beyond or on an axis limit, it will be positioned at the axis limit. If any 2 points coincide, the position of coincidence will be marked with a 2. If 3 points coincide, the position of coincidence will be marked with a 3. (See figure 6-23). The end of the plot will contain a summary of the minimum and maximum values of each curve.

In figure 6-23, we see that there was no memory shortfall at all between 12:49 and 14:09 (points A & B are coinciding with the left axis; i.e., a 2 is output). At 14:19 and 14:29, both curves continue to coincide, but there is now a shortfall of 48K (12th point times delta of 4). At 14:39, the memory shortfall of the CALC increased to 92K but the memory shortfall of the CALC plus the PALC queue increased to 116K.

In order to obtain a continuous curve from the plots, the user needs only to connect the corresponding letter points (see figure 6-23).

6.3.12.1 Plot 1 - Available Memory vs. Outstanding Demand in Core Allocator Queue vs. Outstanding Demand in Core Allocator Queue + Peripheral Allocator Queue. This three parameter plot provides an overview of the time dependence of both the system load and memory availability. It can aid in better balancing the workload across the day and in determining when memory shortfalls or supluses exist. The addition of memory demand waiting in the Peripheral Allocator is an attempt to give a truer picture of how much additional memory could be properly utilized, if available. As long as the B and C points fall to the left of the A points, a memory surplus exists. If the B and C points fall to the right of the A points, a memory shortfall is present.

6.3.12.2 Plot 2 - Memory Shortfall in Core Allocator vs. Memory Shortfall in Core Allocator + Peripheral Allocator. This plot is obtained from the previous plot by simply calculating the actual shortfall and plotting the shortfall points.

6.3.12.3 Plot 3 - Number of Activities in Core Queue vs. Number of Activities in Peripheral Allocator Queue. In this plot, the number of activities waiting for memory is displayed, instead of their memory demand.

6.3.12.4 Plot 4 - Average Size of TSS, FTS, NCP. This plot displays the average size of TSS, FTS and NCP as they change their demand for memory over time. These three programs are those whose memory demand would significantly change over time. The FTS and NCP programs are part of the WIN subsystem.

DISTRIBUTION COLLECTED ON SYSTEM NMCC2 AT 12:39:46 ON 80-12-15

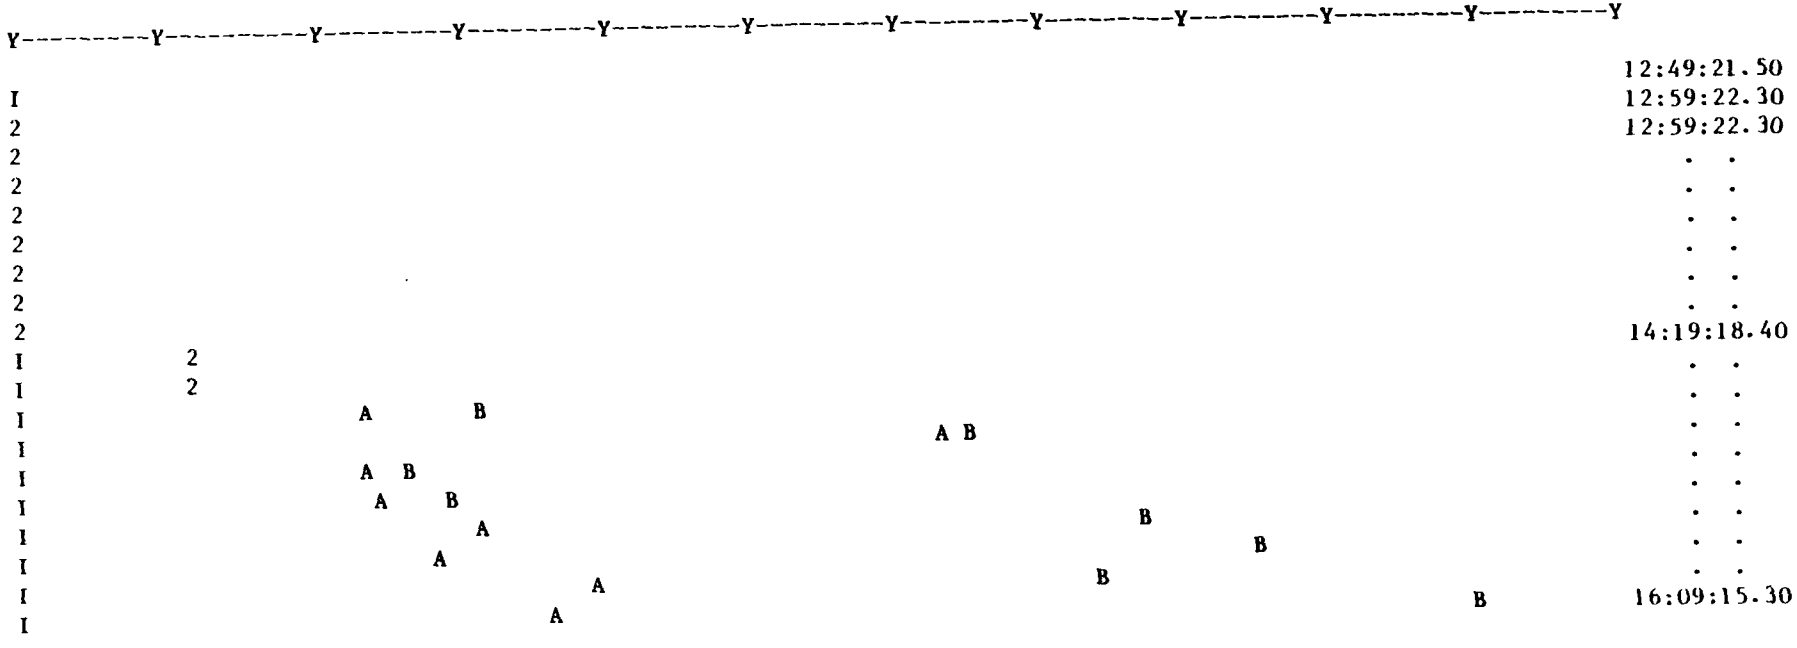
MEMORY SHORTFALL IN CALC VS SHORTFALL IN CALC+PALC PLOT- 2

DELTA= 4.0 A= CALC B=PALC

TIME OF DAY

MEMORY (K WORDS)

0. 112.0 228.0 340.0



CURVE	MINIMUM	MAXIMUM
CALC	0.	0.1950E 03
+ PALC	0.	0.3000E 03

Figure 6-23. Standard Plot

6.3.13 Memory Statistics Report. This report is produced after the histograms. It details all the information needed to start a system analysis as detailed in section 14. The report is shown in figure 6-24. The values for this report are obtained as described in section 14.6.3.

6.3.14 Special Job Memory Reports. This report details the memory demands made by a series of jobs specially requested by the analyst (see SPECL option). Figures 6-25 and 6-26 display the format for these reports. Figure 6-25 displays the Memory Demand Report. Every time one of the specially requested jobs issues a MME GEMORE for memory, gets swapped out/into memory, or release memory, an entry is made to this report. In figure 6-25, we see that FTS issued a GEMORE for 1K of memory at 11:47:03. The time is presented in hours and fraction of an hour as well as in hours, minutes and seconds. At the time of the GEMORE, FTS was at 35K. The -99999 in the Time to Satisfy column indicates that this was the time when the GEMORE was issued (and the request has not yet been satisfied or denied). The GEMORE was satisfied at 11:47:03 after a wait of 0 tenths of a second. The fact that the request was satisfied is indicated by the word MET under the DEMAND TYPE column and also the fact that FTS is now at 36K. The last column of the report can be used to determine how many cycles through the memory allocator were required before the request was satisfied or denied. The GEMORE request was made on the 82nd call to the allocator and was satisfied on the 83rd call to the allocator. At 11:47:09, FTS issued another GEMORE request that was DENIED 10 tenths of a second (1 second) later. This denial was immediately followed by a SWAP which lasted 2 tenths of a second. When FTS returned from the SWAP, it had received the 4K of memory that had been denied from the earlier GEMORE request. Finally, at 11:47:39, FTS released (GEMREL) 1K (-1) of memory and its size was reduced to 39K. At preselected intervals (same interval constraints used to produce the plots), a summary line is printed indicating the total wait time accumulated by the job during that interval of time. It should be noted that this option can cause the Memory Map to generate output even if the Memory Map Report has not been activated (see subsection 6.1.24).

6.3.15 Distribution of Urgency Over Time Report. This reports (figure 6-27) is always produced and cannot be turned off. An entry is this report is made under the same time interval contraints as the memory plots. The report displays the average distribution of urgencies present in the system during each interval of time. Therefore, between 14:06 and 14:16, 39.9% of all activities active in the system had an urgency level of 5, while 5.8% had an urgency level of 30. Urgencies are grouped in classes of 5 (i.e., 0-4 reported as 0, 5-9 reported as 5, etc.). The urgency classes of 55-60 and 60-65 are further subdivided into user activites (u) and system activities (s). The summary at the bottom of the page indicates, for each activity processed in the system, what percentage of activities reached a maximum urgency of the indicated level. Therefore, in figure 6-27, we see that 40.2% of all activities processed had a maximum urgency of 5, while 7.1% of all activities processed were user activities with a maximum urgency between 60-64, and 3.3% of all activities processed were system activities in the same urgency class.

The greatest consideration to an urgent job (40) is given with respect to allocation denial. If a priority job is denied and less than 3 jobs are active, the allocation overdue status is set immediately for that job. Allocation overdue will not be set for a routine job until 50 activities have been allocated. In any case, if a priority job is denied allocation, a flag is set signifying this. Until all urgent jobs have been allocated, all nonpriority new job or first activity entries in the job stack will be ignored. New activity entries will be processed, but it can be seen that needed peripherals could go unused for great lengths of time.

There are two algorithms available for determining when a program may receive processor service. The algorithm used at a given site is dependent upon a site option patch that is included in the start-up deck.

Option 1 is called "urgency throughput" and involves jobs being placed in the processor's queue based on their priority. Priority is computed as the job's urgency divided by four. With WIN and other system programs running at an urgency of 61 and user jobs running at urgencies of 40 and above, their relative priorities are very close (15 for WIN and other system programs and 10 or higher for the user activities). Each time a job is dispatched, its priority is decremented by one, and upon reaching zero is recomputed. Based upon this algorithm, WIN programs will theoretically have a priority higher than the user jobs for only one-third of the cycle. The other two-thirds of the time, when its priority drops below 10, WIN programs will be competing for the processor with many other less critical jobs as an equal.

Option 2 is called "I/O priority" and involves giving priority to heavy I/O type jobs. The reasoning behind this is that an I/O bound job will not require the processor for extended periods of time so, therefore, the system will give this job the processor so that it can perform its minimal CPU functions, issue an I/O request, and be removed from the processor queue.

Since WIN software (with the possible exception of FTS) does not perform much I/O, this puts any I/O bound slave job at a higher priority than the WIN software. This means that this I/O bound job will always go into the dispatcher's queue before the WIN software and most other critical software.

6.3.16 Zero Urgency Job Report. According to the GCOS software documentation, the Core Allocator will swap out any job with a zero urgency, if there is an outstanding demand for memory. During several system studies performed by CCTC, it has been observed that this swapping of zero urgency jobs does not always appear to be occurring. This report lists all those jobs which were observed as having a zero urgency, but were not swapped by the Core Allocator. A job is entered in this report only when the following conditions exist:

- o job has a zero urgency
- o there is an outstanding demand for memory
- o the size of this zero urgency job is sufficient to satisfy at least the smallest memory demand currently outstanding.

The report indicates the SNUMB and Userid of the job, the start and stop times of the job, the total amount of time (.1 sec) that the job satisfied all conditions stated above, and finally, the total percent of time that the job was in memory under the above conditions. Other reports would indicate the size of the jobs listed in this report and the level of outstanding demand that was present on this system, while this job was in execution. Figure 6-28 shows a sample of this report.

6.4 Error Messages

All error messages are self-explanatory or else followed by the words "For Information Only." In this case, the message can be ignored and processing will continue, without error.

MEMORY STATISTICS TABLE (DHAH66

DATE	START	STOP	dDURS	CPU/IO RATIO	AVG ACT SIZE(K)	USER MEM AVAILABLE	SYSTEM MEMORY	EXCESS MEMORY (+/-) CALC	PALC
811217	1001	1130	1.48	0.336	23.7L	21d	294	67	66

AVG # USER ACT WAITING MEM	AVG # SYSTEM ACT WAITING MEM	AVG # USER ACT IN MEM	AVG # SYSTEM ACT IN MEM	RATIO OF DURATION VS MEMORY TIME	% SLAVE MEM USED	TIME USER ACT SWAPPED
1.000	0.-	3.765	8.771	1.150	75.757	0.010 SEC

WAIT TIME FOR ORIGINAL ALLOCATION	DURATION OF USER ACTIVITY	# USER SWAPS	ACT/HOUR (THROUGHPUT)	USER ACT/HOUR (THROUGHPUT)	SWAPS PER HOUR
25.500	70.050 SEC	125	177.920	121.770	185.361

Figure 6-24. Memory Statistics Report

SPECIAL JOB MEMORY DEMAND REPORT ON SYSTEM NMCC ON 82-06-18

TIME	DEMAND TYPE	TIME TO SATISFY IN .1 SEC	SNUMB	SIZE OF DEMAND IN K	CURRENT SIZE IN K	CALL
11.78394(114703)	GEMORE	-99999	FTS	1	35	82
11.78417(114703)	MET	0	FTS	1	36	83
11.78578(114709)	GEMORE	-99999	FTS	4	36	86
11.78628(114710)	DENIED	10	FTS	4	36	88
11.78633(114710)	SWAP	2	FTS	40	40	91
11.79422(114739)	GEMREL	-99999	FTS	-1	39	93

**** SINCE LAST PRINT OUT FTS HAS ACCUMULATED 205 .1 SEC WAIT TIME, CURRENT TIME IS 1148

Figure 6-25. Special Job Memory Demand Report

SPECIAL JOB MEMORY SIZE REPORT ON SYSTEM NMCC ON 82-03-01

TIME	SNUMB	AVG SIZE	SNUMB	AVG SIZE	SNUMB	AVG SIZE
1415	FTS	63	NCP	29	TS1	124
1426	FTS	56	NCP	29	TS1	120
1436	FTS	47	NCP	29	TS1	81
1446	FTS	38	NPC	29	TS1	114

Figure 6-26. Special Job Memory Size Report

DISTRIBUTION OF URGENCY REPORT OVERTIME

[illegible]	0	5	10	15	20	25	30	35	40	45	50	55U	60U	55S	60S
14:06	0.	41.9	1.5	1.2	1.3	0.7	0.9	0.	0.1	0.	0.	16.7	9.1	8.7	18.1
14:16	0.	39.9	0.	0.	0.2	0.	5.8	0.	0.	0.	0.	20.6	8.6	8.4	16.5
14:26	0.	44.3	0.	0.	0.3	0.	2.2	0.	0.	0.	0.	12.5	10.6	9.6	20.6
14:36	0.	40.3	0.	0.	0.1	0.1	5.2	0.1	1.4	0.	0.	9.6	9.0	7.1	27.0
14:46	0.	46.5	0.8	1.1	2.3	0.1	1.9	0.	0.1	0.5	4.2	6.4	6.9	6.2	22.9
14:56	0.	40.6	0.	0.	0.1	0.	6.1	0.	0.0	0.	7.2	11.9	7.1	5.5	21.4
15:06	0.	45.3	3.1	0.	0.2	0.	3.4	0.	0.	0.5	1.1	9.2	8.9	5.5	22.8
15:16	0.	46.0	0.2	1.1	2.9	0.	5.7	0.	0.	0.	0.	7.8	8.3	5.1	22.8
15:26	0.	45.0	0.	0.	3.5	0.	7.4	0.	0.	0.	0.	11.1	11.9	5.1	16.2
15:36	0.	41.9	0.	0.	0.0	0.	7.3	0.	0.	0.4	6.3	12.3	8.3	6.4	17.0
15:46	0.	48.5	0.	0.	0.2	0.	5.5	0.	0.0	1.0	3.1	10.1	6.0	5.8	19.8
15:56	0.	48.5	0.	0.	0.4	0.	5.9	0.	0.	1.3	0.	12.5	6.8	6.6	18.0

SUMMARY OF ACTIVITIES OVER THE TOTAL DURATION

	0.	40.2	0.	0.	10.0	0.4	3.3	0.	2.9	7.1	2.9	22.2	7.1	0.4	3.3

Figure 6-27. Distribution of Urgency Over Time Report

ZERO URGENCY JOB REPORT FOR SYSTEM NMCC ON 82-08-25

SNUMB-ACT	USERID	START	STOP	TIME AT ZERO URGENCY AND DEMAND FOR MEMORY WAS PRESENT (.1 SEC)	% OF MEMORY TIME JOB WAS IN ZERO URGENCY STATE AND MEMORY DEMAND PRESENT
$PALC - 1		1630	1715	835	3%
$RTIN - 0		1630	1715	75	0%
$IOGN - 0		1630	1715	51	0%
TLCF - 1		1630	1715	15	0%
DMSTA - 2	OPNSUTIL	1630	1715	9	0%
FSS - 1	DJ8XI342RBR	1714	1715	81	62%
NACE - 1	OPNSUTIL	1630	1715	1167	4%

Figure 6-28. Zero Urgency Job Report

6.5 Multireel Processing

If more than a single reel of data has been collected, a series of messages will be printed at the computer console informing the operator that a new data reel is required. The following are the messages produced.

a. DISMOUNT REEL #XXXXX THEN MOUNT REEL NUMBER YYYYY ON ZZZZZ

In this case, XXXXX is the old reel number, YYYYY is the new reel number, and ZZZZZ is the tape drive ID.

If the operator fails to mount the new tape, the above message will be repeated three times, after which the program will terminate, and all reports will be produced.

b. IS TAPE XXXXX MOUNTED ON DRIVE ID YYYYY (Y/N)

In this case, XXXXX is the tape number being requested and YYYYY is the appropriate tape drive ID.

This message occurs when the data reduction program finds the wrong tape has been mounted (by comparing internally generated tape labels). If the operator answers N, the message in (c) below is produced. If the operator answers Y, the data reduction program will terminate and all reports will be produced. In this case, the data reduction program is unable to process the tape. Even though the operator is mounting the correct tape, the internal label on the new tape does not match that being requested by the old tape. The user should check the data collection session to insure that the operator did not respond with an incorrect tape number during multireel change.

After entering the Y or N, the operator will need to hit the EOM key twice in order for the response to be transmitted.

c. WRONG REEL JUST MOUNTED, DISMOUNT AND MOUNT REEL XXXXX ON ZZZZZ

In this case, XXXXX is the new reel number, and ZZZZZ is the tape drive ID.

d. CAN TAPE XXXXX BE MOUNTED ON DRIVE YYYYY (Y/N)

In this case, XXXXX is the new desired reel and YYYYY is the tape drive ID.

If the operator fails to answer this message, it will be repeated until he responds with a "Y" for YES or "N" for NO. If he types in "Y", then message (a) will be repeated. If he types in "N", then the program will be terminated and all reports will be produced.

SPECIAL JOB MEMORY SIZE REPORT ON SYSTEM NMCC ON 82-03-01

TIME	SNUMB	AVG SIZE	SNUMB	AVG SIZE	SNUMB	AVG SIZE
1415	FTS	63	NCP	29	TS1	124
1426	FTS	56	NCP	29	TS1	120
1436	FTS	47	NCP	29	TS1	81
1446	FTS	38	NPC	29	TS1	114

Figure 6-26. Special Job Memory Size Report

DISTRIBUTION OF URGENCY REPORT OVERTIME

URGENCY	0	5	10	15	20	25	30	35	40	45	50	55U	60U	55S	60S
14:06	0.	41.9	1.5	1.2	1.3	0.7	0.9	0.	0.1	0.	0.	16.7	9.1	8.7	18.1
14:16	0.	39.9	0.	0.	0.2	0.	5.8	0.	0.	0.	0.	20.6	8.6	8.4	16.5
14:26	0.	44.3	0.	0.	0.3	0.	2.2	0.	0.	0.	0.	12.5	10.6	9.6	20.6
14:36	0.	40.3	0.	0.	0.1	0.1	5.2	0.1	1.4	0.	0.	9.6	9.0	7.1	27.0
14:46	0.	46.5	0.8	1.1	2.3	0.1	1.9	0.	0.1	0.5	4.2	6.4	6.9	6.2	22.9
14:56	0.	40.6	0.	0.	0.1	0.	6.1	0.	0.0	0.	7.2	11.9	7.1	5.5	21.4
15:06	0.	45.3	3.1	0.	0.2	0.	3.4	0.	0.	0.5	1.1	9.2	8.9	5.5	22.8
15:16	0.	46.0	0.2	1.1	2.9	0.	5.7	0.	0.	0.	0.	7.8	8.3	5.1	22.8
15:26	0.	45.0	0.	0.	3.5	0.	7.4	0.	0.	0.	0.	11.1	11.9	5.1	16.2
15:36	0.	41.9	0.	0.	0.0	0.	7.3	0.	0.	0.4	6.3	12.3	8.3	6.4	17.0
15:46	0.	48.5	0.	0.	0.2	0.	5.5	0.	0.0	1.0	3.1	10.1	6.0	5.8	19.8
15:56	0.	48.5	0.	0.	0.4	0.	5.9	0.	0.	1.3	0.	12.5	6.8	6.6	18.0
SUMMARY OVER THE TOTAL DURATION															
	0.	40.2	0.	0.	10.0	0.4	3.3	0.	2.9	7.1	2.9	22.2	7.1	0.4	3.3

Figure 6-27. Distribution of Urgency Over Time Report

THIS PAGE LEFT INTENTIONALLY BLANK

6.5 Multireel Processing

If more than a single reel of data has been collected, a series of messages will be printed at the computer console informing the operator that a new data reel is required. The following are the messages produced.

a. DISMOUNT REEL #XXXXX THEN MOUNT REEL NUMBER YYYYY ON ZZZZZ

In this case, XXXXX is the old reel number, YYYYY is the new reel number, and ZZZZZ is the tape drive ID.

If the operator fails to mount the new tape, the above message will be repeated three times, after which the program will terminate, and all reports will be produced.

b. IS TAPE XXXXX MOUNTED ON DRIVE ID YYYYY (Y/N)

In this case, XXXXX is the tape number being requested and YYYYY is the appropriate tape drive ID.

This message occurs when the data reduction program finds the wrong tape has been mounted (by comparing internally generated tape labels). If the operator answers N, the message in (c) below is produced. If the operator answers Y, the data reduction program will terminate and all reports will be produced. In this case, the data reduction program is unable to process the tape. Even though the operator is mounting the correct tape, the internal label on the new tape does not match that being requested by the old tape. The user should check the data collection session to insure that the operator did not respond with an incorrect tape number during multireel change.

After entering the Y or N, the operator will need to hit the EOM key twice in order for the response to be transmitted.

c. WRONG REEL JUST MOUNTED, DISMOUNT AND MOUNT REEL XXXXX ON ZZZZZ

In this case, XXXXX is the new reel number, and ZZZZZ is the tape drive ID.

d. CAN TAPE XXXXX BE MOUNTED ON DRIVE YYYYY (Y/N)

In this case, XXXXX is the new desired reel and YYYYY is the tape drive ID.

If the operator fails to answer this message, it will be repeated until he responds with a "Y" for YES or "N" for NO. If he types in "Y", then message (a) will be repeated. If he types in "N", then the program will be terminated and all reports will be produced.

6.6 Tape Error Aborts

During the course of processing, it is possible that the operator will be required to abort the data reduction program due to an irrecoverable tape error. If such a condition occurs, the operator should abort this job with a "U" abort. This will allow the data reduction program to enter its wrap-up code processing and produce all reports generated prior to the tape error.

SECTION 7. MASS STORE MONITOR DATA REDUCTION PROGRAM (MSMDRP)

7.1 Introduction

The Mass Store Monitor Data Reduction Program is a FORTRAN program that sequentially processes data the Mass Store Monitor collected and wrote on tape. MSMDRP produces a number of reports depicting the physical and logical usage of the mass storage subsystem during the monitoring period. A list of these reports is found in table 7-1 and report descriptions are presented in subsection 7.5.

The Mass Store Monitor and its Data Reduction Program were conceived as a response to the need for information on the rate and characteristics of usage of mass store subsystems. The information collected is applicable in the following areas:

- Discovery of improper configurations (software or hardware), e.g., not alternating usage of dual physical channels configured on a subsystem.
- Discovery of improper device utilization, e.g., use of only a small number of the devices configured instead of having the activity spread over all configured devices.
- Discovery of open file allocation on a device which can cause long and frequent arm movement.
- Identification of files (either system or user data base) which are frequently accessed and quantification of the rates of access to them.

There are two inputs to the MSMDRP. The first is the data tape produced by the MSM in the General Monitor Collector. The second input is a set of report option control cards used to alter the reports in some way other than the standard default. The various user input options and their formats are described in subsection 7.6. The actual reports produced by the MSMDRP are described in subsection 7.5.

The Mass Store Monitor Data Reduction Reports can be used to assist the analyst in applying models in the area of system performance evaluation by providing computer system resource usage information in a format conducive to defining the system workload. Models of computing systems are frequently used to evaluate configurations which are not currently available or are perhaps unrealized. The usefulness of this type of model depends upon two basic characteristics: (1) how faithfully the model represents the operation of the system being modeled, and (2) how faithfully the input parameters to the model represent the workload to be placed on the system. MSM and MSMDRP are key contributors to such model development.

Table 7-1. MSM/MSMDRP Reports (Part 1 of 2)

1 System Configuration and Channel Usage Report (File 42)

2 System Summary Report (File 42)

3 System Traces Captured by Monitor Report (File 42)

4 Channel Status Changes Report (File 29)

5 Physical Device, Device ID Correlation Table Report (File 42)

6 Device Space Utilization Report (File 42)

7 Device Seek Movement Report (File 42)

8 Head Movement Efficiency Report (File 42)

9 System File Use Summary Report (File 21) (SYSFILES)

10 Individual Module Activity Report (File 21) (SYSFILES)

11 SSA Module Usage Report by Job (File 21)

12 File Code Summary Report (File 23) (FILECODE)

13 CAT/File String Report (File 23)

14 Connect Summary Report by Userid/SNUMB (File 23)

15 Activity Summary Report (File 24) (ACTIVITY)

16 Device Area File Code Reference Report (File 22) (AREA)

17 Device File Use Summary Report (File 21) (FILEUSE)

18 Chronological Device Utilization Report (File 26) (CHRONO)

19 FMS Cache Report (File 21)

20 Connects Per 10 Minute Report (File 20) (RATE)

21 Proportionate Device Utilization Report (File 42)

22 Elapsed Time Between Seeks Report (File 42)

23 Data Transfer Size Report (File 42)

24 Data Transfer Sizes for TSS Swap Files (File 42)

7.2 Data Collection Methodology

The MSM in the General Monitor Collector processes GCOS trace types 7 and 15 and collects information to monitor the usage of the entire disk subsystem. The information collected on the occurrence of the above traces enables the MSMDRP to identify the activity issuing the I/O request, the file being accessed, the disk pack upon which the file is located, arm movement required in order to accomplish the requested file accessing, and the type of accessing being requested; e.g., read, write, write verify, etc.

If the system being monitored by the MSM is configured with SSA Cache Core, the MSM will create two direct transfer traces (types 73 and 76) in order to collect data to analyze the effectiveness of SSA Cache Core. The method for generating these new direct transfer traces is described in subsection 5.2.2, and the formats for the MSM generated records used by the MSMDRP are described in subsection 5.4.3.

Finally, if the system being monitored by the MSM is configured with FMS Catalog Cache or is utilizing disk in core space tables, a data record is generated so that the MSMDRP can report on the effectiveness of FMS Catalog Cache or in core space table buffering.

7.3 Analytical Methodology

An evaluation of the Mass Storage Subsystem reports produced by the MSMDRP requires concurrent use of the reports produced by the Channel Monitor Data Reduction Program (CMDRP). Chapter 14 provides a detailed description of the procedure to be followed in such an evaluation. Subsection 8.3 provides a detailed description of the entire I/O process, and the traces generated during the processing of an I/O request. In general, the CMDRP is used to identify channels and/or devices which are acting as bottlenecks to the efficient operation of the system, while the MSMDRP reports are used to determine the exact activities, files, and file codes that are causing the contention uncovered by the CMDRP reports. The MSMDRP reports will also identify those devices experiencing seek elongation problems and the files upon these devices which are responsible for the seek elongation. Finally, the MSMDRP reports will identify those files that are candidates for device relocation or placement into Hard Core or SSA Cache Buffer space.

Before a user conducts a Mass Storage Subsystem Evaluation, it is important to have an understanding of the entire I/O process. Subsection 8.3 provides a detailed description of the entire I/O process and all traces generated during the processing of an I/O request. In this subsection, a description of only the connect (trace type 7) event will be presented.

Each time a system program or application program issues an I/O request (read disk/tape, write disk/tape, seek, etc. . .) the GCOS system will generate a trace type 7 (connect event). Upon the occurrence of this event, several internal tables are updated and it is these tables that the

MSM references in order to generate its data record. A program's SSA area contains tables for the Peripheral Assignment Table (PAT) and the PAT Pointer. These are used to describe the device and space allocation for a particular file and the file code to correlate a user file code to the PAT and the device on which that file is allocated. The .CRIO and .CRCT tables contain descriptive information concerning device and channel configuration. Finally, the program's SSA area also contains an area which is used for I/O entries. These entries are each 11 words long and contain detailed information concerning the I/O just requested. They are referred to as the 11 word I/O queue entry.

I/O requests can be of two types (single or multicommand). Multicommands are of the type seek-read, seek-write, or seek-write verify. Single commands can be status requests of certain types, or reads/writes, where seeks are not required. These different types of I/O commands are processed and reported in different fashions by the various MSMDRP reports (see individual output reports). Finally, whenever the system generates a multi I/O command, it is necessary for the system to record the actual seek address being requested. Normally, this seek address is stored in I/O queue word number 4. Whenever the MSM processes a multicommand, it expects to find a valid seek address at this location. However, there are certain occurrences when a multicommand is issued and I/O queue word 4 does not contain a valid seek address. In these cases, Bit 32 of I/O queue word 2 is set to a 0. The Computer Performance Evaluation Office is currently conducting studies to determine exactly when this nonstandard procedure occurs and if there is any manner in which the MSM can determine the correct seek address. Under the current processing procedures, the MSMDRP recognizes the seek address as being invalid and performs certain special processing in order to recover from this event, thereby affecting several of the output reports (see individual output reports). This software release has a code included to correct this problem. Therefore, no occurrences of this condition should be reported in any of the MSM output.

7.4 Data Reduction Methodology

The MSMDRP currently uses random I/O (File 58) to process histogram data for the Device Space Utilization and Device Seek Movement reports. This feature allows the MSMDRP to process an unlimited number of devices with a minor increase in memory requirements. As delivered, the MSMDRP will process data describing 75 mass storage devices and 40 mass storage channels. It will produce 64 unique histograms with no random I/O. If the number of channels or devices is insufficient, the user will need to edit file B29IDPXO/SOURCE/MSM. The user should enter the edit subsystem and process the following command:

```
B RS:/NRDEVXX=XX75/;*:/NRDEVXX=XX new number of devices/

B RS:/NRCHANXX=XX40/;*:/NRCHANXX=XX number of new channels/
```

For each additional device, the size of the program will increase by 10 words and for each additional channel, the program will increase by 45 words. For the above edit, the character "X" signifies a space.

The next variable that will need to be changed is RPTCNT. This number represents the total number of histograms and reports that will be processed with no random I/O. To calculate the value required, the following formula should be used.

(number of devices actually configured)*2+8

If this value is less than 72 (32 disk devices), no change is required. If the value required is greater than 72, the user may alter this value. This will help to limit the amount of random I/O being performed but will increase storage by 80 words for each increment above 72. This trade-off between CPU/IO time and memory must be made at the discretion of the user. In order to change this value, the following edit function should be performed:

B RS:/RPTCNTX=X72/;*:/RPTCNTX=Xnew value/ (11 occurrences)

As in the earlier edit example, the character "X" should not be typed, but is being used to represent a blank column.

After performing the above edits, the user should recompile the source program by entering the card subsystem and issuing a run command.

7.5 MSMDRP Output

The MSMDRP p[illegible]uces a series of 24 reports listed in table 7-1 over which the user has [illegible]imited control. Those eight reports with a NAME=codename designation offer greater parameter control to the user. This parameter control will be described in subsection 7.6. In table 7-1, the file nn designation indicates the file code used to record the given report and is of no real concern to the user. In addition, a series of messages are produced which supply the user with information concerning special processing events that occurred during the execution of the data reduction program. Most of these processing messages are for information only, and can be ignored. The following subsections will describe all the reports listed in table 7-1, and subsection 7.5.25 will describe the processing messages that may be produced during the course of data reduction.

7.5.1 System Configuration and Channel Usage Report (File 42). This report documents the system identification, configuration, and the date and time of the monitoring period, as well as reporting the usage of all configured I/O channels. Figure 7-1 is an example of this report. The heading line indicates the software version number that corresponds to this document. The version number should be 01-82. The first line after the heading provides the tape number(s) the report was generated from, the system identification, the date (in the form year, month, and day - YYMMDD), and the start and stop times (HH:MM:SS) of the MONITORING SESSION. The next several lines of output describe the overhead of all GMF

```
**************************************************************
* * *       M A S S   S T O R A G E   M O N I T O R      * * *
* * *                  VERSION  01-82                     * * *
**************************************************************

          TAPE # 24502
SYSID                DATE        START TIME     STOP TIME

OSCC2 ,6.4.ID        82-01-08      13:49:04       14:31:40       FOR A TOTAL OF 0.71 HOURS

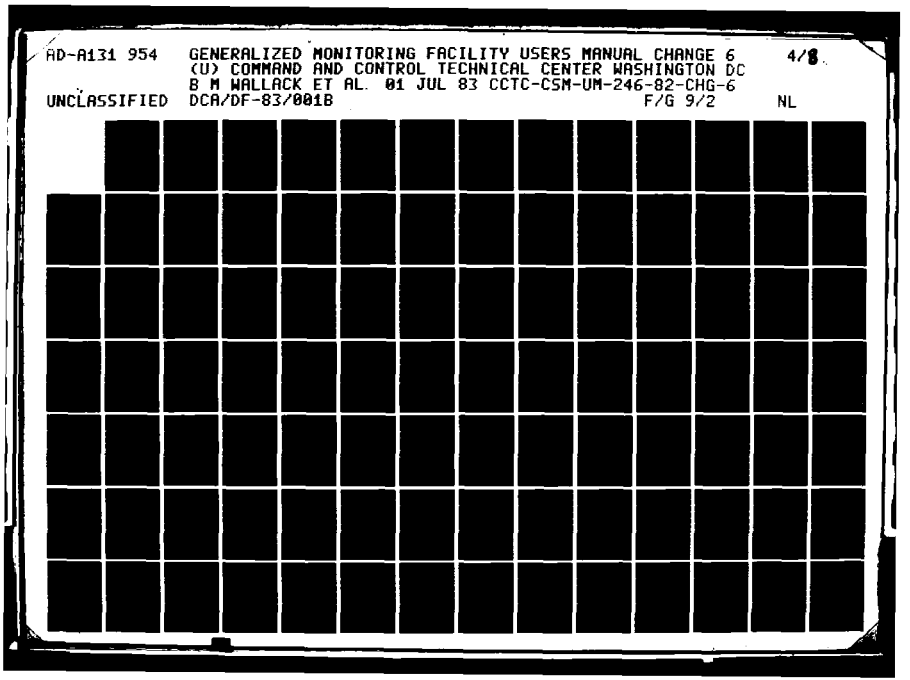
          MONITOR              TIME(SEC)          % OVERHEAD
            EXEC                 100                1.96
            MUM                  120                2.35
            MSM                  160                3.13
            NAME                  40                 .78
            FMS                   50                 .98
            TOTAL                                    8.42

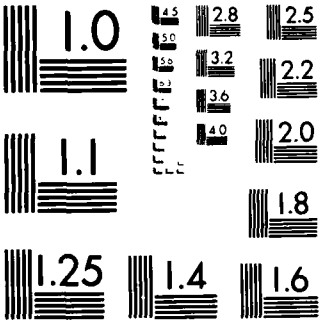
CONFIGURATION:  DUAL    PROCESSOR 6680,  DUAL  IOM,  512K  MEMORY - 52 OF WHICH WERE HCM
                THE HCM DOES NOT INCLUDE 3K FOR .CALC AND 6-8K FOR FILSYS
              CPUS ACTUALLY CONFIGURED = 2.00  CPUS ACTUALLY AVAILABLE = 1.75
              MAX # OF SYSTEM INTERCOM IOS ALLOWED ARE  128
              MAX # OF OUTSTANDING SYSTEM INTERCOM IOS GENERATED WAS   26
                                  IOM NUMBER 0

CHANNEL          TYPE          CROSSBAR        CONNECTS
 0-08            .DS450          0-09            47750
                                 1-08
                                 1-09
 0-09            .DS450        SEE ABOVE         16851
 0-12            .DS191          0-13            55524
                                 1-12
                                 1-13
 0-13            .DS191        SEE ABOVE          2479
 0-16            .DS181          0-17             8290
                                 1-16
                                 1-17
 0-18            .DMTA9          NONE               62
 0-19            .DMTA9          NONE             9147

(IN THE ACTUAL REPORT A SUMMARY OF IOM NUMBER 1 REPORT WOULD FOLLOW)
```

Figure 7-1. System Configuration and Channel Usage Report


AD-A131 954
GENERALIZED MONITORING FACILITY USERS MANUAL CHANGE 6
(U) COMMAND AND CONTROL TECHNICAL CENTER WASHINGTON DC
B M WALLACK ET AL. 01 JUL 83 CCTC-CSM-UM-246-82-CHG-6
UNCLASSIFIED
DCA/DF-83/001B
F/G 9/2
NL
4/8



MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS 1963-A

monitors that were active during data collection. The monitor name is given, its CPU time in seconds, and its overhead as a function of total processor power. The GMF executive overhead is separated from the actual monitors and is listed as "EXEC". The monitor "NAME" is an area of code within the Mass Store Monitor and even though listed separately it is also included under the monitor "MSM". The monitor "FMS" is also an area of code within the Mass Store Monitor, but in this case it has not been included under the monitor "MSM".

Monitor "CM" in this report describes the processor overhead of subroutine T4 (terminate processing) and subroutine T22 (start I/O processing). Monitor "MSM" in this report describes the processor overhead of subroutine T7 (connect processing). Therefore, if the Channel Monitor was active, but the Mass Store Monitor was not, this report will still list both "CM" and "MSM" as contributing to the processor overhead. The total Channel Monitor overhead will be found by adding the overhead of the "CM" monitor to the overhead of the "MSM" monitor, to the overhead of the "FMS" monitor.

If both the Channel Monitor and Mass Store Monitor were active, then the combined overhead of both monitors can be found as the sum of "MSM" + "CM" + "FMS".

For purposes of this report, % overhead is computed as:

$$\frac{\text{(CPU TIME Used by Monitor)}}{\text{(TOTAL Elapsed Time)x(Number of Processors)}}$$

Following the overhead description are five lines of configuration information describing the number of processors, IOMs, and amount of memory configured to the system. In addition, the size of GCOS Hard Core, the size of the Core Allocator and the size of FILSYS is also presented. The third line of the configuration data indicates the number of processors actually configured and actually available. These numbers might be different than shown on the first line due to the assigning and releasing of processors. In figure 7-1, we see that one processor was released for a period of time (i.e., CPUs actually available is equal to 1.75). The actual time that processors were available or released is indicated in the status message printouts (see subsection 7.5.25).

The final two lines of the report indicate how many simultaneous intercom I/Os are permitted and the maximum number of outstanding intercom I/Os recorded during the monitoring session. The intercom I/O feature is the means by which two programs residing in the H6000 can pass data to one another. This capability is used extensively by FTS and TELNET to pass data to and from NCP. If the system exhausts all available intercom I/O buffers, then the programs requiring this facility will be delayed. These two figures would provide an indication that the system did indeed exhaust all available intercom I/O buffer space whenever both lines contain the same number. Whenever this occurs, it shows that the entire buffer table was exhausted and therefore, the probability is high that jobs are being delayed while they are waiting for buffer space to free.

In addition, whenever the number of outstanding intercom I/Os is found to be equal to the total buffer capability, a warning message will be printed on the status message printout file. This message will indicate the time of day that the buffer pool was exhausted and a second printout will occur whenever the number of outstanding intercom I/Os falls back below the maximum allowed.

The next portion of the report documents the channel configuration by IOM, listing each configured channel number, the device type configured to that channel, and the channel crossbarring. The crossbar column shows those channels that are crossbarred to the channel identified under the channel column. If SEE ABOVE is found, the crossbarring has been displayed on a preceding channel. The I-CC format of each channel description identifies the IOM and the channel number being discussed. The last column of this report displays the number of all connect types issued over that channel. This section will be repeated for each IOM configured to the system. Figure 7-1 only displays IOM 0 activity. This report is always generated and cannot be turned off.

THIS PAGE LEFT INTENTIONALLY BLANK

```
**********************************************************************
***        M S M   S Y S T E M   S U M M A R Y   R E P O R T       ***
**********************************************************************

TOTAL CONNECTS TO DSS181          8540 OF    207117 ( 4%) AT A RATE OF    12024. PER HOUR
TOTAL CONNECTS TO DSS191         76949 OF    207117 (37%) AT A RATE OF   108341. PER HOUR
TOTAL CONNECTS TO MS0450        107273 OF    207117 (51%) AT A RATE OF   151036. PER HOUR
TOTAL CONNECTS TO MPC-CONTROL       12 OF    207117 ( 0%) AT A RATE OF       16. PER HOUR

                    COMMANDS PER CHANNEL
                       IOM NUMBER 0

   CHANNEL        8        9       12       13       16      17
   COUNT      47750    16851    55524     2479     8290       0
                       IOM NUMBER 0

   CHANNEL        8        9       12       13       16      17
   COUNT      38189     4487    18773      177      254       0

                 COMMAND          COUNT
                 01 CONTRL           12
                 17 FMT TK        15485
                 25 READ         126899
                 26 RD CR           192
                 31 WRITE         35483
                 33 WR-VER        12862
                 40 RST ST         1841
                 TOTAL           192774
```


Figure 7-2. MSM System Summary Report

7.5.2 System Summary Report (File 42). The System Configuration and Channel Usage Report and the System Summary Report may be used to assess overall system utilization. Figure 7-2 is an example of the System Summary Report. The first set of lines shows the number of connects to the monitored mass storage subsystems compared to the total connects issued (TAPE+DISK) and the connect rate per hour over the subsystem. Most systems will show a small number of Control Connects being generated by the MPCs configured to the system. These Control Connects will be summed together and listed as a separate subsystem line. Analysis on a Shared Mass Storage System shows the number of MPC connects generated to be a significant percentage of the total connects generated. The next lines show the breakdown of the mass storage connects by the IOM channel over which they were issued. The final part of this report is a list of the commands (octal code and mnemonic) issued to the mass storage subsystem and the count of each issued during the monitoring session. This report is always generated and cannot be turned off.

A well performing system, under a heavy workload, should show a high utilization of the configured resources. Figure 7-2 shows that the I/O activity is predominantly on the MSU450 subsystem configured on channels 8 and 9 of IOM 0 and 1 (see figure 7-1). The MSU450s are receiving 51% of all connects and, therefore, should be the major area of concern. The access rate for every subsystem is reported on the top of the System Summary Report and it can be seen that the MSU450s have an access rate significantly higher than the other subsystems. All signs indicate that if system throughput is being affected by disk activity, then the MSU450s would be the probable cause of such problems.

The next item to check on these two reports should be the channel usage. The two highest used logical channels of any subsystem should be on a separate PSI channel of a two-PSI channel subsystem. Referring to figure 7-2, one can see that logical channel 8 of IOM 0 and IOM 1 has the highest usage, and this is the proper configuration. If the highest used logical channels are not on separate PSI channels, the $ XBAR card in the startup configuration section is suspected as the cause. The channels are used in the order given on the $ XBAR card (i.e., if the primary channel is busy, the next channel tried is given on the crossbar). The alternate use of PSI channels for maximum simultaneity must, therefore, be appropriately specified in the boot deck. Subsection 8.3 provides a detailed explanation for analyzing the correctness of the crossbar configuration.

While looking at the System Summary Report, it is also of interest to note the ratio of READ commands to WRITE commands (over two to one in this example). This gives an indication of the nature of the usage of the mass storage space. A quick look at the number of write/verify (WR-VER) commands executed is also of interest as they are essentially double (WRITE, then READ) data transfer commands which require more device and channel time.

The general fraction of utilization for each logical channel gives an indication of the degree of simultaneity of access to the subsystem. If only N of the configured logical channels have nonzero counts, then there were never more than N accesses being performed simultaneously by the

subsystem. The proportional relationships among the counts of accesses made over each of the logical channels are quantitative indications of the frequency of occurrence of specific levels of simultaneity. As an example, if we look at figure 7-2, we see that only 4487 connects out a total of 107,273 connects went to Channel 9, IOM 1. This means that only 4487 times during the measuring period were all four MSU450 disk channels being utilized simultaneously. In this example, channel queuing (i.e., shortage of channel power) would not appear to be a problem. This is not to infer that device queuing is not a problem, just that channel queuing does not appear to be a problem. If the number of accesses to the lowest priority channel is a larger percentage of the total accesses, then channel queuing needs to be examined. Queuing for devices and/or channels can be analyzed by running the Channel Monitor Data Reduction Program (see section 8).

7.5.3 System Traces Captured by Monitor Report (File 42). This report contains the number of occurrences of each specific trace type record on the data collector tape processed by the MSMDRP (figure 7-3). This report provides little, if any, information required by the user for his analysis. This report is always generated and cannot be turned off.

7.5.4 Channel Status Changes Report (File 29). This report lists the initial status for all tape and disk channels configured to the system (figure 7-4). If, during the course of the monitoring session, a given channel or IOM was dropped or added to the system (dynamic reconfiguration) a new report will be produced indicating the activation of deactivation changes and the time that the change occurred. Finally, this report will indicate whether the SSA cache option and FMS cache option are active, and if so, will indicate their initial status and any changes that occur to that status. If a given option is not active, a zero will be reported for each of the values. This report is always generated and cannot be turned off.

7.5.5 Physical Device, Device ID Correlation Table (File 42). Each mass storage device configured in the system is listed with a unique device ID. A typical report is presented in figure 7-5. This unique device is needed since different devices can have the same device number on the Honeywell 6000. (See Device ID 1, Device ID 7, and Device ID 18 of figure 7-5). These unique numbers are referenced in several reports produced by the MSMDRP. This report is always generated and cannot be turned off.

7.5.6 Device Space Utilization Report (File 42). The device space utilization histogram report is produced for every device on the mass storage subsystem and shows the distribution of access to the device space. Figure 7-6 is an example. It should be noted that the name of the device is also given. This example presents all connects made to the device with the name RF5. If an exchange took place and the RF5 disk pack was moved from 0-08-05 to 0-08-01 the data reduction program will account for that exchange and any connects that are made to 0-08-01 will be

CHANNEL STATUS CHANGES REPORT FOR NMCC2 ON 80/09/20

IOM	CHANNEL	DEACTIVATE/ACTIVATE CHANGE	TIME
0	08	INITIAL ACTIVE	16:50:49.7
0	09	INITIAL ACTIVE	16:50:49.7
0	12	INITIAL ACTIVE	16:50:49.7
0	13	INITIAL ACTIVE	16:50:49.7
0	14	INITIAL ACTIVE	16:50:49.7
0	15	INITIAL ACTIVE	16:50:49.7
0	16	INITIAL ACTIVE	16:50:49.7
0	17	INITIAL ACTIVE	16:50:49.7
0	18	INITIAL ACTIVE	16:50:49.7
1	08	INITIAL ACTIVE	16:50:49.7
1	09	INITIAL ACTIVE	16:50:49.7
1	12	INITIAL ACTIVE	16:50:49.7
1	13	INITIAL ACTIVE	16:50:49.7
1	14	INITIAL ACTIVE	16:50:49.7
1	15	INITIAL ACTIVE	16:50:49.7
1	16	INITIAL ACTIVE	16:50:49.7
1	17	INITIAL ACTIVE	16:50:49.7

INITIAL VALUES FOR SSA CACHE - LAL, MBA, SIZE
000000200000 000000001566 10

INITIAL VALUES FOR FMS CACHE - ABS ADDR, MBA, OPTION WORD, #320 WORD BUFFERS
000000000000 000000000000 000000000000 0

Figure 7-4. Channel Status Changes Report

SYSTEM TRACES CAPTURED BY MONITOR REPORT

TRACE	COUNT	TRACE	COUNT	TRACE	COUNT	TRACE	COUNT	TRACE	COUNT	TRACE	COUNT	TRACE	COUNT	TRACE	COUNT
00	0	01	0	02	0	03	0	04	3214	05	0	06	0	07	3214
10	129	11	0	12	0	13	0	14	636	15	0	16	0	17	0
20	0	21	450	22	3214	23	0	24	0	25	0	26	0	27	0
30	0	31	0	32	0	33	0	34	0	35	0	36	0	37	0
40	0	41	0	42	0	43	0	44	0	45	0	46	15	47	0
50	21	51	0	52	0	53	0	54	0	55	0	56	0	57	0
60	0	61	0	62	0	63	0	64	0	65	0	66	0	67	0
70	0	71	0	72	0	73	0	74	0	75	0	76	0	77	0

Figure 7-3. System Traces Captured by Monitor Report

THE PHYSICAL DEVICE, DEVICE ID CORRELATION TABLE

```
DEVICE ID - 1 IS FOUND ON IOM-0 PUB-08 AND IS DEVICE #01
DEVICE ID - 2 IS FOUND ON IOM-0 PUB-08 AND IS DEVICE #02
DEVICE ID - 3 IS FOUND ON IOM-0 PUB-08 AND IS DEVICE #03
DEVICE ID - 4 IS FOUND ON IOM-0 PUB-08 AND IS DEVICE #04
DEVICE ID - 5 IS FOUND ON IOM-0 PUB-08 AND IS DEVICE #05
DEVICE ID - 6 IS FOUND ON IOM-0 PUB-08 AND IS DEVICE #06
DEVICE ID - 7 IS FOUND ON IOM-0 PUB-12 AND IS DEVICE #01
DEVICE ID - 8 IS FOUND ON IOM-0 PUB-12 AND IS DEVICE #02
DEVICE ID - 9 IS FOUND ON IOM-0 PUB-12 AND IS DEVICE #03
DEVICE ID -10 IS FOUND ON IOM-0 PUB-12 AND IS DEVICE #04
DEVICE ID -11 IS FOUND ON IOM-0 PUB-12 AND IS DEVICE #05
DEVICE ID -12 IS FOUND ON IOM-0 PUB-12 AND IS DEVICE #06
DEVICE ID -13 IS FOUND ON IOM-0 PUB-12 AND IS DEVICE #07
DEVICE ID -14 IS FOUND ON IOM-0 PUB-12 AND IS DEVICE #08
DEVICE ID -15 IS FOUND ON IOM-0 PUB-12 AND IS DEVICE #09
DEVICE ID -16 IS FOUND ON IOM-0 PUB-12 AND IS DEVICE #10
DEVICE ID -17 IS FOUND ON IOM-0 PUB-12 AND IS DEVICE #11
DEVICE ID -18 IS FOUND ON IOM-0 PUB-16 AND IS DEVICE #01
DEVICE ID -19 IS FOUND ON IOM-0 PUB-16 AND IS DEVICE #02
DEVICE ID -20 IS FOUND ON IOM-0 PUB-16 AND IS DEVICE #03
DEVICE ID -21 IS FOUND ON IOM-0 PUB-16 AND IS DEVICE #04
DEVICE ID -22 IS FOUND ON IOM-0 PUB-16 AND IS DEVICE #05
DEVICE ID -23 IS FOUND ON IOM-0 PUB-16 AND IS DEVICE #06
DEVICE ID -24 IS FOUND ON IOM-0 PUB-16 AND IS DEVICE #07
```

Figure 7-5. Physical Device, Device ID Correlation Table

DISTRIBUTION COLLECTED ON SYSTEM NMCC2 AT 18:36:38 ON 81-03-13

SPACE UTILIZATION OF IOM-0,PUB-08,DEVICE-05-- MSO450 RF5

INDIV. NUMBER	CUMUL. NUMBER	CUMUL. PRC.	INDIV. PRC.	CYLNDR NUMBER	PERCENT OF OCCURRENCE 00 05 10 15 20 25 30 35 40 45 50	REPORT 49
					I....I....I....I....I....I....I....I....I....I....I	
564	564	27.606	27.606	0- 16	Ixxxxxxxxxxxxxxxxxxxxxxxxxxx	
67	631	30.886	3.279	17- 33	Ixxx	
207	838	41.018	10.132	34- 50	Ixxxxxxxxx	
0	838	41.018	0.	51- 67	I	
0	838	41.018	0.	68- 84	I	
0	838	41.018	0.	85- 101	I	
0	838	41.108	0.	102- 118	I	
21	859	42.046	1.028	119- 135	Ix	
.	.	.	.	. .	.	
.	.	.	.	. .	.	
.	.	.	.	. .	.	
2	998	48.850	0.098	663- 679	I	
220	1218	59.618	10.768	680- 696	Ixxxxxxxxxxx	
0	1218	59.618	0.	697- 713	I	
213	1431	70.044	10.426	714- 730	Ixxxxxxxxxx	
0	1431	70.044	0.	731- 747	I	
612	2043	100.000	29.956	748- 764	Ixxxxxxxxxxxxxxxxxxxxxxxxxxxx	

2043 ENTRIES TOTAL AVERAGE = 408.11160 VARIANCE = 121836.060 STANDARD DEVIATION = 349.050

Figure 7-6. Device Space Utilization Report

reported on this histogram and not to the 0-08-01 histogram. Entries in the column headed CYLNDR NUMBER give the range of cylinders which form each histogram bucket. The number of cylinders in each bucket is a function of the device type. The entries in the column headed INDIV. NUMBER give the number of accesses made to that device within the physical space defined by the range of cylinders.

Similarly, the columns headed INDIV. PRC and CUMUL. PRC give the individual and cumulative percentages of all accesses to that device which were made within each cylinder range. The graphic portion of the display gives a visual indication of the percentage of accesses which were made for each range of the device space. This helps to quickly assess the access pattern of the usage for the device, i.e., whether the device is totally allocated and used or locally used. Figure 7-6 shows a device whose usage is split between two extremes. Such a situation should be investigated.

In the upper right hand corner of the report, a report number is indicated. This report number is used only to distinguish one histogram from another and in no way indicates the device to which the report refers. In addition, report numbers may not appear sequentially and this, in no way, is indicative of a problem. This report is always generated and cannot be turned off.

7.5.7 Device Seek Movement Report (File 42). The seek movement histogram is produced for devices in the mass storage subsystem being analyzed and provides the distribution of distance traveled by the read mechanism. Figure 7-7 is an example. The data used to generate this report is the absolute value of the difference between the cylinder addresses of each successive access to the given device. The column headed CYLNDR MOVED contains the range of seek movement distance for each line of the report. The column headed INDIV. NUMBER contains the counts of the number of accesses which caused the arm to be moved that distance. Figure 7-7 shows 854 accesses caused no arm movement (the same cylinder was successively accessed) for IOM-0 Device 05 on PUB 8. The INDV. PRC and CUMUL. PRC columns give the individual and cumulative percentages of the accesses to that device which resulted in a particular range of seek movement.

Figure 7-7 shows, for example, that 263 (12.8 percent) of the accesses caused arm movement in the range of 714 to 730 cylinders and that 1313 (64.2 percent) of the arm movements were a distance of 16 cylinders or less. The statistics at the bottom of this report give the characterization of the seeking characteristics of the files being accessed on this device during the monitoring period. The minimum movement, maximum movement, average, variance, and standard deviation are included. This report is always generated and cannot be turned off.

The Device Space Utilization Report (figure 7-6) shows that most of the accessing was concentrated in two areas on the device (cylinders 0-50 and

DISTRIBUTION COLLECTED ON SYSTEM NMCC2 AT 18:36:38 ON 81-03-13

SEEK MOVEMENT OF IOM-0,PUB-08,DEVICE-05-- MS0450 RF5

INDIV. NUMBER	CUMUL. NUMBER	CUMUL. PRC.	INDIV. PRC.	CYLNDR NUMBER	PERCENT OF OCCURRENCE 00 05 10 15 20 25 30 35 40 45 50	REPORT 50
					I....I....I....I....I....I....I....I....I....I....I	
854	854	41.801	41.801	0- 0	Ixxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx	
459	1313	64.268	22.467	1- 16	Ixxxxxxxxxxxxxxxxxxxx	
118	1431	70.044	5.776	17- 33	Ixxxxxx	
63	1494	73.128	3.084	34- 50	Ixxx	
17	1511	73.960	0.832	51- 67	Ix	
5	1516	74.205	0.245	68- 84	I	
7	1523	74.547	0.343	85- 101	I	
10	1533	75.037	0.489	102- 118	I	
1	1534	75.086	0.049	119- 135	I	
18	1552	75.967	0.881	136- 152	Ix	
.	.	.	.	. .	.	
.	.	.	.	. .	.	
.	.	.	.	. .	.	
10	1633	79.931	0.489	663- 679	I	
36	1669	81.694	1.762	680- 696	Ixx	
2	1671	81.791	0.098	697- 7.3	I	
263	1934	94.665	12.873	714- 730	Ixxxxxxxxxxxxx	
2	1936	94.763	0.098	731- 747	I	
107	2043	100.000	5.237	748- 764	Ixxxxx	

2043 ENTRIES TOTAL AVERAGE = 167.63534 VARIANCE = 85389.356 STANDARD DEVIATION = 292.215

Figure 7-7. Device Seek Movement Report

680-764). This is not necessarily bad, but if "interlaced" accessing was between these two areas, many seeks of about 700 cylinders would result. If, on the other hand, accessing to one area was completed before accessing to another began, there would be no large number of movement seeks. In this case, the example in figure 7-7 shows 18 percent of the seeks to be across a distance of 714-764 cylinders. This represents a situation where the particular files and jobs may need to be identified so that methods for reducing the number of movement seeks may be found and implemented for performance improvement. Further confirmation of this problem could be found by analyzing the Head Movement Efficiency Report (see subsection 7.5.8 and figure 7-8). If a problem exists then the connects/arm movement column for this particular device should approach a value of one. This figure indicates how many connects are issued between each movement of the arm. Therefore, even though we may have long seeks occurring on the device, if a large number of connects are being processed between these seeks, this would tend to lessen the impact of the long seeks.

For the Device Space Utilization Report and Device Seek Movement Report, an entry is made only for multi-command connects (see subsection 7.3) such as a seek/read or seek/write. If the first command of an IO connect is not a seek, or pre-seek, then an entry will not be made to this set of histograms. For this reason, the number of connects reported in these reports, for a given device, may be somewhat lower than that reported in the Proportionate Device Utilization Histogram, described in subsection 7.5.20. In addition to not recording non-multicommands such as controller commands and reset status commands, there are also a significant number of multicommands that are issued by the system for which the system does not generate a "valid" seek address in the "normal" manner. These connects will not be reported in the Space Utilization or Seek Movement histograms, but will be reported in the Proportionate Device Utilization Histogram. Research is currently underway to determine under what conditions these "nonvalid" seek address I/O requests are generated and whether there is any procedure by which the MSM could determine the correct seek address (see subsection 7.3). This software release has a code included to correct this problem. Therefore, no occurrences of this condition should be reported in any of the MSM output.

7.5.8 Head Movement Efficiency Report (File 42). This report displays how many connects are issued per arm movement of the device. Figure 7-8 is an example. The first three columns give the IOM, Channel, and Device number of the device. This is followed by the number of connects issued to that device and the number of times any arm movement was required (size of the seek is not considered). The final column indicates the ratio of connects to arm movements. The larger this ratio is, the more efficient is the device (i.e., the larger is the number of connects being handled between each arm movement of the device). Following the breakdown of arm movement by individual device, a summary is presented for arm movement within each individual mass storage subsystem. This is followed by three lines of output summarizing the overall efficiency of the entire disk subsystem. The first line presents the total number of connects issued, the second

MSM HEAD MOVEMENT EFFICIENCY REPORT FOR SYSTEM NMCC ON 81-12-07

IOM	PUB	DEVICE	CONNECTS	ARM MOVES	CONNECTS/ARM MOVEMENT
0	08	01	34848	28491	1.223
0	08	02	21479	8673	2.477
0	08	03	30030	18410	1.631
0	08	04	31128	19449	1.600
0	08	05	0	1	0.
0	08	06	220	1	220.000
0	12	01	6449	3138	2.055
0	12	02	2520	86	29.302
0	12	03	2048	549	3.730
0	12	04	3303	562	5.877
0	12	05	4818	1887	2.553
0	12	06	2929	754	3.885
0	12	07	220	1	220.000
0	12	08	4824	1479	3.262
0	12	09	9145	3303	2.769
0	12	10	3249	690	4.709
0	12	11	2285	754	3.031
0	16	01	0	1	0.
0	16	02	0	1	0.
0	16	03	12	1	12.000
0	16	04	3754	168	22.345
0	16	05	260	12	12.667
0	16	06	220	1	220.000
0	16	07	220	1	220.000

DEVICE TYPE	CONNECTS	ARM MOVES	CONNECTS/ARM MOVEMENT
.DS181	4466	185	24.141
.DS191	41790	13203	3.165
.DS450	117705	75025	1.569

TOTAL MASS STORAGE CONNECTS	163961
TOTAL ARM MOVEMENTS	88413
HEAD MOVEMENT EFFICIENCY	1.854

Figure 7-8. Head Movement Efficiency Report

line presents the total number of arm movements, irrespective of the number of cylinders traversed by each movement, and the final line presents the overall head movement efficiency (line 1 divided by line 2). The only connects reported in this report are those multicommand connects for which the MSM was able to determine a "valid" seek address (see subsection 7-3). By examining figure 7-8, the user will observe that the MSU450s appear to be least efficient, with the exception of device number 2. This report is always generated and cannot be turned off.

7.5.9 System File Use Summary Report (File 21). This report indicates where each system file is located and to what extent it was accessed across the measurement session. Only those files accessed are displayed. Figure 7-9 shows an example of this report. This report is produced by default but can be turned off with the input option (OFF) (subsection 7.6.9). The sum of accesses to all system files is then expressed as a percentage of all mass storage accesses. The system files are those files defined, via startup in the .CRDIT table. As can be seen in figure 7-9, the file names listed under the File Name column are not the actual file name, but rather relative file names indicative of their position description within the startup deck. Actual file names can be output in this report if the user selects the input option described in subsection 7.6.4. The connects reported are multicommand in nature and contain a "valid" seek-address as captured by the MSM. (See subsection 7.3).

This list of system files is then followed by a list of modules which reside in hard core because they are hard core modules or because they have been loaded into hard core by system personnel in order to save on I/O processing. If a system module is not loaded in hard core, and is required for some processing, then the system must perform an IO function to read this module from disk into a user's SSA work space. A significant amount of such system IO can cause severe system degradation. This degradation can be reduced by placing additional system modules in hard core or else by increasing the size of SSA Cache Memory. The System File Use Summary Report and the following Individual Module Activity Report should provide sufficient information to determine whether user action is required to reduce system IO overhead. If the percentage of system IO reported in this report is greater than 5-7%, then some user action is probably required. If addtional hard core space is available, the user should move as many system modules as possible into hard core. Each hard core module requires 1/2K (512 words) of memory, and there is 64K of hard core memory available. The System Configuration and Channel Usage Report (subsection 7.5.1) indicates the amount of hard core currently in use. The Individual Module Activity Report (subsection 7.5.10) can be used to indicate which system modules should be transferred to hard core. If sufficient hard core is not available, then the size of SSA cache should probably be increased. Once again, the Individual Module Activity Report should be referenced to see if this type of action would aid in reducing system IO.

MSM SYSTEM FILE USE SUMMARY REPORT FOR SYSTEM NMCC ON 81-12-07

FILE NUMBER	FILE NAME	IOM-PUB-DEV	STARTING SECTOR/CYLINDER	LENGTH (SECTOR)	ACCESSES
1	SYSTEM FILE1	0- 8- 1	107480/ 141	500	1
2	SYSTEM FILE2	0- 8- 2	40/ 0	5000	2
3	SYSTEM FILE3	0- 8- 2	6300/ 8	2500	62
4	SYSTEM FILE4	0- 8- 1	11040/ 14	4000	16318
5	SYSTEM FILE5	0- 8- 2	13300/ 17	10000	2569
6	SYSTEM FILE6	0-12- 7	40/ 0	4500	14
7	SYSTEM FILE7	0- 8- 1	15040/ 19	7600	421
9	SYSTEM FILE9	0-12- 7	9500/ 12	16000	854
10	SYSTEM FILE10	0-12- 7	25500/ 33	12000	491
				TOTAL	21678(13%)

.MBRT1	LOADED BY STARTUP	TYPE IS	HRD CORE MOD
.MCPIO	LOADED BY STARTUP	TYPE IS	HRD CORE MOD
.MDISP	LOADED BY STARTUP	TYPE IS	HRD CORE MOD
.MDNET	LOADED BY STARTUP	TYPE IS	HRD CORE MOD
.MDUMP	LOADED BY STARTUP	TYPE IS	HRD CORE MOD
.MFALT	LOADED BY STARTUP	TYPE IS	HRD CORE MOD
.MGEPR	LOADED BY STARTUP	TYPE IS	HRD CORE MOD
.MSCM1	LOADED BY STARTUP	TYPE IS	HRD CORE MOD
.MGP01	LOADED BY STARTUP	TYPE IS	HRD CORE MOD
.MGPIO	LOADED BY STARTUP	TYPE IS	HRD CORE MOD
.MIDSC	LOADED BY STARTUP	TYPE IS	HRD CORE MOD
.MIOS	LOADED BY STARTUP	TYPE IS	HRD CORE MOD
.MMTAP	LOADED BY STARTUP	TYPE IS	HRD CORE MOD
.MPRIO	LOADED BY STARTUP	TYPE IS	HRD CORE MOD
.MROUT	LOADED BY STARTUP	TYPE IS	HRD CORE MOD
.MSYOT	LOADED BY STARTUP	TYPE IS	HRD CORE MOD
.MTYPE	LOADED BY STARTUP	TYPE IS	HRD CORE MOD
.MDSX1	LOADED BY STARTUP	TYPE IS	HRD CORE MOD
.MDSX5	LOADED BY STARTUP	TYPE IS	HRD CORE MOD
.MDSX6	LOADED BY STARTUP	TYPE IS	HRD CORE MOD
.MFSIO	LOADED BY STARTUP	TYPE IS	HRD CORE MOD
.MSECR	LOADED BY STARTUP	TYPE IS	HRD CORE MOD

Figure 7-9. System File Use Summary Report

7.5.10 Individual Module Activity Report (File 21). This report shows the accessing done to each system module (figure 7-10). The report presents the system file the module resides in, followed by the module name and type. The module location, access count, and percentage of system file usage is then given. The last two entries give the total number of SSA CACHE buffer hits and disk loads this module accumulated (these values are 0 if SSA CACHE is not active). A minus one in column indicates a nonstandard SSA module. The Number of Accesses column reports the number of connects made to this SSA module as determined by the issuance of a trace type 7. The Disk Load column reports the number of times the SSA CACHE logic claims to have issued a connect to this SSA module. In those cases where no lost data occurred during the monitoring session and a GMC termination record was generated (i.e., GMC terminated correctly), these two columns should display equal values. Figure 7-10 shows this to be the case for almost all of the SSA modules. However, there are some exceptions. Module .MALC6 shows 109 accesses but only 104 disk loads (i.e., a difference of five). This apparent inconsistency has been reported to Honeywell, and an explanation requested.

If lost data occurred during the monitoring session (i.e., trace type 7 data lost), the Number of Accesses column could be significantly lower than the Disk Load column. If, on the other hand, GMC did not abort cleanly and a termination record was not generated, the Disk Load column could be significantly lower than the Number of Accesses column.

Finally, it should be noted that data from the last two columns pertains only to "STANDARD SSA" modules (see the Type column). Modules that are typed as "ABSOLUTE" or "EXCEPT PROC" are not placed into SSA Cache Core and therefore do not generate values for the last two columns. This report is produced by default but can be turned off with the input option (OFF) (subsection 7.6.9).

When generating this report, the data reduction program creates a temporary file (file code 55) which is used to produce a Job SSA Module Usage Report (subsection 7.5.11). If this report is desired, the report must be requested via input option (MODULE) (subsection 7.6.6).

At the bottom of this report, a summary line is produced indicating the percentage of buffer hits and disk loads. If the percentage of buffer hits is less than 90, then the size of SSA cache should be increased. For each 1K increase, 2 additional modules will be loaded into the SSA Cache memory.

When using this report to determine which additional SSA modules should be placed into GCOS Hard Core, the user should reference the "% of Activity" column. Those modules with the largest reported figure would be candidates for movement. In figure 7-10, .MALC6, .MALC9, .MFS03, .MFS04 would be candidates for movement.

MSM INDIVIDUAL MODULE ACTIVITY REPORT FOR NMCC ON 81-12-07

SYSTEM FILE	MODULE NAME	TYPE	IOM-PUB-DEVICE	SECTOR IN FILE	NUMBER ACCESSES	% OF ACTIVITY	BUFFER HITS	DISK LOADS
SYSTEM FILE4	.MAC02	STANDARD SSA	0- 8- 1	180	44	1	42	44
SYSTEM FILE4	.MALC2	STANDARD SSA	0- 8- 1	452	11	0	23	11
SYSTEM FILE4	.MALC5	STANDARD SSA	0- 8- 1	473	21	0	18	21
SYSTEM FILE4	.MALC6	STANDARD SSA	0- 8- 1	482	109	2	77	104
SYSTEM FILE4	.MALC9	STANDARD SSA	0- 8- 1	501	101	2	83	95
SYSTEM FILE4	.MBRT2	STANDARD SSA	0- 8- 1	670	12	0	3	12
SYSTEM FILE4	.MBRT3	STANDARD SSA	0- 8- 1	679	3	0	0	0
SYSTEM FILE4	.MBRT5	ABSOLUTE PRG	0- 8- 1	697	2	0	0	0
SYSTEM FILE4	.MBRT6	ABSOLUTE PRG	0- 8- 1	729	26	0	0	0
SYSTEM FILE4	.MCAL1	STANDARD SSA	0- 8- 1	831	20	0	858	20
SYSTEM FILE4	.MCAL2	STANDARD SSA	0- 8- 1	840	1	0	0	1
SYSTEM FILE4	.MFLT1	STANDARD SSA	0- 8- 1	1332	78	1	1402	77
SYSTEM FILE4	.MFS03	STANDARD SSA	0- 8- 1	1361	101	2	182	93
SYSTEM FILE4	.MFS04	STANDARD SSA	0- 8- 1	1370	99	2	30	93
SYSTEM FILE4	.MFS06	STANDARD SSA	0- 8- 1	1388	19	0	4	19
SYSTEM FILE4	.MFS08	STANDARD SSA	0- 8- 1	1406	18	0	37	18
SYSTEM FILE4	.MFS09	STANDARD SSA	0- 8- 1	1415	57	1	526	52
.								
.								
.								
.								
			TOTAL				76357	33455
			% HITS, LOADS/HITS + LOADS				69.53	30.47

Figure 7-10. Individual Module Activity Report

7.5.11 SSA Module Usage Report by Job (File 21). When requested by the user, this report will produced a listing for every job run during the monitoring period, showing all SSA modules referenced by that job and the number of such references. An example is shown in figure 7-11. This is the best method to use when determining which SSA modules should be softloaded into TSS core. This also provides an excellent means for studying the usage of SSA modules in general. This report is off by default and must be requested with a user input option (MODULE) (subsection 7.6.6).

7.5.12 File Code Summary Report (File 23) (NAME=FILECODE). The File Code Summary Report lists, by each activity, the files allocated to mass storage, their location and size, and the number of accesses made to each in the system during the monitoring period. Figure 7-12 is an example of this report. The activities in this report are in the same order as they appear in the Activity Summary Report (see subsection 7.5.15).

Each activity is identified by its SNUMB, activity number, and $ IDENT and USERID cards. There are as many data lines as necessary to describe each mass storage file used by the activity and the number of times the file was accessed. There is one line per file, and the file is described by its two-character file code, the device on which it was allocated (ALLOCATED DEVICE), its origin on that device (FILE ORIGIN) in units of LLINKS (320 words) and cylinders relative to the beginning of the device, and the size of the file (FILE SIZE) in LLINKS and cylinders. The column headed CONNECTS gives the count of the number of accesses made to the file.

There are several special file codes that will appear in this report. The following file codes will appear for almost every activity that is processed:

00 - Mass storage accesses made without a normal PAT entry; e.g., accesses made by the operating system as part of job intialization which are done without a PAT for efficiency

-- - All accesses to SYSOUT

=9 - File grow, PAT refresh, permanent allocation

=1 - Temporary allocation

0% - Load, pop, or push of an SSA

*, - FILSYS catalog search connect

It should be realized that these file codes are not used for a unique file request, and therefore may actually reference several different files. Therefore, in figure 7-12 it can be seen that activity 1 for SNUMB 2802T

DATA COLLECTED ON 81-12-25 AT 18:06

SNUMB	MODULE	# ACCESSES	MODULE	# ACCESSES	MODULE	# ACCESSES	MODULE	# ACCESSES
3625T	.MCAL1	7	DIRXLT	1	DDLXLT	1	.MFLT1	15
3625T	.MFLT2	11	.MMORE	5	.MAC10	7	.MFS09	17
3625T	.MFS03	8	.MFS36	4	.MFS15	4	.MFS61	8
3625T	.MFS73	10	.MFS37	4	.MRELS	4	.MREL1	3
3625T	.MFS04	24	.MFS13	2	.MSY11	11	.MSYT2	2
3625T	.MSYT3	2	.MIOS1	3	.MBRT2	3	.MFS27	2
3625T	.MGP23	2	.MBRT6	2	.MALC9	3	DBUG	1
3625T	.MMOR1	1	.MFS19	1	.MFS34	1	.MALC5	1
3625T	.MFS49	1	.MALC7	1	.MFS23	1	.MFS18	1
3625T	.MSNP1	2						
3637T	.MCAL1	6	ECFR	1	.MFS09	4	.MFLT1	2
3637T	.MFLT2	1	.MFS08	6	.MFS68	2	.MSYT2	5
3637T	.MSYT3	2	.MIOS1	1	.MBRT2	1	.MSYT1	4
3637T	.MFS27	1	.MGP23	1	.MBRT6	1	.MALO9	2

Figure 7-11. SSA Module Usage Report By Job

FILE CODE SUMMARY REPORT FOR SYSTEM NMCC ON 81-12-01

SNUMB,ACTVY #,IDENT,USERID 2802T- 1 1820513/10/4804/FORDI,KONKEL 674IDP00

FILE CODE	CONNECTS	FILE SIZE (LLINK/CYLINDER)	ALLOCATED DEVICE	FILE ORIGIN (LLINK/CYLINDER)
00	479	0/ 0	1-10- 2	0/ 0
*1 T S N F	76	48/ 1	0-12- 1	44361/291
B* T S N F	145	72/ 1	0-12- 9	26085/171
-- P R N F	53	0/ 0	0-12- 7	0/ 0
*Z T R N F	83	60/ 1	0-12- 4	48778/320
*C P S C F	44	139/ 1	1- 8-13	23523/154
S* T S N F	157	72/ 1	0-12- 4	48263/317
K* T S N F	36	600/ 4	2- 9-10	50649/333
D* P S N F	8	1/ 1	1- 8- 9	46522/306
*5 T S N F	38	36/ 1	0-12-10	29611/194
R* P S C F	61	65/ 1	0-13- 8	23364/153
*4 T S N F	2	12/ 1	0-12-11	48662/320
*J T R N F	1	0/ 0	1-10- 9	46518/306

SNUMB,ACTVY #,IDENT,USERID 2760T- 3 1820615/30/4955/GFEH* 220IDP11

FILE CODE	CONNECTS	FILE SIZE (LLINK/CYLINDER)	ALLOCATED DEVICE	FILE ORIGIN (LLINK/CYLINDER)
20 T R N F	1	420/ 3	0- 8- 4	51921/341
00	81	0/ 0	1- 8- 2	0/ 0
02 T R N F	384	180/ 2	0- 8- 8	61573/405
05 P S C F	3	14/ 1	0- 8- 7	25274/166
03 T S N F	20	1200/ 8	0- 8- 7	49663/326
01 T R N F	36	240/ 2	0- 8- 7	47671/313
-- P R N F	44	0/ 0	0- 8-12	0/ 0
10 T R N F	290	240/ 2	1-10-10	22349/147
04 P S N F	1	1/ 1	0- 8- 7	49383/324
12 T R N F	2	240/ 2	0- 8- 4	51681/340
09 T S N F	2	1200/ 8	0- 8- 8	46470/305
18 T S N F	2	60/ 1	1- 8-13	49109/323
15 T S N F	2	60/ 1	0-10-10	25549/168
19 T R N F	1	300/ 2	0- 8- 2	48910/321
*J T R N F	1	0/ 0	1- 8- 7	49378/324

Figure 7-12. File Code Summary Report

made 479 connects to file code 00 located on device 1-10-2. In actuality, the activity may have connected to several different devices, each time referencing this special file code. Instead of reporting these connects as separate entries, all the requests are grouped, and the allocated device that is reported is the the one that was referenced by the first of the 479 connects. There is a user option available (AREA) that expands the usage definition of these file codes. This option will then display all files using these file codes rather than grouping all of the unique file codes into a single collection file code.

The file code report for the TSS subsystem will also contain several special file codes. Following is a list of these:

TU - user files requested under TSS subsystems.
<< - file code referenced by TSS that contains a character that cannot be printed. The << symbol is being used instead.
JO - TSS I/O to this file code is JOUT processing.

This report will record all accesses to a file that is made with a multicommand request. No check is made to determine if a valid seek address has been recorded. (See subsection 7.3).

This report may produce excessive output and therefore is not produced under default conditions. The user must explicitly request this report with the use of an input option (ON) (see subsection 7.6.10).

Each file code entry has information which indicates the type of file by a "P" for permanent and "T" for temporary. This entry is located on the right side of the file code characters. Immediately to the right of the file type character is the access characteristics of the file. This is denoted by an "R" if random, or an "S" for a sequential file. The next character defines a permanent file as cataloged, "C", or noncataloged, "N". This character for a temporary file will be a blank as it is neither cataloged nor noncataloged. GCOS files are usually defined as noncataloged, permanent files and are treated at startup time much like permanent files by the file system. They are not, however, given the allocation treatment during normal operation time that is given to permanent files. The final character is an "F" for fixed or an "R" for removable.

7.5.13 Cat/File String Report (File 23). Immediately following the File Code Summary Report, the user will find the CAT/File String Report (figure 7-13). This report will provide the CAT/File String for every permanent user file referenced during execution. It will not list the CAT/File String for any special system files. In addition, it will indicate the total number of connects required to locate the file (catalog searching) every time the system was required to search for that file via a MME GEFSYE. Finally, it will indicate the number of connects required to

CAT/FILE STRING REPORT

SNUMB-ACT	FILE CODE	CONNECTS	CONN/GFYS	CAT/FILE STRING
8174T- 0	H*	70	10.000	A16IDPXO/HSTAR/H.DMS/
8174T- 0	H*	10	10.000	A16IDPXO/HSTAR/H.MSFMAC/
8174T- 0	H*	9	9.000	A16IDPXO/HSTAR/H.MSFMOD/
8174T- 0	PL	12	4.000	A16IDPXO/LINK/PLUSLIB/
8174T- 3	H*	20	10.000	A16IDPXO/HSTAR/H.DMS/
8174T- 4	H*	20	10.000	A16IDPXO/HSTAR/H.DMS/
8174T- 5	H*	20	10.500	A16IDPXO/HSTAR/H.DMS/
8174T- 5	PL	9	4.500	A16IDPXO/LINK/PLUSLIB/
8174T- 6	H*	20	10.000	A16IDPXO/HSTAR/H.DMS/
8174T- 7	H*	20	10.000	A16IDPXO/HSTAR/H.DMS/
8174T- 8	H*	20	10.000	A16IDPXO/HSTAR/H.DMS/
8174T- 9	H*	20	10.000	A16IDPXO/HSTAR/H.DMS/
8174T- 9	PL	9	4.500	A16IDPXO/LINK/PLUSLIB/
8174T-10	H*	20	10.000	A16IDPXO/HSTAR/H.MSFMAC/
8174T-10	PL	9	4.500	A16IDPXO/LINK/PLUSLIB/
8174T-11	H*	18	9.000	A16IDPXO/HSTAR/H.MSFMOD/
TEST1- 0	20	6	6.000	B29IDPXO/OBJECT/CAM/
TEST1- 0	30	6	6.000	B29IDPXO/OBJECT/GRT/
TEST1- 0	10	6	6.000	B29IDPXO/OBJECT/MUM/
TEST1- 2	20	13	6.500	B29IDPXO/OBJECT/CAM/
TEST1- 2	30	13	6.500	B29IDPXO/OBJECT/GRT/
TEST1- 2	10	13	6.500	B29IDPXO/OBJECT/MUM/
TSS - 0		13	13.000	A16IDPXO/DATA/TESTDATA/CKPTBAD/
TSS - 0		14	3.500	A16IDPXO/DATA/TESTDATA/DECELCMD/
TSS - 0		7	7.000	A16IDPXO/JOBS/CKPT/
TSS - 0		9	9.000	A16IDPXO/JOBS/DCAPS-EX
TSS - 0		8	8.000	A16IDPXO/RUNJCL/R.CANAL2/
TSS - 0		8	8.000	A16IDPXO/RUNJCL/R.FLASH/
TSS - 0		8	8.000	A16IDPXO/RUNJCL/R.STF/
TSS - 0		8	8.000	A16IDPXO/SYS-UTIL/KISSBK/
TSS - 0		7	7.000	A16IDPXO/SYS-UTIL/LISTJOBS
TSS - 0		7	3.000	A16IDPXO/SYS-UTIL/LIST-R./
TSS - 0		2	2 000	A16IDPXO/JOBS/DCAPS-EX/
TSS - 0		24	8.000	CMDLIB/LSIT/
TSS - 0		10	5.000	DATAMGT/DSTATF/
TSS - 0		15	5.000	DR8XI332XO/CMDLIB/LSIT/

Figure 7-13. Cat/File String Report

locate the file per MME GEFSYE. It will also provide the file code assigned to the file so that a cross reference can be made to the File Code Summary Report. This report is off by default and must be explicitly requested with the user input option (CAT) (see subsection 7.6.17).

In creating a structure, a user must be aware of the accesses that are required to actually obtain a file. Since the SMC is divided into 32 sections and each section is accessed via a hash, a user name can usually be found with one disk access. This is true if the total number of different user names is no more than 800 and all userids are evenly distributed over the 32 sections. It should also be noted that SMC contention (FMS attempting to access more than one SMC entry in the same section at the same time) can be significantly reduced by an even distribution of user names across the 32 sections.

The UMC is arranged so that access to the file description for one of the first four files or catalogs cataloged for a user requires one disk access, to one of the first 18 files requires one or two disk accesses, to one of the first 37 requires one, two or three accesses and so on.

If a large number of files are cataloged under the same user, allocation of the most of these files will be slow, since locating their file descriptions requires many disk accesses. For example, if there are 1000 files cataloged under a single user, allocating any one of the 500 files listed last requires between 28 and 54 disk accesses.

When it is necessary to include a large number of files under one user name, using subordinate catalogs each to include a fraction of the total can reduce the number of disk accesses.

Most the connects being displayed by this report are not attributed to the particular job in the File Code Summary Report. Rather the large portion of these connects are attributed to the Peripheral Allocator, file (*,), since it is the Peripheral Allocator that is responsible for issuing the MME GEFSYEs required to locate a given user file.

In figure 7-13, we see that job 8174T, activity 0, has three references to file code H*, where each H* is to a different Cat/File string. The reason for this is that during allocation (activity 0) the peripheral allocator must insure that all files referenced in the entire job are available to be used by this job. For job 8174T file H.DMS is asigned to H* in activities 3, 4, 5, 6, 7, 8 and 9, while file H.MSFMAC is assigned to H* in activity 10 and file H.MSFMOD is assigned to H* in activity 11. Therefore, since three different files are assigned to file code H* during the course of the job, activity 0 of that job will display three different references to file code H*. The reader will notice that file H.DMS is actually used in activities 3, 4, 5, 6, 7, 8 and 9, and that during activity 0 the peripheral allocator actually searched for this file 7 times for a total of

70 connects, or 10 connects per search. Then during the activity in which it was actually used, an additional two catalog searches (20 connects) were also required.

7.5.14 Connect Summary Report By Userid/SNUMB (File 23). If the user does not want to produce a complete File Code Summary Report, he may request that a report be produced for only certain USERIDs and/or SNUMBS. In this case, a much smaller File Code Summary Report (subsection 7.5.12) report will be produced. In addition, the user will receive a Connect Summary Report which will indicate, for the requested items, the number of times that USERID or SNUMB was referenced, the total number of connects made by that individual and the % of total connects represented by that item. If a requested SNUMB has a USERID equal to a requested USERID, then it will be reported twice in this report. See figure 7-14 for a sample of this report. This report is not produced by default and must be requested by a user input option (PROJ) (subsection 7.6.11).

7.5.15 Activity Summary Report (File 24). The Activity Summary Report lists each activity processed during the monitoring period and summarizes the activities as characterized by the eight variables: CP Time, Mass Storage Connects, Total Connects, and CP Time Per Connect (Mass Storage/Total) (see figure 7-15). The report lists the SNUMB-ACTIVITY number, the maximum number of I/O queues assigned to the activity, the maximum number of I/Os in process (transmission and/or queuing), the maximum number of intercom I/Os outstanding, the CP TIME (in milliseconds) charged by accounting to the job during the monitoring period, the number of connects issued to Mass Storage, the number of connects issued to Mass Storage and Tape, and the ratio of CP time over accesses for both Mass Storage Accesses and Mass Storage and Tape Accesses in the column headed CP TIME PER CONNECT. The bottom line of this report is titled TOTALS and gives the total charged processor time, the total connects, and the ratio of these totals.

System Schedular information for each activity (activity 0) is not reported separately. All activity 0 data is accumulated and reported as a single entry under the SNUMB $GENB. Any job whose SNUMB begins with a $ will not be considered as the System Scheduler, even if it processes as activity zero. In addition, NACE DERAIL jobs and BASHT jobs will also be processed as regular activities, even though they run under an activity number of zero. Finally, there is an input option (ZERO) that allows the user to define other zero activity jobs, that should not be considered as part of the System Scheduler.

Every activity that is processed is assigned buffer space so as to be able to initiate I/O. The standard default value for I/O queue buffer space is 5 I/O queues. However, methods are available for activities to acquire additional I/O queue space. The number of queues allocated to a given activity is outputted on this report. Following this column are two columns that provide an indication as to the amount of I/O being issued by a given activity. If the maximum number of I/O in process column displays

a number equal to the value in the I/O queue column, then this activity is probably being delayed because of a lack of sufficient I/O queue tables. If the maximum intercom outstanding column displays a value equal to one less than the value in the I/O queue column, then this activity is generating intercom I/Os at such a rate as to be exhausting the I/O queue table space. In this case also, the activity will be delayed.

Intercom I/O is the means by which two programs residing in the H6000 can pass data back and forth. TELNET and FTS use this technique extensively to pass and receive data from NCP. Whenever an activity exhausts its available I/O queue table allocation, a warning message will be printed on the status message report. This warning will indicate the time of day and the activity name that has exhausted its I/O queue resource. When the limit is no longer being reached, another warning message will be written indicating the activity is no longer at its resource limit.

The mass storage connects are displayed along with the total connects. For example, $PALC issued 906 mass storage connects and 1030 total connects. This represents 3.53 milliseconds of CPU time per mass store connect and only 3.10 milliseconds of CPU time between any connect type.

A line of asterisks are output when the monitor terminates in order to signify that the jobs that follow did not necessarily terminate. Information known about each job at the monitor termination is output. This report is on by default but may be turned off with a user input option (OFF) (subsection 7.6.9).

The report is useful in two applications. First, a quantitative feel for the CPU I/O balance of the system operation may be obtained from the TOTALS ratio of CP TIME PER CONNECT. Secondly, particular jobs which use excessive amounts of CPU or I/O time may be identified by SNUMB by scanning the list. More details on file usage of each activity in the Activity Summary Report are given on the File Code Summary Report. The $IDENT card of the job can also be found there for a more complete job identification.

THIS PAGE LEFT INTENTIONALLY BLANK

CONNECT SUMMARY REPORT BY USERID/SNUMB FOR SYSTEM NMCC2 ON 81-09-20

USERID/SNUMB	# OF ACTIVITIES	# OF CONNECTS	% OF TOTAL CONNECTS
DJ8XI70202	0	0	0.
DJ8XI70203	3	587	1.106
OPNSUTIL	1	3	0.006
DMSTA	1	2088	3.932
NEWRM	0	0	0.
8174T	11	7352	13.846
TSS	1	1891	3.561
$PALC	1	2224	4.188

Figure 7-14. Connect Summary Report By USERID/SNUMB

ACTIVITY SUMMARY REPORT FOR SYSTEM NMCC ON 82-06-25

JOB	MAX IO QUEUES	MAX IO IN PROCESS	MAX INTERCOM OUTSTANDING	CP TIME(MS)	CONNECTS MASS	CONNECTS TOTAL	CP TIME PER CONNECT	
$CALC- 0	20	9	1	5157	440/	440	11.72/	11.72
$PALC- 1	5	1	0	3195	906/	1030	3.53/	3.10
$SYOT- 0	17	6	0	0	512/	512	0. /	0.
$RTIN- 0	40	1	0	5141	3/	3	1713.67/	1713.67
TS1 - 1	105	7	0	51245	3281/	3281	15.62/	15.62
$LOGN- 0	50	1	0	1485	13/	13	114.23/	114.23
NCP - 1	87	9	5	76381	1114/	1114	68.56/	68.56
TLNT - 1	8	1	0	0	1/	1	0. /	0.
FTS - 1	183	20	18	83149	3151/	3151	26.39/	26.39
TLCF - 1	90	1	0	2	3/	3	0.67/	0.67
7088T- 1	5	2	0	2946	143/	151	20.60/	19.51
FSPUN- 2	8	1	0	0	2/	2	0. /	0.
MUM - 1	0	0	0	284	0/	47	0. /	6.04
SCHED- 0	0	0	0	10453	135/	149	77.43/	70.15
			TOTALS	239438	9704/	9897	24.67/	24.19

**** END OF REPORT SET NUMBER 1 1130-1141 *****

Figure 7-15. Activity Summary Report

This report will display all connects issued by a job with no regard to the type of I/O command or the validity of the seek address (see subsection 7.3).

7.5.16 Device Area File Code Reference Report (File 22). This report is generated to provide details on the jobs accessing a specific device area with their file codes. Figure 7-16 displays an example. The devices and areas to be listed are defined by the user when requesting input option Area (subsection 7.6.1). In figure 7-16, there are 10 areas requested for investigation. Each activity that accessed a device area is displayed in the report. At the end of the report, the number of connects found to each requested area is also given. This report is identical in format to the File Code Summary Report (subsection 7.5.12) except that this report contains only the file codes which referenced the specific area of the desired devices. The AREA N of each file code specifies within which area of the possible set of requested areas this particular file code fell. When this option is selected, the file code reference will automatically be expanded and special system file codes will be reported only if they actually referenced the requested area (see subsection 7.5.12). In addition, if the special system file codes referenced multiple areas, these file codes will appear multiple times within this report. In figure 7-16, it can be seen that activity 3 of job 52323 has multiple references to file code 0%. In the standard File Code Summary Report, all these references would have been grouped as a single reference, but in this report, they are expanded within each unique area requested by the user.

A complete explanation of the special file codes can be found in subsection 7.5.12. This report is not produced under default conditions and must be requested with a special user input option (AREA) (subsection 7.6.1).

7.5.17 Device File Use Summary Report (File 21). This report shows the device use by the accesses per file class (temporary or permanent). Figure 7-17 is an example of this report. Each of these classes of allocation is subdivided into sequential and random files and their corresponding percentage of the total file use is presented in the report. File 00 accesses are not included in this report. This report will reflect only multicommand connects, but will make no check as to the validity of the seek address (see subsection 7.3). The device numbers being reported under the "DEVICE" column are the unique set of device numbers generated by the MSMDRP (see subsection 7.5.5). This report is on by default but may be turned off with a user input option (OFF) (subsection 7.6.9).

7.5.18 Chronological Device Utilization Report (File 26). This report provides a chronological listing of the six most active disk devices, by device number and their probability of utilization (see figure 7-18). This report is so designed that any time quantum can be set in the report. By varying the time quantum parameter, the user may select integer values from 1 to n (where n is a positive value in seconds). A time quantum variation is requested with a user input option (TIMEQ) (subsection 7.6.14).

DEVICE AREA FILE CODE REFERENCE FOR SYSTEM NMCC ON 81-21-01

SNUMB,ACTVY #,IDENT,USERID 52323- 3 1820011/10/4887/ 498IDPOJ

FILE CODE	CONNECTS	FILE SIZE (LLINK/CYLINDER)	ALLOCATED DEVICE	FILE ORIGIN (LLINK/CYLINDER)	
0% P R N F	19	0/ 0	1-14-11	162/ 1	AREA 4
0% P R N F	8	0/ 0	0-12-11	1052/ 6	AREA 4

SNUMB,ACTVY #,IDENT,USERID 7862T- 1 1833333/10,WOANFF 498IDPOJ

FILE CODE	CONNECTS	FILE SIZE (LLINK/CYLINDER)	ALLOCATED DEVICE	FILE ORIGIN (LLINK/CYLINDER)	
0% P R N F	2	0/ 0	0-12-11	92/ 0	AREA 4
0% P R C F	706	16384/108	0-12-11	3446/ 22	AREA 8

SNUMB,ACTVY #,IDENT,USERID 52323- 4 1820011/10/4887 498IDPOJ

FILE CODE	CONNECTS	FILE SIZE (LLINK/CYLINDER)	ALLOCATED DEVICE	FILE ORIGIN (LLINK/CYLINDER)	
0% P R N F	10	0/ 0	0-12-11	1153/ 7	AREA 4

NUMBER OF CONNECTS PER AREA

1888	10	325	2766	844	1368	1368	706	0	4723

Figure 7-16. Device Area File Code Reference Report

MSM DEVICE FILE USE SUMMARY REPORT FOR SYSTEM NMCC ON 81-12-01

DEVICE	TEMPORARY ALLOCATION (32%) ACCESSES					PERMANENT ALLOCATION (67%) CATALOGED ACCESSES (27%)					NON-CATALOGED ACCESSES (40%)				
	SEQUENTIAL	- %	RANDOM	- %	TOTAL	SEQUENTIAL	- %	RANDOM	- %	TOTAL	SEQUENTIAL	- %	RANDOM	- %	TOTAL
1	389	5	6800	94	7189	19	2	618	97	637	0	0	21406	100	21406
2	192	3	5856	96	6048	80	100	0	0	80	0	0	5134	100	5134
3	0	0	0	0	0	0	0	0	0	0	0	0	230	100	230
4	78	44	96	55	174	0	0	274	100	274	0	0	567	100	567
5	8541	59	5842	40	14383	8	3	237	96	245	0	0	1026	100	1026
6	917	89	102	10	1019	20	5	333	94	353	0	0	199	100	199
7	0	0	5295	100	5295	1	10	9	90	10	0	0	447	100	447
8	0	0	6075	100	6075	25	32	53	67	78	0	0	224	100	224
9	0	0	851	100	851	28	1	1676	98	1704	0	0	231	100	231
10	960	86	156	13	1116	13	10	111	89	124	0	0	4215	100	4215

Figure 17. Device File Use Summary Report

CHRONOLOGICAL REPORT OF DEVICE UTILIZATION FOR SYSTEM OSCC2 ON 79-10-25

TIME	DEVID	PROB	DEVID	PROB	DEVID	PROB	DEVID	PROB	DEVID	PROB	DEVID	PROB
11:35:38:10	4	0.36	2	0.24	8	0.12	11	0.08	15	0.08	16	0.08
11:36:38:00	4	0.44	19	0.24	2	0.09	16	0.06	8	0.05	11	0.05
11:37:38:00	4	0.32	20	0.17	16	0.09	2	0.07	11	0.06	12	0.05
11:38:38:00	21	0.49	4	0.07	1	0.07	15	0.07	14	0.04	9	0.04
11:39:38:00	21	0.55	4	0.07	14	0.05	9	0.04	16	0.04	15	0.03
11:40:38:00	21	0.60	4	0.10	1	0.05	2	0.05	14	0.03	8	0.02
11:41:38:00	21	0.73	4	0.09	2	0.06	1	0.03	11	0.01	8	0.01
11:42:38:00	21	0.68	4	0.09	2	0.05	1	0.03	11	0.02	14	0.02
11:43:38:00	21	0.72	4	0.11	2	0.05	8	0.02	1	0.02	14	0.01
11:44:38:00	4	0.21	21	0.21	1	0.13	7	0.11	8	0.07	2	0.06
11:45:38:00	4	0.21	15	0.18	21	0.15	7	0.07	1	0.07	2	0.05
11:46:38:00	21	0.30	4	0.21	8	0.18	1	0.10	14	0.04	16	0.03
11:47:38:00	21	0.26	4	0.23	1	0.15	8	0.08	10	0.04	2	0.03
11:48:38:00	4	0.21	17	0.18	21	0.12	1	0.09	8	0.06	9	0.06

Figure 7-18. Chronological Device Utilization Report

The default time parameter is set for 60 seconds and prints the 6 most active disk devices over each minute of monitoring time. The first column shows the time, starting at the beginning and terminating at the ending time of the tape or timeframe. The remaining columns show six disk devices, by device number with their probability of utilization, consecutively, in descending order, relative to the activeness of all the disk devices. The utilization is the probability of that device being accessed over the time quantum. By examining figure 7-18, one can see that device 4 is the most active disk device with device 21 being a very close second. The device numbers being reported under the "DEVID" column are the unique set of device numbers generated by the MSMDRP (see subsection 7.5.5). This report is not produced by default and must be requested with an input option (ON) (subsection 7.6.10).

7.5.19 FMS Cache Report (File 21). If the system being monitored is configured with FMS cache, the report shown in figure 7-19 will always be generated and cannot be turned off. The data items are internal counters generated by GCOS in its own monitoring of FMS Cache operation. The report presents the number of times that we had a cache hit (CACHE HITS) and the number of times that we requested a catalog look-up and had to read the disk (READS). The ratio of HITS/HITS + READS is reported (HIT RATIO). This ratio should be about 55%, and if it is lower, attempts should be made to increase the size of the FMS cache. The number of FMS cache buffers currently allocated is also displayed on this report. Other columns indicate the number of times the cache area was written to (WRITES) or cleared (CACHE CLEARS). Clearing of cache is done for several reasons, and details will not be covered in this document. The remaining columns of this report indicate just some of the reasons why a cache hit was not obtained and really are not important for purposes of a performance analysis.

This figure will also report on the buffer efficiency of the in core space tables. In GCOS the user has the option of keeping the available space tables (AST) on disk instead of in hard core. Module MAS04 is responsible for buffering and processing AST's when this option is invoked. If the user selects this option called No-In-Memory-Available-Space-Tables (NIAST), GCOS will establish a buffer pool in hard core and will cycle the ASTs through this pool on an as-needed basis. Obviously, the size of this buffer pool will have a direct effect on whether or not the AST is already in memory when required, or must be read into memory from the disk device.

The NIAST option permits a site that has experienced difficulties with the 64K hard core fence limitation to reduce the HCM memory requirements. Use of the NIAST option increases system overhead since I/Os are required to obtain the AST from the device and write the updated AST back to disk. In a shared mass storage environment the NIAST option is required for obvious reasons.

Information on NIAST buffer utilization is kept by .MAS04 in words 0-3. These words are broken down in the following manner (see figure 7-19).

0	ALCBUF	THE NUMBER OF TIMES BUFFER ALLOCATION WAS ATTEMPTED
1	NOBUF	THE NUMBER OF TIMES ALL BUFFERS WERE BUSY
2	TBLINM	THE NUMBER OF TIMES THE AST TABLE WAS ALREADY IN MEMORY
3	WAITBL	THE NUMBER OF TIMES THE TABLE WAS IN MEMORY BUT BUSY

The relationship of ALCBUF to TBLINM provides an insight into how efficiently the ASTs are being processed and the amount of I/O overhead involved (Efficiency Ratio). The relationship between NOBUF and ALCBUF provides information on whether the number of buffers is sufficient to handle the AST processing (Buffer Sufficiency). The relationship between TBLINM and WAITBL may indicate that heavily used devices are creating delays in allocating disk space (Delay Ratio Factor).

These statistics show the contention for buffers (i.e., number attempts to obtain a buffer versus the number of times a buffer was unavailable), and thus indicate whether more or fewer buffers are needed (value of n on $ INFO card). This contention should be balanced against memory usage to determine the number of buffers required. The following procedure defines approximation of memory size for b buffers using NIAST. A subsequent paragraph contains the formula for approximating memory size for non-NIAST operations. The following procedure should be applied to the largest device type (in number of AUs) to determine the amount of memory required for buffers.

1.	Determine AUs per device.	n/a = AU
2.	Determine packing density per AU.	AU/600 = PD
3.	Make PD modulo 8 (MD).	(PD+7)/8 * 8 = MD
4.	Modulo 8 density plus minimum AST plus 1.	MD + 64 + 1 = AST
5.	Add TDT space.	AST * 1.5 = TS (Table Space)
6.	Add buffer control words.	TS + 5 = BS (Buffer Space)
7.	Determine space for b buffers.	BS * b = TBS (Total Buffer Space)
8.	Add maximum .MAS04 size.	TBS + 650 = HCM Required

Where: b - Number of buffers on $ INFO card for NIAST or default value

n - Number of llinks on largest NIAST device on system:

DSS167 - 7999	DSS190 - 47943	MSU0400 - 61864
DSS170 - 14399	DSS191 - 61864	MSU0450 - 61864
DSS180 - 14399	MSU0310 - 14399	MSU0450 - 123272
DSS181 - 14399		

a - AU size in llinks for each device.

The size calculation for NIAST/RMVBL (no TDT required) is the same except that the multiplier in step 8 is 1.0 and .MAS04 size is 450 (not 650) in step 8. For non-NIAST devices, the calculation is

n/((600 * a + 7)/8 * 8 + 74

NOTE: These calculations contain integer division in which remainder is ignored.

The following description of mass storage allocation indicates the relationship between the buffers in .MAS04 and allocation using the AST's on the mass storage devices. This is for a single system (not shared mass storage); for shared mass storage, the AST Read is always performed.

In order to allocate mass storage space on a NIAST device, the AST on the device must be updated. This is done as follows:

- If the AST is not in any of the buffers, the AST must be read from the device, updated, and written to the device. A copy of the AST remains in the buffer.
- If the AST is in a buffer and not busy, it is updated and written to the device; the Read is not required. A copy of the updated AST remains in the buffer.
- If the AST is in a buffer but is busy, .MAS04 must wait until the buffer is not busy. The AST is then updated and written to the device.
- If all buffers are busy with ASTs of devices other than the one requested (i.e., the requested AST is not in a buffer and all buffers are busy), the allocation module must wait until a buffer is not busy and then retry the access. When an unbusy buffer is found, the AST is read into it, updated, and written to the device.

From this, it can be seen that the more buffers in .MAS04, the less the chance of a buffer being unavailable and the more chance that a requested AST will be in a buffer. However, the buffers are memory-resident, and in the interest of conserving memory, a site should try to operate with few buffers as possible without adversely affecting allocation I/O.

The PERM/RMVBL mix of devices on the system is another factor that must be considered in determining the number of buffers for NIAST operations. The site should consider the following:

- Only permanent files are allocated on RMVBL devices; i.e., no temporary files are allocated on RMVBL devices.
- Since an AST is used only in allocating (creating) files, not in normal read/write access, the number of buffers required for efficient operation will generally be small for a system containing a high percentage of RMVBL devices.
- In general, if the number of RMVBL devices is small compared to the number of PERM devices, the ALL option should be used (this option is implicit in the shared mass storage environment). If the number of PERM devices is small compared to the number of RMVBL devices and operation is nonshared, the RMVBL option should be used.

It may appear that the ideal number of buffers to ensure availability of a buffer and an AST in memory would be one buffer per device. However, this uses an excessive amount of memory and is unnecessary since buffer usage for some buffers (particularly those associated with RMVBL devices) is very low. Note that multiple buffers per device cannot be used because of system gating. NIAST operation with one buffer per device uses more memory than non-NIAST operation because .MAS04 and the buffers are memory resident.

In most cases, use of the NIAST option requires less memory space for mass storage allocation because the ASTs are maintained on the individual devices. However, the mass storage allocation module and its buffers are memory resident, and for some small mass storage configurations, these may use more memory space than would the ASTs.

It is recommended that a site wanting to use the NIAST option, or getting this option implicitly in a shared mass storage environment, consider the following procedure:

- Initially use the system-supplied (default) value for n (one buffer for each four NIAST devices).
- Periodically evaluate the NIAST allocation performance statistics to determine the efficiency of the mass storage allocation function; i.e., allocation attempts vs. buffer busy and AST available in memory. This will indicate whether there is a need for more or fewer buffers.
- If adjustment is needed, either to reduce buffer contention or to reduce memory usage, change the value of n on the $ INFO NIAST card on the next startup.

7.5.20 Proportionate Device Utilization Report (File 42). This report shows the proportionate utilization of each device configured on the mass storage subsystem. Figure 7-20 is an example. This histogram identifies each unique device ID (device number zero is an MPC controller) and provides both a count of the number of accesses made to each device (under the column headed INDIV. NUMBER) as well as the percent of all accesses which were to each device (under the column headed INDIV. PRC). The histogram shows the proportionate utilization of each device (i.e., the percent of all accesses which went to each device) in a graphical form. The physical device that each "Device ID" of the histogram represents is shown in the Physical Device ID Correlation Table (see figure 7-5). This report is always generated and cannot be turned off. In this report the user is looking for a device or devices which have significantly more utilization than others in the system. This highly used device would then be a potential bottleneck.

It is desirable, but not always practical, to have equal utilization for each device. The user should be reminded that data in figure 7-20 is cumulative over the monitoring period. The actual accessing pattern could have been periodic with the following form: Many accesses to device 4

FMS CACHE REPORT FOR SYSTEM NMCC ON 81-12-01

CACHE HITS	WRITES	READS	NON CC READS	NON 320 READS	SSTAK SKIPS	CACHE CLEARS	HIT RATIO
1903	2382	3889	195	2	199	102	48.933

THE NUMBER OF FMS CACHE BUFFERS IS SET AT 8
THE MAX ALLOWED IS 32 AND A GOOD HIT RATIO IS 55%

NIAST BUFFER EFFICIENCY REPORT

NUMBER OF BUFFERS ALLOCATED	BUFFER ATTEMPTS	BUFFER BUSIES	AST TABLE IN MEMORY	AST TABLE IN MEMORY BUT WAS BUSY
12	18	0	8	0

EFFICIENCY RATIO - INDICATION OF IO OVERHEAD 44.444

RATIO OF BUFFER SUFFICIENCY 100.000

DELAY RATIO FACTOR - ARE HEAVILY USED DEVICES CAUSING AN ALLOCATION DELAY 0.

Figure 7-19. FMS Cache Report

DISTRIBUTION COLLECTED ON SYSTEM OSCC2 AT 13:49:04 ON 81-09-28

DEVICE UTILIZATION BY DEVICE NUMBER

REPORT 1

INDIV. NUMBER	CUMUL. NUMBER	CUMUL. PRC	INDIV. PRC	DEVICE ID		PERCENT OF OCCURRENCE 00 05 10 15 20 25 30 35 40 45 50
						I....I....I....I....I....I....I....I....I....I....I....I
12	12	0.000	0.000	0-	0	I
18947	18959	9.835	9.828	1-	1	Ixxxxxxxxxx
8206	27165	14.091	4.257	2-	2	Ixxxx
20269	47434	24.606	10.514	3-	3	Ixxxxxxxxxx
34033	81467	42.260	17.654	4-	4	Ixxxxxxxxxxxxxxxxx
238	81705	42.384	0.124	5-	5	I
25580	107285	55.653	13.269	6-	6	Ixxxxxxxxxxxx
.	.	.	.	.	.	.
.	.	.	.	.	.	.
.	.	.	.	.	.	.

192774 ENTRIES 0 OUT OF RANGE

Figure 7-20. Proportionate Device Utilization Report

followed by many accesses to device 3 followed by many accesses to device 2 followed by many accesses to device 1, etc. Each device could have been a bottleneck for a subperiod of the total monitoring period. This could also have been the case if the proportionate utilization of each device was equal. The Channel Monitor can be used to uncover this cyclic type of usage. In addition, the Chronological Device Utilization Report (see subsection 7.5.18) was designed to uncover this type of problem by breaking down device utilization over time, rather than by utilizing a histogram. Nevertheless, when a single or small number of devices has a disproportionately large share of the accesses, they are potential bottlenecks and their usage should be further analyzed.

This report will show all connects that were issued to a given device. This includes all read/write connects, as well as any command type connects issued to a given device.(See subsection 7.3).

7.5.21 Elapsed Time Between Seeks Report (File 42). This is a histogram report for the frequency of occurrence of elapsed time intervals between the issuance of mass storage access connects. Figure 7-21 presents a sample. The elapsed time is calculated as the time difference between successive mass storage connects from the central system and thus is representative of the workload. It does not provide any meaningful information on the subsystem service capabilities.

The data presented give the count (INDIV. NUMBER) and percentage (INDIV. PRC) of elapsed time between accesses which fell within each time range. The column headed TIME MSECS gives the time range in milliseconds. Thus, the data of the row with a time of 18 gives the count and fraction of elapsed time intervals in the range of 17+ to 18 milliseconds. The columns headed CUMUL. NUMBER and CUMUL. PRC. give the accumulated counts and percentage and are useful in describing the mass storage rates, e.g., 75.4 percent of the accesses occur less than 21 ms after the last access.

The bottom of the report provides a statistical summary of the data in the report. Statistics given include average, variance, and standard deviation. These statistics apply to all data points that were measured. The statistics concerning OUT OF RANGE are for those data points which fall outside the range of the histogram. OUT OF RANGE points are included in the previous statistics. This report is always generated and cannot be turned off.

7.5.22 Data Transfer Size Report (File 42). A sample histogram report on the frequency of occurrence of sizes of the data blocks transferred between mass storage and main memory is given in figure 7-22. Refer to subsection 7.5.12 for a description of the histogram format. This report has increments of 64 words, and the number in the column headed NUMBER WORDS is the upper value. The occurrence of certain data transfer sizes should be anticipated. For example, 64-word blocks are used for catalog accessing; in other parts of GCOS, standard system format is 320 words. SSA modules are usually slightly less than 512 words. When the Timesharing Subsystem

DISTRIBUTION COLLECTED ON SYSTEM NMCC2 AT 18:36:38 ON 81-03-13

ELAPSED TIME BETWEEN SEEKS

INDIV. NUMBER	CUMUL. NUMBER	CUMUL. PRC	INDIV. PRC	TIME MSECS		PERCENT OF OCCURRENCE 00 05 10 15 20 25 30 35 40 45 50	REPORT 2
						I....I....I....I....I....I....I....I....I....I....I	
0	0	0.	0.	0-	0	I	
5935	5935	2.066	2.066	1-	1	Ixx	
7879	13814	4.809	2.743	2-	2	Ixxx	
10424	24238	8.438	3.629	3-	3	Ixxxx	
13957	38195	13.297	4.859	4-	4	Ixxxxx	
13675	51870	18.058	4.761	5-	5	Ixxxxx	
14260	66130	23.023	4.965	6-	6	Ixxxxx	
14658	80788	28.126	5.103	7-	7	Ixxxxx	
13735	94523	32.908	4.782	8-	8	Ixxxxx	
13192	107715	37.500	4.593	9-	9	Ixxxxx	
12161	119876	41.734	4.234	10-	10	Ixxxx	
11550	131426	45.755	4.021	11-	11	Ixxxx	
10956	142382	49.569	3.814	12-	12	Ixxxx	
10396	152778	53.189	3.619	13-	13	Ixxxx	
9580	162358	56.524	3.335	14-	14	Ixxx	
9380	171738	59.789	3.266	15-	15	Ixxx	
8894	180632	62.886	3.094	16-	16	Ixxx	
8319	188951	65.782	2.896	17-	17	Ixxx	
7945	196896	68.548	2.766	18-	18	Ixxx	
7182	204078	71.048	2.500	19-	19	Ixxx	
6454	210532	73.295	2.247	20-	20	Ixx	
.	.	.	.	.	.	.	
.	.	.	.	.	.	.	
.	.	.	.	.	.	.	
567	278696	97.026	0.197	49-	49	I	

287238 ENTRIES TOTAL AVERAGE = 16.89080 VARIANCE = 35200.476 STANDARD DEVIATION = 187.618

Figure 7-21. Elapsed Time Between Seeks Report

DISTRIBUTION COLLECTED ON SYSTEM NMCC2 AT 18:36:38 ON 81-03-13

DATA TRANSFER SIZE

INDIV. NUMBER	CUMUL. NUMBER	CUMUL. PRC	INDIV. PRC	NUMBER WORDS	PERCENT OF OCCURRENCE 00 10 20 30 40 50 60 70 80 90 100	REPORT 3
					I....I....I....I....I....I....I....I....I....I....I	
10195	10195	3.493	3.493	0- 63	Ixx	
13634	23829	8.163	4.671	64- 127	Ixx	
1994	25823	8.846	0.683	128- 191	I	
2392	28215	9.666	0.819	192- 255	I	
15622	43837	15.018	5.352	256- 319	Ixxx	
173872	217709	74.583	59.565	320- 383	Ixxxxxxxxxxxxxxxxxxxxxxxxxxxxxx	
12159	229868	78.748	4.165	384- 447	Ixx	
15620	245488	84.009	5.351	448- 511	Ixxx	
10051	255539	87.543	3.443	512- 575	Ixx	
129	255668	87.587	0.044	576- 639	I	
8206	263874	90.398	2.811	640- 703	Ix	
80	263954	90.426	0.027	704- 767	I	
35	263989	90.438	0.012	768- 831	I	
.	.	.	.	. .	.	
.	.	.	.	. .	.	
.	.	.	.	. .	.	
29	283963	97.280	0.010	3136-3199	I	

291902 ENTRIES TOTAL AVERAGE = 710.98554 VARIANCE = 6275221.500 STANDARD DEVIATION = 2505.039

Figure 7-22. Data Transfer Size Report

swaps user subsystems, it will swap the entire subsystem under a single I/O request. Therefore, I/O transfer sizes of 1024 times the size of the subsystem will be recorded in this report. The TSS swap activity is individually recorded in the Data Transfer For TSS Swap Files histogram (see subsection 7.5.23). Special user application data base structures may use a different but observable block size, and this report will give an indication of the relative frequency of reference to that data base. It should be noted that this data is derived by the monitor from a scan of the DCW list at connect time. Any I/O which uses an "embedded DCW" list technique, or includes a transfer DCW in the list, may not have its size correctly recorded by the monitor. The size recorded will be less than the actual size for these special cases. This report is always generated and cannot be turned off.

7.5.23 Data Transfer Sizes For TSS Swap Files (File 42). The distribution length of the data transfer sizes to TSS swap files is displayed in figure 7-23. TSS has four different possible swap files (#S, #T, #U and #V), and activity to all four swap files are is recorded in this single histogram. In addition, if there are multiple copies of Timesharing, all of the various copies of Timesharing will record their swap activity to this histogram. The report has increments of 1024 words (1K) so that each individual bucket represents a TSS subsystem of 1K, 2K, 3K, etc. Therefore, this histogram can be used not only to determine the swapping pattern of TSS, but also the relative sizes of the various subsystems being used under Timesharing.

7.5.24 Connects Per Minute Report (File 20). This report (figure 7-24) provides a time-oriented summary of the number of connects issued by all jobs executing in the system, by the Timesharing Subsystem, and by specially defined user jobs. If multiple copies of Timesharing are present, I/O activity generated by all copies will be recorded under the single entry of TSl. It is not possible with the current program to obtain separate entries for each of the copies of Timesharing that are executing.

This report would be useful in a study of Time Sharing Response. If Time Sharing Response was known to be bad during a given time period, this report would provide some indication as to whether or not the rate of I/O was excessive during that period. This report is off by default and must be turned on with a user input option (ON) (subsection 7.6.10). The time quantum control has a default value of 10 minutes but can be varied by a user input option (RATECH) (subsection 7.6.16). If specially defined user jobs are to be reported, the user input option RATE must be invoked (subsection 7.6.18).

7.5.25 Special FTS File Access Time Report (File 42). This report (figure 7-24.1) provides a list of all file codes used by the WIN File Transfer System for the purpose of transferring user files. The report can be used to determine the rate at which the file transfer system is reading or writing a given file within the host system. It should be realized that a file cannot be transferred across the network any faster than FTS is able

to read that file within the host system. Similarly, FTS cannot write a file on the host system any faster than it is receiving data from the network. The actual catalog/file string being transferred can be determined by obtaining the CAT/File String Report (see subsection 7.5.13). The two reports can be correlated based on time and file code. It should be noted that all FTS Catalog File Strings are separated from the remaining Catalog File Strings and outputted on a separate page of the CAT File String Report. In the FTS File Access Report, the time of day is printed in HHMMSS.SS format. This report will automatically be produced and cannot be disabled. Furthermore, if any two consecutive reads or writes to a given file take more than 30 seconds, a warning message will be produced on the Special Processing Message page. This warning message states:

*** Special FTS Message - FTS took xx seconds between consecutive IOS to the same file. TOD was .HHMMSSSS E 06

There appears to be a problem with the current version of this report and research is being conducted to generate a solution. In the current software, it is possible for the "Size-LL" column to report that a file is 50LL long and the "Trans-LL" column to report that FTS transferred 75LL. It would appear to be impossible for FTS to read/write more data than currently indicated to exist on the file. A fix will be released immediately upon testing and verification. In a majority of cases, however, the data being reported is accurate.

7.5.26 TSS Swap File Usage Over Time Report (File 42). This report is obtained whenever the user selects the Connects Per Minute Report (see subsection 7.5.24). The time quantum of the report is controlled by manipulating the time quantum control of the Connect Per Minute Report. This report will indicate the rate of TSS swapping that is occurring and the swap files being used. It will aid the user in determining if TSS swapping is a cause of bad TSS response and if the user should add additional swap files. If a period of bad TSS response has been observed or reported, this report can be used to determine if the rate of TSS swapping is significantly higher during this time period. If a correlation if found, then a possible explanation of the bad TSS response would be the subsystem swapping being performed by TSS. Figure 7-24.2 shows a sample of this report. If multiple copies of Timesharing are present, all swap activity going to #S, #T, etc., will be reported under the individual file, irregardless of which copy of TSS is responsible for the I/O activity.

7.5.27 Device Seek Movement Summary Report (File 29). It is possible for the user to turn off all histograms (see subsection 7.6.19), but still receive a report summarizing the seek movement activity to each device. This report is merely the last-line information that would have been reported on the individual histograms, had the histograms been obtained. Figure 7-24.3 is a sample of this report. In order to obtain this report, the user should refer to the input option LIMITS in subsection 7.6.19.

7.5.28 Special Processing Messages. During the course of processing, several special processing messages may be generated by the MSMDRP. Most of these are for information purposes only and can be ignored by the analyst. Following is a list and brief explanation of the most common messages. These messages will most normally occur immediately in front of the System Configuration and Channel Usage Report.

- o MONITOR MUM WAS ACTIVE . . . .
  This message is produced for each monitor that was active during the monitoring session.

- o MSM DATA REDUCED FROM CHANNEL MONITOR . . .
  The MSM was not active during the monitoring session but the Channel Monitor was active. Sufficient information is collected so that all reports from the MSMDRP can be generated with the exception of the SSA cache portion of the Individual Monitor Activity Report (see subsection 7.5.10).

- o RUN BEING TERMINATED. DATA FOR MONITOR . . .
  Neither the MSM or CM were active and, therefore, no reports can be produced.

- o PROCESSOR # N IS (AVAILABLE/RELEASED) AT (TIME)
  This message will indicate the assignment or releasing of

- o SPECIAL FTS MESSAGE - FTS TOOK . . .
  The message is explained in subsection 7.5.25.

- o WARNING WARNING SYSTEM INTERCOM IO . . .
  Refer to subsection 7.5.1

- o WARNING WARNING I/O QUEUE TABLE . . .
  Refer to subsection 7.5.15

DISTRIBUTION COLLECTED ON SYSTEM NMCC2 AT 18:36:38 ON 81-03-13

DATA TRANSFER SIZES FOR TSS SWAP FILES

INDIV. NUMBER	CUMUL. NUMBER	CUMUL. PRC	INDIV. PRC	NUMBER WORDS	PERCENT OF OCCURRENCE 00 05 10 15 20 25 30 35 40 45 50
					I....I....I....I....I....I....I....I....I....I....I....I
2407	2407	25.406	25.406	64-1024	Ixxxxxxxxxxxxxxxxxxxxxxxx
3162	5569	58.782	33.376	1025-2048	Ixxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx
161	5730	60.481	1.699	2049-3072	Ixx
30	5760	60.798	0.317	3073-4096	I
.	.	.	.	.	.
.	.	.	.	.	.
.	.	.	.	.	.

REPORT 5

9474 ENTRIES TOTAL AVERAGE = 9065.15894 VARIANCE = 59641796.000 STANDARD DEVIATION = 12634.943

Figure 7-23. Data Transfer Sizes For TSS Swap Files Report

MSM CONNECTS PER 5 MINUTE REPORT FOR SYSTEM DNAH66 ON 81-12-17

TOD	2783T	TS1	TOTAL
10:10:00.00	4641	1160	17988
10:15:00.00	12168	1565	15590
10:20:00.00	11492	1126	25680
10:25:00.00	9766	1451	20420
10:30:04.50	5132	758	10550

Figure 7-24. Connect Per 5 Minute Report

SPECIAL FTS FILE ACCESS TIME REPORT

** THIS REPORT LIST ALL FILE CODES USED BY FTS TO TRANSFER USER FILES.

** IT ALSO INDICATES THE TOTAL AMOUNT OF TIME IT TAKES TO ACCOMPLIST THE READ/WRITE OF THE FILES.

** THIS INCLUDES NO NETWORK TRANSFER TIME ON READS.

SNUMB	START	STOP	SIZE-LL	TRANS-LL	RATE LL/MIN	FC	TIME-SEC	COMMAND
FTS	114019.30	114019.40	2	2	120.000	PB	1.00	READ (025)
FTS	114042.20	114043.80	4	4	148.814	PD	1.61	READ (025)
FTS	114658.60	114658.60	2	2	120.000	PD	1.00	READ (025)
FTS	114114.40	114738.40	54	54	8.439	PE	383.93	READ (025)
FTS	114701.10	114915.30	34	34	15.198	PF	134.23	READ (025)
FTS	114010.20	115017.10	144	144	14.236	PA	606.90	WRITE (031)

Figure 7-24.1. Special FTS File Access Time Report

TSS SWAP FILE USAGE OVER TIME REPORT (SIZE OF FILE WRITE IN WORDS) FOR SYSTEM NMCC2 ON 82-06-10

TOD	FILE #S		FILE #T		FILE #U		FILE #V	
	CONNECTS	AVG SIZE	CONNECTS	AVG SIZE	CONNECTS	AVG SIZE	CONNECTS	AVG SIZE
1303	406	7067.862	1476	8507.461	0	0.	0	0.
1313	311	8988.502	1192	9531.279	0	0.	0	0.
1323	174	8547 448	949	6986.984	0	0.	0	0.
1333	218	8774.789	1067	10412.626	0	0.	0	0.
1343	495	11218.626	1308	11293.547	0	0.	0	0.
1353	366	7144.787	1294	8085.091	0	0.	0	0.

7-44.2

Figure 7-24.2. TSS Swap File Usage Over Time Report

DEVICE SEEK MOVEMENT SUMMARY FOR SYSTEM NMCC2 ON 82-06-10

DEVICE LOCATION	DEVICE NAME	CONNECTS	AVERAGE SEEK LENGTH
SEEK MOVEMENT OF IOM-0,PUB-08,DEVICE-01-- MS0450	ST1	33559	56.20373
SEEK MOVEMENT OF IOM-0,PUB-08,DEVICE-02-- MS0450	DC2	8505	76.86820
SEEK MOVEMENT OF IOM-0,PUB-08,DEVICE-03-- MS0450	DC3	4630	68.88812
SEEK MOVEMENT OF IOM-0,PUB-08,DEVICE-04-- MS0450	DC4	5979	51.75682
SEEK MOVEMENT OF IOM-0,PUB-08,DEVICE-05-- MS0450	RF5	4213	37.23190
SEEK MOVEMENT OF IOM-0,PUB-08,DEVICE-07-- MS0450	RF7	278	0.00719

Figure 7-24.3. Device Seek Movement Summary Report

processors and provides the processor number involved as well as the time of day the assignment/releasing of processors occurred.

- o FOLLOWING PRINTS ARE THE INPUT OPTIONS . . .
  An echo print of all nonstandard input options selected will be produced. If any input option is described incorrectly, an error message will be generated indicating the type of error and the card number in error. The user should correct the error and resubmit the job for processing.

- o FOR INFORMATION ONLY
  This message will then be followed by several lines of output describing special record types that have been processed, or special processing events that have been executed by the MSMDRP. In most cases, the message can be ignored. Those messages which are important, and reveal an error in processing logic, will be described below. All other messages will not be described.

- o JULIAN DATES ARE BAD - RUN TERMINATED
  Every GMF data record is preceded by the current Julian date. The MSMDRP has found a Julian date that does not agree with the Julian date found on previous records. This can occur when an old GMF data tape is reused without degaussing. Old data is on the tape, and if the new data failed to write an end of file mark on the tape because of a system crash or malfunction, the MSMDRP, after reaching the end of the new data, would attempt to process the old data without realizing that it was old. The check on the Julian date prevents this from happening. The MSMDRP will terminate cleanly and all reports will be produced.

- o HAVE INCREASING OR BAD SEQUENCE NUMBERS . . .
  A problem has occurred in reading the data tape. If the run is reprocessed, the error may disappear. If it reoccurs, then the tape was generated with an error. In most cases, the MSMDRP will recover and processing will not be significantly affected. If it occurs often, contact CCTC/C751.

- o PROCESSING TERMINATED BY NXTRECRD . . .
  MSMDRP has requested the operator to mount a new tape and the operator responded that he did mount the new tape. However, MSMDRP is unable to match the inital record on the new tape with the last record on the old tape. User should check the data collection procedure to insure that correct tapes were mounted during the data collection phase. MSMDRP will terminate cleanly and all reports will be produced.

- o INCURRED A BAD SEEK ADDRESS . . .
  In the logic processing of subroutine SYSTMFIL, an unexpected condition has occurred. MSMDRP will continue processing correctly, but if this occurs frequently, CCTC/C751 would like to be notified. Call 202-695-0856.

7.6 Default Option Alteration

Most users rely upon the standard MSM Report formats and their default values as these suit a wide range of needs. A capability to change the reports is built into MSMDRP. The general form for all option requests are as follows: The first card contains an action code describing the action to be taken. Subsequent cards modify report parameters for some of the action codes. All input cards are free format with the only requirement being that at least one blank space separates multiple input parameters. The very last input card must have the word "END" entered in it. This card must be present whether or not any other input options are selected.

There is no specific order required of the options, and multiple entries of each are permissible. If several inputs refer to the same report, the last one encountered will have precedence. If a report is turned off by default and is modified, it will be turned on through the request for modification. The chart below shows the available actions: the mnemonic code for the user to identify the action; the function; and the default.

Mnemonic	Function	Default (indicated in parentheses)
AREA	Request file code references made to a specific area of a specific device	(not provided)
DEBUG	Debug	(no debug)
ERROR	Do not stop on Input Error	(stop)
FILDEF	Define system files by name	(no names used)
END	This card must be present	
MODULE	Produce the SSA Module Usage Report by Job	(no report produced)
NCONN	Process a limited number of connects	(total tape processed)
NREC	Process a limited number of tape records	(total tape processed)
OFF	Turn reports off	(all reports ON except reports 12,16, 18,20 - see table 7-1)
ON	Turn reports on	(all reports on except reports 12,16, 18,20 - see table 7-1)
PROJ	Produce the Connect Summary Report by Userid/SNUMB	(no report produced)

RN	This option must be selected when .MSMDRP is used to process WW6.4 data or when MSMDRP is executed on a WW6.4 system (process WW7.2/4JS data on a WW7.2/4JS system)
TIME	Set a timespan for measurement (no time criterion)
TIMEQ	Change time quantum for Chronological Device Utilization Report (report is off - default value is 60 seconds)
USERID	Suppress userids from reports (userids printed)
RATECH	Change time quantum for Connects Per 10 Minute Report (report is off - default value is 10 minutes)
CAT	Turn on the Cat/File String Report (report off)
RATE	Request the Connect Per 10 Minute Report for specific user jobs.
LIMITS	Limit the amount of processing performed and reports produced.
ZERO	Process jobs with an activity number of zero.

7.6.1 Monitor a Specific Device Area (Action Code AREA). This option allows a user to specify specific areas of a device for which all jobs referencing this area are to be highlighted. The format of the display is that of a File Code Summary and contains those jobs and file codes that reference the area of interest.

The device to be investigated is identified via the PUB and IOM number. The specific areas of interest are identified as beginning at the starting address defined in llinks. The length of the area is also in llinks, with a zero meaning the end of the device. A total of ten possible areas are allowed. The format for this card is shown in figure 7-25.

See subsections 7.5.12 and 7.5.16 for complete details on the report format generated with this user option. This report is off by default and will be activated by the processing of this action code.

7.6.2 System Debug (Action Code DEBUG). This is a restricted option for GMF system developers. DEBUG should only be used with guidance received by CCTC/C751.

7.6.3 Continue Data Reduction After an Input Option Error (Action Code ERROR). This option allows data reduction to continue when an error has

been detected and reported in an input option request. The default value reports the error and aborts the data reduction procedures. The format for this option is the word ERROR on the data card.

7.6.4 Specify System File Names (Action Code FILDEF). This option allows the user to specify the name of each system file displayed in the

Card 1 = A
Card 2 = N
Card 3 = B C D E F

Where

A = The word AREA
N = The number of areas to be specified. A maximum of ten areas are permitted. A Card #3 must be present for each area requested.
B = IOM number
C = Pub number
D = Device number
E = Starting address in llinks
F = Length of area in llinks

The following definitions apply to this option.

Device Type	Numbers Cylinders	Number Sectors/ Cylinder	Number Sectors/ Block (LLINK)
180	200	360	5
181	200	360	5
190	407	589	5
191	407	760	5
450	811	760	5

Figure 7-25. Specific Device Area Report Card Input

System File Use Summary Report discussed in subsection 7.5.9. This option is specified with a set of data cards. The first data card contains the word FILDEF. The second data card contains the number of system files to be described on the following cards. The following cards each contain a single pair of data points separated by at least one blank. The first data point is the system file number and the second data point is the desired system file name.

The standard output of the System File Use Summary Report is to label each system file as System File 1, System File 2, etc., corresponding to GCOS-HI-USE, GCOS-LO-USE, etc. In order to know the correct order of the file names, the user should check the $ FILES section of the startup deck. The order of the files in the $ FILES section of the startup deck is the order they are referenced in the report.

7.6.5 End Card (Action Code END). This card must be present at all times and must be the last data card supplied. It consists of the word END entered on the card.

7.6.6 Produce the SSA Module Usage Report by Job (Action Code MODULE). This option allows the user to produce the SSA Module Usage Report. This report will list every SSA module used by every job that was run during the monitoring session. See subsection 7.5.11 for details concerning this report. This report is off by default and cannot be turned on by using the ON option. This report can be activated only by entering MODULE on the data card.

7.6.7 Record Limitation by Connects (Action Code NCONN). This option allows a user to process only a specific number of connects. This option is especially useful if the tape contains an error on it and cannot be completely processed. Using this option, the user can process the tape or tapes up to the point where the tape error exists. This option requires two data cards. The first data card contains the word NCONN with the second card containing the number of connects to be processed.

7.6.8 Record Limitation by Records (Action Code NREC). This option allows a user to process only a specific number of tape records. This option is especially useful if the tape contains an error on it and cannot be completely processed. Using this option, the user can process the tape or tapes up to the point where the tape error exists. This option requires two data cards. The first data card contains the word NREC with the second card containing the number of tape records to be processed.

7.6.9 Turn a Report Off (Action Code OFF). This option allows a user to turn a report off that is on by default. In MSMDRP, all reports are on except report numbers 12, 16, 18 and 20 (see table 7-1). Only those reports in table 7-1 that have a name in parentheses () can be turned off with this option. Two data cards are required to use this option. The first card contains the word OFF and the second card contains the name of the report as displayed in the parentheses () in table 7-1.

7.6.10 Turn a Report On (Action Code ON). This option allows a user to turn a report on that is off by default. In MSMDRP, all reports are on except report numbers 12, 16, 18 and 20 (see table 7-1). Only those reports in table 7-1 that have a name in parentheses () can be turned on with this option (#9,10,12,15,16,17,18,20). Two data cards are required to use this option. The first card contains the word ON and the second card contains the name of the report as displayed in the parentheses () in table 7-1.

7.6.11 Produce Connect Summary Report by Userid/SNUMB (Action Code PROJ). This option allows the user to specify up to a total of 40 USERIDs and SNUMBs for which he wants the Connect Summary Report by Userid/SNUMB produced. The number of SNUMBs requested cannot exceed 10. In addition, the user can request the entire File Code Summary Report, or the user may want to see the File Code Summary Report only for a prespecified set of jobs or USERIDs, or the user may not want the File Code Summary Report at all. For example, the user can request 35 different USERIDs and 5 SNUMBs or 40 different USERIDs and 0 SNUMBs or 30 different USERIDs and 10 SNUMBs or 3 different USERIDs and 6 SNUMBs, etc. The format for this option is shown in figure 7-26. If values of zero are desired, they must be punched on the card. A blank is not equivalent to a zero. The Connect Summary Report will indicate for each requested USERID or SNUMB the number of connects made by the job or USERID. If a requested SNUMB also has a requested USERID, the number of connects issued by that job will be reported twice in the summary report. Refer to subsection 7.5.14 for a description of the report to be produced with this option. If the user desires to see a File Code Summary Report, it will be turned on via this option. The user does not need to use the "ON" input option.

7.6.12 Reduce WW6.4 Data or Process MSMDRP on a WW6.4 System (Action Code RN). This option requires two cards. The first card has the letters RN and the second card one of the following numbers:

1 - WW6.4/2H system processing WW6.4/2H data
2 - WW6.4/2H system processing WW7.2/4JS data
3 - WW7.2/4JS system processing WW6.4/2H data

The default is a WW7.2/4JS system with WW7.2/4JS data.

7.6.13 Set a Timespan of Measurement (Action Code TIME). The timespan of data collection can cover many hours of which only a few may be of interest. This option allows a user to specify the timespan (or spans) to displayed in all reports. For example, the user may specify that he wants to collect data from 0500 to 2200 and wants to display data only from 0900 to 1700 in all reports.

If the entire reduction will have a set timespan, the name "TOTAL" is used. Histogram reports cannot be individually timespanned. All timespans of "other" reports will be bounded by the overall report timespan, if one will be used. Up to five timespans for each report type may be specified,

Card 1 = A
Card 2 = B C D
Card 3+ = E
Card 4+ = F

Where

A = The word PROJ
B = 0 if Connect Summary Report is desired, but no File Code Summary Report is desired. The 0 must be punched on the card. A blank is not equivalent to a 0.
B = 1 if Connect Summary Report is desired and a complete File Code Summary Report is wanted
B = 2 if Connect Summary Report is desired and only a partial File Code Summary Report is wanted
C = Number of Userids (30 MAX)
D = Number of SNUMBS (10 MAX)
E = A total of C Userids with one Userid per card
F = A total of D SNUMBS separated by at least one blank space. All SNUMBS should fit on one card.

Figure 7-26. Limited File Code Summary Input Card Format

and they must be serially ordered. Names for other reports may be found in table 7-1. Only those reports in table 7-1 that have the NAME=specification can be time controlled. If a report does not have this specification, it cannot be time controlled. If the entire reduction is to be time spanned, the name that should be entered on the data card is TOTAL. The format for this action is given in figure 7-27. All times are expressed as four character fields with no intervening blanks. Time is based on a 24-hour clock. If the user wants to request the time 4:07 he must input a "0407".

If a start time but no stop time is desired, no characters should be entered after the minutes of the start time. If stop time is requested, there must be a start time corresponding to it.

The File Code Summary Report, the Activity Summary Report and the Device Area File Code Reference Report are all considered as a single unit under this option. Whenever a time frame is reached for any of these reports, all reports will be spanned. As an example, suppose that the user requested the following:

- File Code Summary Report for 1500-1600 and 1700-1800.
- Device Area File Code Reference Report for 1530-1730.
- Activity Summary Report for 1400-1530

For the above requests, the following report spans would be produced:

- At 1530, the File Code Summary Report would be produced for the period 1500-1530 and the Activity Summary Report would be produced for the period 1400-1530.
- At 1600, the File Code Summary Report would be produced for the period 1530-1600, and the Device Area File Code Reference Report would be produced for the period 1530-1600.
- At 1730, the File Code Summary Report would be produced for the period 1700-1730 and the Device Area File Code Reference Report would be produced for the period 1600-1730.
- Finally, at 1800, the File Code Summary report would be produced for the period 1730-1800.

7.6.14 Change the Time Quantum Value For the Chronological Device Utilization Report (Action Code TIMEQ). The user can change the time quantum value used to produce the Chronological Device Utilization Report by inputting the quantum in seconds. The default value is 60 seconds. Two cards are required. The first card contains the word TIMEQ and the second card contains the new quantum in seconds.

7.6.15 Suppress the USERIDs (Action Code USERID). The user can suppress the printing of USERIDs on the File Code Summary Report by entering the word USERID on a data card.

```
Card 1 A
Card 2 B C
Card 3 D E F G . . .
```

Where

A = The word TIME
B = Report name to be time spanned, or 'TOTAL'
C = Number of different times appearing on Card 3
D,E,F = Time used to define a timespan. Individual times must be separated from each other by at least one blank column. All times are considered to be on a 24-hour clock and must be expressed as a 4-digit field.

Figure 7-27. Input Option TIME Card Format

7.6.16 Change the Time Quantum Value for the Connect Per 10 Minute Report (Action Code RATECH). The user can change the time quantum value used to produce the Connect Per 10 Minute Report by inputting the quantum in seconds. Two cards are required. The first card contains the word RATECH and the second card contains the new quantum in minutes. The default value is 10 minutes.

7.6.17 Turn on the Cat/File String Report (Action Code CAT). This option, consisting of the word CAT on the data card, will turn on the Cat/File String Report (see subsection 7.5.13).

7.6.18 Request the Connect Per 10 Minute Report for Specific User Job (Action Code RATE). This option will allow the user to obtain the Connect Report for specific jobs as well as for the TimeSharing Subsystem and the Total System (see subsection 7.5.24). Card number 1 contains the word RATE, card number 2 the number of jobs desired (a maximum of 8 are permitted), and card number 3 the SNUMBs of the jobs desired. In addition to the requested jobs, the TimeSharing Subsystem as well as the Total System will also be reported. If multiple copies of TSS are in use, all activity will be reported under the single job name of TS1. If the user wants to obtain this report for only TimeSharing and the Total System, then he simply needs to use the "ON" input option using the name "RATE" for the required report ID.

7.6.19 Limit the Processing and Output (Action Code LIMITS). This option will allow the user to control the amount of output produced and the amount of record processing performed. Card number 1 contains the word LIMITS and card number 2 contains either the word ONLYSP or the word NOHIST or the word SUMARY. If the word ONLYSP is used then the Mass Store Monitor program will process only those data records that are generated by the SNUMBs requested under the RATE input option (see subsection 7.6.18). All other data will be ignored. The user must take care when examining the histograms and reports that are produced. The user must remember that only a limited amount of data has been processed. If the word NOHIST is used then no seek or space utilization histograms will be produced. This option can be used in conjunction with the ONLYSP option (must have two LIMITS input cards) or can be used by itself. In the latter case, all data will be analyzed, but no histograms will be produced. Finally, the user can request a summary of the seek movement activity. He can obtain this summary whether or not he selects to produce the set of histograms. To obtain the summary report, he must type SUMARY on a card immediately following the LIMITS card. A summary listing will not be produced for the space histograms, as this summary information is meaningless for this set of histograms.

7.7 JCL

The data reduction procedures consist of a single FORTRAN program having a main level and multiple subroutines.

A description of the more important JCL cards is presented below (see figure 7-28).

The $:LIMITS card should be studied to meet user needs. The run time (99) and output limit (30K) may both need to be altered as required by the duration of the monitoring run. The MSMDRP requires 55K of memory in order to execute plus an additional 2K for SSA space. During the initial loading process, MSMDRP will actually require 68K of memory, but 11K will be released immediately upon loading.

The statement:

```
$    DATA          I*
```

is used to identify the data cards that follow as described in subsection 7.6. At least one data card is required, that being an "END" request.

7.8 Multireel Processing

If more than a single reel of data has been collected, a series of messages will be outputted to the console informing the operator that a new data reel is required. The following are the messages produced.

a. DISMOUNT REEL #XXXXX THEN MOUNT REEL NUMBER YYYYY ON ZZZZZ

   In this case, XXXXX is the old reel number, YYYYY is the new reel number, and ZZZZZ is the tape drive ID.

   If the operator fails to mount the new tape, the above message will be repeated three times, after which the program will terminate, and all reports produced.

b. IS TAPE XXXXX MOUNTED ON DRIVE ID YYYYY (Y/N)

   In this case, XXXXX is the tape number being requested for mounting and YYYYY is the tape drive ID.

   This message occurs when the data reduction program finds the wrong tape has been mounted (by comparing internally generated tape labels). If the operator answers N, the message in (c) below is produced. If the operator answers Y, the data reduction program will terminate and all reports will be produced. In this case, the data reduction program is unable to process the tape. Even though the operator is mounting the correct tape, the internal label on the new tape does not match that being requested by the old tape. The user should check the data collection session to insure that the operator did not respond with an incorrect tape number during multireel change.

```
Col 1       8           16
$           IDENT       1820251/30/3044
$           SELECT      B29IDPX0/OBJECT/MSM
$           TAPE        01,X1D,,12345
$           LIMITS      99,55K,-4K,30K
$           DATA        I*
            Data cards - at least an "END" card must be present
$           ENDJOB
```

Figure 7-28. DRP MSM JCL

After entering the Y or N, the operator will need to hit the EOM key twice in order for the response to be transmitted.

c. WRONG REEL JUST MOUNTED, DISMOUNT AND MOUNT REEL XXXXX ON ZZZZZ

In this case, XXXXX is the new reel number, and ZZZZZ is the tape drive ID.

d. CAN TAPE XXXXX BE MOUNTED ON DRIVE YYYYY (Y/N)

In this case, XXXXX is the new desired reel and YYYYY is the tape drive ID.

If the operator fails to answer this message it will be repeated until he responds with a "Y" for YES or "N" for NO. If he types in "Y", then message (a) will be repeated. If the types in "N", then the program will be terminated and all reports will be produced.

7.9 Tape Error Aborts

During the course of processing it is possible that the operator will be required to abort the data reduction program due to an irrecoverable tape error. If such a condition occurs, the operator should abort the job with a "U" abort. This will allow the data reduction program to enter its wrap-up code processing and produce all reports generated prior to the tape error.

SECTION 8. CHANNEL MONITOR DATA REDUCTION PROGRAM (CMDRP)

8.1. Introduction

The Channel Monitor Data Reduction Program (CMDRP) is a FORTRAN program that sequentially processes data the Channel Monitor and Mass Store Monitor collected and wrote on tape. CMDRP produces a number of reports depicting the usage of the mass storage subsystem (i.e., connect time) and the conflicts being caused by multiple accesses to the subsystem. Report descriptions are presented in subsection 8.5.

When making I/O requests, system and user programs are required to provide system software with sufficient information in the request to complete the I/O transaction. In the case of mass store requests, a multiprogram environment will usually generate multiple requests for the same I/O device, causing the I/O system to create a queue for the waiting requests. Since an I/O channel is also a required resource to complete the I/O request, a channel queue may also develop and become a bottleneck in the I/O process. Knowledge of any type of queue in the I/O system is needed to determine if a bottleneck exists and if so, to determine where it is. The Channel Monitor has been developed to collect system events which audit the I/O request process through the I/O system. Associated with the data collection software is an offline reduction program to display the collected data in a meaningful form. Using this information, decisions can be made concerning mass store configuration of channels and devices. In addition, information is also reported concerning the queuing on tape channels. This set of reports, used in conjunction with those provided by the Mass Store Monitor Monitor Data Reduction Program (see section 7), will provide the user with a complete picture of mass store usage. It should be noted that when the Channel Monitor is active, sufficient data is collected to run not only the Channel Monitor Data Reduction Program, but also the Mass Store Monitor Data Reduction Program. However, unless the Mass Store Monitor was active, no reports will be produced on SSA cache performance, FMS catalog cache performance, or Cat/File string structure and accessing (see section 7).

There are two inputs to the CMDRP. The first is the data tape produced by the CM in the General Monitor Collector. The second input is a set of report option control cards used to alter the reports in some way other than the standard default. The various user input options and their formats are described in subsection 8.6. The actual reports produced by CMDRP are described in subsection 8.5.

8.2 Data Collection Methodology

The CM in the General Monitor Collector processes GCOS trace types 4, 7, and 22, which include the following:

- o Terminate I/O event

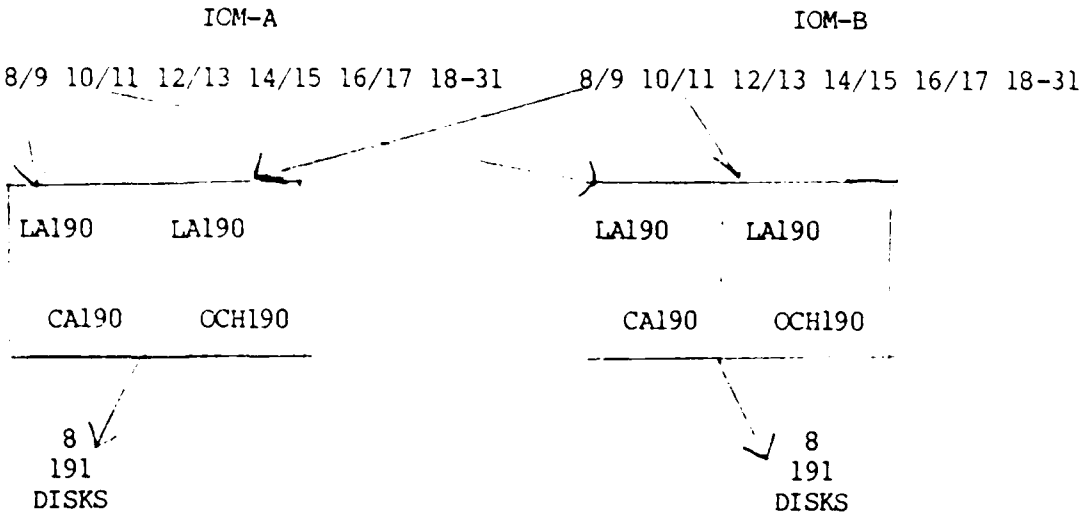


Figure 8-1. IOM Configuration

- o Connect I/O to IOM event
- o Link I/O to I/O chain event

The processing of a type 4 trace records all terminate I/O events generated by Input/Output Supervisor (IOS). The data generated by the processing of this trace event is sufficient to provide the information required by the reduction program for I/O termination processing. The format of the CM generated record for this trace event is described in subsection 5.4.6.1.

The processing of a type 7 trace records all connect events generated by IOS. Additional data are collected by the monitor to be used by both the Channel and Mass Store Reduction programs. These additional data for the Channel Monitor include the I/O queue address located in each monitored program's SSA. The IOM, channel, and device number are also collected for both monitors. The format of the generated record for this trace event is described in subsection 5.4.3.1.

The processing of a type 22 trace records all link I/O queue (i.e., I/O initiation) events generated by IOS. Additional data collected by the CM include the IOM, channel, and device number. The format of the generated record for this trace event is described in subsection 5.4.6.3.

8.3 Analytical Methodology

Three GCOS trace events are used to audit the I/O requests as they are processed through the I/O sequence. These trace types are type 4, type 7, and type 22 which denote the occurrence of a terminate interrupt, I/O connect, and start I/O, in the given order. The collection of these events in the time sequence in which they occur will simulate the I/O entries as they are queued to the various resources of the Mass Store Subsystem. This simulation of the I/O sequence is executed by the data reduction program.

In order to understand the reports produced by the CMDRP, the user should understand the I/O process as performed under GCOS. In the example shown in figure 8-1, channels 8 and 9 of IOM A and B are connected to one set of 8 drives. Channels 10 and 11 of IOM A and B are connected to another set of 8 drives. Each IOM can support up to 24 channels when describing a mass store configuration, and both physical and logical channels can be described. A physical channel is an actual channel over which data may travel. A logical channel is nothing more than a buffer area which is used to set up and store all information required for initiation of the actual I/O event. However, the actual transfer of data cannot be performed until a physical channel is available. However, the establishment of logical channels does allow significant processing to occur prior to the actual transfer of data. Both physical and logical channels require the

use of an IOM port. In figure 8-1, channels 8 and 10 on both IOMs are physical channels while 9 and 11 are logical. Every physical channel may have three other logical channels associated with it; i.e., a total of four channels.

A normal sequence of events produced by a request for I/O would be a type 22, type 7, and type 4 event, in that order. The trace type 22 indicates an I/O request has been issued. The order of the channel numbers, given on the crossbar card of the startup deck, determines the order in which GCOS searches for empty mailboxes and starts I/Os. The first problem area preventing effective I/O throughput will occur if two consecutive logical channel numbers on the same PSIA are sequentially in order on the crossbar card. GCOS, upon receiving an I/O request, will fill the mailboxes and start the first I/O. However, a second I/O cannot be started until the first I/O is complete (data can only travel over a physical channel, not a logical channel). The crossbar card should be designed so as to stagger PSIAs.

The second area preventing effective I/O throughput is the method in which the Link Adapters (LA) are utilized. Each PSIA channel in the MPC shares a LA with another PSIA channel. Each LA can service only one PSIA at a time. The PSIAs are paired as 0-1, 2-3, etc. The crossbar card should be designed so as to prevent two I/Os in succession from going to the same LA.

The following procedure should be followed to determine if the crossbar card is properly designed:

- o List in a vertical column the order of the IOM/channels on the crossbar card.
- o Analyze the MPC card of the startup deck and next to each IOM/channel written in step 1, write the MPC/PSIA associated with that channel.
- o A problem exists if the same MPC/PSIA numbers appear consecutively. This is an indication of consecutive logical channels (problem type 1 above).
- o An indication of the Link Adapter conflict is the occurrence of two consecutive entries in the MPC/PSIA column with the same MPC number and consecutive PSIA numbers (0-1,2-3).

The following is an example of the above procedure:

Cross Bar Cards

```
$ XBAR    IOM-0,PUB-10,IOM-1,PUB-10,
          IOM-2,PUB-10,IOM-2,PUB-11,
          IOM-1,PUB-11,IOM-0,PUB-11
```

```
$ XBAR    IOM-1,PUB-14,IOM-0,PUB-14
          IOM-2,PUB-14,IOM-2,PUB-13,
          IOM-0,PUB-13,IOM-1,PUB-13
```

MPC Cards

```
$ MPC-0   PSI-2,IOM-0,PUB-10,PSI-1,IOM-1,PUB-10,
          PSI-3,IOM-2,PUB-10,PSI-3,IOM-2,PUB-11,
          PSI-1,IOM-1,PUB-11,PSI-2,IOM-0,PUB-11

$ MPC-2   PSI-0,IOM-1,PUB-14,PSI-2,IOM-0,PUB-14
          PSI-1,IOM-2,PUB-14,PSI-1,IOM-2,PUB-13
          PSI-0,IOM-0,PUB-13,PSI-2,IOM-1,PUB-13
```

Chart

IOM/Channel	MPC/PSIA	IOM/Channel	MPC/PSIA
0-10	0-2	1-14	2-0
1-10	0-1	0-14	2-2
2-10	0-3 *	2-14	2-1
2-11	0-3	2-13	1-1 **
1-11	0-1	0-13	1-0
0-11	0-2	1-13	1-2

* problem 1
** problem 2

The problems described by the above procedures could be solved by redesigning the crossbar cards in the following manner:

```
$ XBAR    IOM-0,PUB-10,IOM-1,PUB-10,
          IOM-2,PUB-10,IOM-0,PUB-11,
          IOM-1,PUB-11,IOM-2,PUB-11

$ XBAR    IOM-1,PUB-14,IOM-0,PUB-14,
          IOM-2,PUB-14,IOM-1,PUB-13,
          IOM-0,PUB-13,IOM-2,PUB-13
```

If the I/O request cannot be granted, because either the channel or device being requested is currently busy, the request will be queued. This request will only be serviced when both a channel and device are free. When queuing occurs for a channel, GCOS will indicate the request queued over the primary channel. A primary channel is that channel which appears first on the $XBAR card, for a given string of devices. Therefore, all channel queue histograms are presented only for primary channels. However, a queue on the primary channel actually means that all channels, both physical and logical, connected to the desired device were busy. When the request is finally granted, a trace type 7 is issued.

A table is used to hold the device number, channel number, IOM number, I/O queue entry address, and the time the T22 trace event occurred. With the occurrence of each T22 event, the table entry is filled to mark the linking of the I/O requests. At this time, the required computations for determining the channel and device queue length are made. Channel queue histograms are produced for both tape and mass store devices, while device queue histograms are produced only for mass storage devices. The channel and device queue time also begins at this point and will be updated with the occurrence of the trace 7 event for this I/O request.

With the eventual occurrence of the trace 7 event for the I/O request, several updates are required to the common tables. The I/O queue time data are generated for the channel and device and collected for the appropriate histograms. It should be noted that it is possible for a device to show no queuing, but yet will display I/O queue time in its queue time histogram. The reason for this is that the queue time histogram is reporting the time difference between a T7 trace and a T22 trace. The T7 trace will not be issued until the actual I/O is initiated (i.e., a logical channel becomes available). Even though a logical channel might be available, several milliseconds might pass before a physical channel becomes available and the T7 trace is generated. Even if a physical channel is available, upon the occurrence of a T22 trace, several milliseconds might pass before the system generates the T7 event. This is especially true on a very busy system. A connect queue entry is now filled with data to be used for the I/O service time histograms. This connect queue holds the IOM, channel, device number, and the time of the trace 7 event. The channel and devices status table entries are also marked busy at this point. As a confidence test, the channel status is sensed at the start of processing for each trace 7 event for a nonbusy status. If it is busy, a lost interrupt is considered to have occurred since it is impossible for a connect to be issued to a busy device or channel. Device access histogram data and an IOM command execution count are also generated at trace 7 event time.

The next logical event for the I/O process is the termination interrupt originated by the IOM at the I/O data transfer completion. The signal for this event is transmitted by the IOM to the processor through the SCU as a request for the processor to service the I/O completion. The type 4 trace event contains the IOM number and channel for I/O termination. These data are used to determine the I/O service time by finding the time difference of the connect event and the terminate event. The time difference is collected and displayed in histogram form for each mass store and tape channel as well as for all mass store devices. The channel and device queue length are also adjusted at this point to reflect the absence of a queue being serviced for this channel and device.

It must be noted that exceptions to the normal I/O process are to be expected and must be accounted for in the reduction program. All the exceptions encountered so far have been diagnosed and coding in the program will allow for exceptions. Some of the exceptions include the following:

- o System programs that avoid the T22 trace by generating their own queue entry and by starting the connect immediately.

- o System programs that manipulate the I/O priority by linking themselves ahead of I/O requests already in the queue.

- o I/O requests for device zero (MPC) which makes a channel busy but not a device.

- o Lost interrupts from an I/O connect which leave the connect table and status table in an active state forever, if not detected.

- o Special controller commands which do not involve the channel or device.

8.4 Data Reduction Methodology

The CMDRP currently uses random I/O (File 58) to process histogram data. This feature allows the CMDRP to process an unlimited number of channels/devices with a minor increase in memory requirements. As delivered, the CMDRP will process 75 mass storage devices and 40 mass storage/tape channels. It will produce 106 unique histograms with a minimum amount of random I/O. If the number of channels or devices is insufficient, the user will need to edit file B29IDPX0/SOURCE/CM. The user should enter the edit subsystem and process the following command:

```
B RS:/NRDEVXX=XX75/;*:/NRDEVXX=XX New number of devices/

B RS:/NRCHANXX=XX40/;*:/NRCHANXX=XX New number of channels/
```

For each additional channel the size of the program will increase by 55 words and for each additional device the program will increase by 25 words. In the above edit the character "X" signifies a single space.

The last variable that will need to be changed is RPTCNT. This number represents the total number of histograms that will be processed with negligible random I/O. To calculate the total number of histograms that will be produced under your configuration, the following formula should be used.

> (number of mass store devices)*3 + (number of tape and mass store channels - both logical and physical) + (number of tape and mass store physical channels)*2.

If this value is less than 106 no change is required. If the value computed is greater than 106, the user may alter this value. This will help to decrease CPU/IO time but will increase storage by 80 words for each increment above 106. This tradeoff between CPU/IO time and memory must be

made at the discretion of the user. In order to change this value, the following edit function should be performed:

```
B RS:/RPTCNTX=106/;*:/RPTCNTX=New value
```

As in the earlier edit example, the character "X" should not be typed, but is being used to represent a blank column. After performing the above edits, the user should recompile the source program by entering the card subsystem and issuing a run command.

8.5 CMDRP Output

Reports generated from the Channel Monitor may vary from one collection period to another due to the difference in configuration of the hardware. Report numbers are preassigned to the histogram reports which are hardware independent and are dynamically assigned to histograms which denote the channel and device uniquely configured on each IOM.

The following subsections will describe all reports produced by the CMDRP and subsection 8.6 will describe the user input options.

8.5.1 System Configuration and Channel Usage Report (File 57). This report documents the system identification, configuration, and the date and time of the monitoring period, as well as reporting the usage of all configured I/O channels. Figure 8-2 is an example of this report. The heading line indicates the software version number that corresponds to this document. The version number should be 01-82. The first line after the heading provides the tape number(s) the report was generated from, the system identification, the date (in the form year, month, and day - YYMMDD), and the start and stop times (HH:MM:SS) of the MONITORING SESSION. The next several lines of output describe the overhead of all GMF monitors that were active during data collection. The monitor name is given, its CPU time in seconds, and its overhead as a function of total processor power. The GMF executive overhead is separated from the actual monitors and is listed as "EXEC". The monitor "NAME" is an area of code within the Mass Store Monitor and even though listed separately it is also included under the monitor "MSM". The monitor "FMS" is also an area of code within the Mass Store Monitor, but in this case it has not been included under the monitor "MSM".

Monitor "CM" in this report describes the processor overhead of subroutine T4 (terminate processing) and subroutine T22 (start I/O processing). Monitor "MSM" in this report describes the processor overhead of subroutine T7 (connect processing). Therefore, if the Channel Monitor was active, but the Mass Store Monitor was not, this report will still list both "CM" and "MSM" as contributing to the processor overhead. The total Channel Monitor overhead will be found by adding the overhead of the "CM" monitor to the overhead of the "MSM" monitor, to the overhead of the "FMS" monitor.

```
**************************************************************
* * *          C H A N N E L   M O N I T O R          * * *
* * *                VERSION  01-82                   * * *
**************************************************************

          TAPE # 24502
SYSID                    DATE       START TIME     STOP TIME

OSCC2 ,6.4.ID        82-01-08       13:49:04       14:31:40      FOR A TOTAL OF 0.71 HOURS

MONITOR OVERHEAD - MONITOR          TIME(SEC)      % OVERHEAD
                   EXEC               100            1.96
                   MUM                120            2.35
                   MSM                160            3.13
                   NAME                40             .78
                   FMS                 50             .98
                   TOTAL                             8.42

CONFIGURATION:  DUAL   PROCESSOR 6680,  DUAL  IOM,  512K  MEMORY - 52 OF WHICH WERE HCM
                THE HCM DOES NOT INCLUDE 3K FOR .CALC AND 6-8K FOR FILSYS
              CPUS ACTUALLY CONFIGURED = 2.00  CPUS ACTUALLY AVAILABLE = 1.75
                                IOM NUMBER 0

   CHANNEL          TYPE          CROSSBAR        CONNECTS
    0-08           .DS450           0-09           47750
                                    1-08
                                    1-09
    0-09           .DS450        SEE ABOVE         16851
    0-12           .DS191           0-13           55524
                                    1-12
                                    1-13
    0-13           .DS191        SEE ABOVE          2479
    0-16           .DS181           0-17            8290
                                    1-16
                                    1-17
    0-18           .DMTA9           NONE              62
    0-19           .DMTA9           NONE            9147

(IN THE ACTUAL REPORT A SUMMARY OF IOM NUMBER 1 WOULD FOLLOW)
```

Figure 8-2. System Configuration and Channel Usage Report

If both the Channel Monitor and Mass Store Monitor were active, then the combined overhead of both monitors can be found as the sum of "MSM" + "CM" + "FMS".

For purposes of this report, % overhead is computed as:

$$\frac{\text{(CPU TIME Used by Monitor)}}{\text{(TOTAL Elapsed Time)x(Number of Processors)}}$$

Following the overhead description are three lines of configuration information describing the number of processors, IOMs, and amount of memory configured to the system. In addition, the size of GCOS Hard Core, the size of the Core Allocator and the size of FILSYS is also presented. The third line of the configuration data indicates the number of processors actually configured and actually available. These numbers might be different than shown on the first line due to the assigning and releasing of processors. In figure 8-2, we see that one processor was released for a period of time (i.e., CPUs actually available is equal to 1.75). The actual time that processors were available or released is indicated in the status message printouts (see subsection 8.5.15).

The next portion of the report documents the channel configuration by IOM, listing each configured channel number, the device type configured to that channel, and the channel crossbarring. The crossbar column shows those channels that are crossbarred to the channel identified under the channel column. If SEE ABOVE is found, the crossbarring has been displayed on a preceding channel. The I-CC format of each channel description identifies the IOM and the channel number being discussed. The last column of this report displays the number of all connect types issued over that channel. This section will be repeated for each IOM configured to the system. Figure 8-2 only displays IOM 0 activity. This report is always generated and cannot be turned off.

8.5.2 System Summary Report (File 57). The System Configuration and Channel Usage Report and the System Summary Report may be used to assess overall system utilization. Figure 8-3 is an example of the System Summary Report. The first set of lines shows the number of connects to the monitored mass storage subsystems compared to the total connects issued (TAPE+DISK) and the connect rate per hour over the subsystem. Most systems will show a small number of Control Connects being generated by the MPCs configured to the system. These Control Connects will be summed together and listed as a separate subsystem line. Analysis on a Shared Mass Storage System shows the number of MPC connects generated to be a significant percentage of the total connects generated. The final part of this report is a list of the commands (octal code and mnemonic) issued to the mass storage subsystem and the count of each issued during the monitoring session. This report is always generated and cannot be turned off.

```
**************************************************************************
   * * *       S Y S T E M   S U M M A R Y   R E P O R T       * * *
**************************************************************************
```

TOTAL CONNECTS TO DSS181	8540 OF	207117 (4%)	AT A RATE OF	12024. PER HOUR
TOTAL CONNECTS TO DSS191	76949 OF	207117 (37%)	AT A RATE OF	108341. PER HOUR
TOTAL CONNECTS TO MSO450	107273 OF	207117 (51%)	AT A RATE OF	151036. PER HOUR
TOTAL CONNECTS TO MPC-CONTROL	12 OF	207117 (0%)	AT A RATE OF	16. PER HOUR

	COMMAND	COUNT
01	CONTRL	12
17	FMT TK	15485
25	READ	126899
26	RD CR	192
31	WRITE	35483
33	WR-VER	12862
40	RST ST	1841
TOTAL		192774

Figure 8-3. System Summary Report

A well performing system, under a heavy workload, should show a high utilization of the configured resources. Figure 8-3 shows that the I/O activity is predominantly on the MSU450 subsystem configured on channels 8 and 9 of IOM 0 and 1 (see figure 8-2). The MSU450s are receiving 51% of all connects and, therefore, should be the major area of concern. The access rate for every subsystem is reported on the top of the System Summary Report and it can be seen that the MSU450s have an access rate significantly higher than the other subsystems. All signs indicate that if system throughput is being affected by disk activity, then the MSU450s would be the probable cause of such problems.

The next item to check on these two reports should be the channel usage. The two highest used logical channels of any subsystem should be on a separate PSIA channel of a two-PSIA channel subsystem (see subsection 8.3). Referring to figure 8-2, one can see that logical channel 8 of IOM 0 and IOM 1 has the highest usage, and this is the proper configuration (refer to detailed description in subsection 8.3). If the highest used logical channels are not on separate PSIA channels, the $ XBAR card in the startup configuration section is suspected as the cause. The channels are used in the order given on the $ XBAR card (i.e., if the primary channel is busy, the next channel tried is given on the crossbar). The alternate use of PSIA channels for maximum simultaneity must, therefore, be appropriately specified in the boot deck.

While looking at the System Summary Report, it is also of interest to note the ratio of READ commands to WRITE commands (over two to one in this example). This gives an indication of the nature of the usage of the mass storage space. A quick look at the number of write/verify (WR-VER) commands executed is also of interest as they are essentially double (WRITE, then READ) data transfer commands which require more device and channel time.

The general fraction of utilization for each logical channel gives an indication of the degree of simultaneity of access to the subsystem. If only N of the configured logical channels have nonzero counts, then there were never more than N accesses being performed simultaneously by the subsystem. The proportional relationships among the counts of accesses made over each of the logical channels are quantitative indications of the frequency of occurrence of specific levels of simultaneity. As an example, if we look at figure 8-2, we see that only 4487 connects out a total of 107,273 connects went to Channel 9, IOM 1. This means that only 4487 times during the measuring period were all four MSU450 disk channels being utilized simultaneously. In this example, channel queuing (i.e., shortage of channel power) would not appear to be a problem. This is not to infer that device queuing is not a problem, just that channel queuing does not appear to be a problem. If the number of accesses to the lowest priority channel is a larger percentage of the total accesses, then channel queuing needs to be examined.

8.5.3 System Traces Captured by Monitor Report (File 57). This report contains the number of occurrences of each specific trace type recorded on the data collector tape processed by the CMDRP (figure 8-4). This report provides little, if any, information required by the user for his analysis. This report is always generated and cannot be turned off.

8.5.4 Channel Status Changes Report (File 57). This report lists the initial status for all tape and disk channels configured to the system (figure 8-5). If, during the course of the monitoring session, a given channel or IOM was dropped or added to the system (dynamic reconfiguration) a new report will be produced indicating the activation or deactivation changes and the time that the change occurred. Finally, this report will indicate whether the SSA cache option and FMS cache option are active, and if so, will indicate their initial status and any changes that occur to that status. If a given option is not active, a zero will be reported for each of the values. This report is always generated and cannot be turned off.

8.5.5 Physical Device, Device ID Correlation Table (File 57). Each mass storage device configured in the system is listed with a unique device ID. A typical report is presented in figure 8-6. This unique device is needed since different devices can have the same device number on the Honeywell 6000. (See Device ID 1, Device ID 7, and Device ID 18 of figure 8-6). These unique numbers are referenced in several reports produced by the CMDRP. This report is always generated and cannot be turned off.

8.5.6 Channel Statistics Report (File 57). The Channel Statistics Report is actually a series of reports used to summarize the queuing that occurred over the channels and devices. These reports are processed as a group and are produced by default. The entire series of reports can be turned off via the use of the "OFF" input option (see subsection 8.6.5).

8.5.6.1 Channel Busy and Device Busy Report. The Channel Busy and Device Busy Report is given in figure 8-7. These data are collected into an array during execution of the reduction program and indicates the number of times channel K and device N were both busy at the I/O request link time. Remember that this report is presented by primary channel. The primary channel is busy when all logical channels associated with the primary channel are busy.

8.5.6.2 Channel Busy and Device Free Report. The Channel Busy and Device Free Report, figure 8-8, is generated in a similar method to the Channel Busy and Device Busy Report. If a channel is busy a sufficient number of times, this will indicate the need for more channel power. If, for example, there are no more IOM ports available, a large amount of channel queuing can be a strong indication of the need for another IOM.

8.5.6.3 Channel Free and Device Busy Report. The Channel Free and Device Busy Report, figure 8-9, is generated in a similar manner to the previous

SYSTEM TRACES CAPTURED BY MONITOR REPORT

TRACE	COUNT	TRACE	COUNT	TRACE	COUNT	TRACE	COUNT	TRACE	COUNT	TRACE	COUNT	TRACE	COUNT	TRACE	COUNT
00	0	01	0	02	0	03	0	04	3214	05	0	06	0	07	3214
10	129	11	0	12	0	13	0	14	636	15	0	16	0	17	0
20	0	21	450	22	3214	23	0	24	0	25	0	26	0	27	0
30	0	31	0	32	0	33	0	34	0	35	0	36	0	37	0
40	0	41	0	42	0	43	0	44	0	45	0	46	15	47	0
50	21	51	0	52	0	53	0	54	0	55	0	56	0	57	0
60	0	61	0	62	0	63	0	64	0	65	0	66	0	67	0
70	0	71	0	72	0	73	0	74	0	75	0	76	0	77	0

Figure 8-4. System Traces Captured by Monitor Report

CHANNEL STATUS CHANGES REPORT FOR NMCC2 ON 80/09/20

IOM	CHANNEL	DEACTIVATE/ACTIVATE CHANGE	TIME
0	08	INITIAL ACTIVE	16:50:49.7
0	09	INITIAL ACTIVE	16:50:49.7
0	12	INITIAL ACTIVE	16:50:49.7
0	13	INITIAL ACTIVE	16:50:49.7
0	14	INITIAL ACTIVE	16:50:49.7
0	15	INITIAL ACTIVE	16:50:49.7
0	16	INITIAL ACTIVE	16:50:49.7
0	17	INITIAL ACTIVE	16:50:49.7
0	18	INITIAL ACTIVE	16:50:49.7
1	08	INITIAL ACTIVE	16:50:49.7
1	09	INITIAL ACTIVE	16:50:49.7
1	12	INITIAL ACTIVE	16:50:49.7
1	13	INITIAL ACTIVE	16:50:49.7
1	14	INITIAL ACTIVE	16:50:49.7
1	15	INITIAL ACTIVE	16:50:49.7
1	16	INITIAL ACTIVE	16:50:49.7
1	17	INITIAL ACTIVE	16:50:49.7

INITIAL VALUES FOR SSA CACHE - LAL, MBA, SIZE
000000200000 000000001566 10

INITIAL VALUES FOR FMS CACHE - ABS ADDR, MBA, OPTION WORD, #320 WORD BUFFERS
000000000000 000000000000 000000000000 0

Figure 8-5. Channel Status Changes Report

THE PHYSICAL DEVICE, DEVICE ID CORRELATION TABLE

```
DEVICE ID - 1 IS FOUND ON IOM-0 PUB-08 AND IS DEVICE #01
DEVICE ID - 2 IS FOUND ON IOM-0 PUB-08 AND IS DEVICE #02
DEVICE ID - 3 IS FOUND ON IOM-0 PUB-08 AND IS DEVICE #03
DEVICE ID - 4 IS FOUND ON IOM-0 PUB-08 AND IS DEVICE #04
DEVICE ID - 5 IS FOUND ON IOM-0 PUB-08 AND IS DEVICE #05
DEVICE ID - 6 IS FOUND ON IOM-0 PUB-08 AND IS DEVICE #06
DEVICE ID - 7 IS FOUND ON IOM-0 PUB-12 AND IS DEVICE #01
DEVICE ID - 8 IS FOUND ON IOM-0 PUB-12 AND IS DEVICE #02
DEVICE ID - 9 IS FOUND ON IOM-0 PUB-12 AND IS DEVICE #03
DEVICE ID -10 IS FOUND ON IOM-0 PUB-12 AND IS DEVICE #04
DEVICE ID -11 IS FOUND ON IOM-0 PUB-12 AND IS DEVICE #05
DEVICE ID -12 IS FOUND ON IOM-0 PUB-12 AND IS DEVICE #06
DEVICE ID -13 IS FOUND ON IOM-0 PUB-12 AND IS DEVICE #07
DEVICE ID -14 IS FOUND ON IOM-0 PUB-12 AND IS DEVICE #08
DEVICE ID -15 IS FOUND ON IOM-0 PUB-12 AND IS DEVICE #09
DEVICE ID -16 IS FOUND ON IOM-0 PUB-12 AND IS DEVICE #10
DEVICE ID -17 IS FOUND ON IOM-0 PUB-12 AND IS DEVICE #11
DEVICE ID -18 IS FOUND ON IOM-0 PUB-16 AND IS DEVICE #01
DEVICE ID -19 IS FOUND ON IOM-0 PUB-16 AND IS DEVICE #02
DEVICE ID -20 IS FOUND ON IOM-0 PUB-16 AND IS DEVICE #03
DEVICE ID -21 IS FOUND ON IOM-0 PUB-16 AND IS DEVICE #04
DEVICE ID -22 IS FOUND ON IOM-0 PUB-16 AND IS DEVICE #05
DEVICE ID -23 IS FOUND ON IOM-0 PUB-16 AND IS DEVICE #06
DEVICE ID -24 IS FOUND ON IOM-0 PUB-16 AND IS DEVICE #07
```


Figure 8-6. Physical Device, Device ID Correlation Table

```
CHANNEL BUSY AND DEVICE BUSY REPORT FOR NMCC2 ON 80-10-06

CHANNEL  08  IOM  2  AND DEVICE 01 BOTH BUSY   928  TIMES

CHANNEL  08  IOM  2  AND DEVICE 02 BOTH BUSY  3270  TIMES

CHANNEL  08  IOM  2  AND DEVICE 03 BOTH BUSY   399  TIMES
```

Figure 8-7. Channel Busy and Device Busy Report

```
CHANNEL BUSY AND DEVICE FREE REPORT FOR NMCC2 ON 80-10-06

CHANNEL  06  ICM  2  BUSY AND DEVICE  01  NOT BUSY  1576  TIMES

CHANNEL  14  ICM  2  BUSY AND DEVICE  02  NOT BUSY     2  TIMES

CHANNEL  14  ICM  2  BUSY AND DEVICE  03  NOT BUSY     1  TIME
```

Figure 8-8. Channel Busy and Device Free Report

```
CHANNEL FREE AND DEVICE BUSY REPORT FOR NMCC2 ON 80-10-06

CHANNEL  08  IOM  0  NOT BUSY AND DEVICE  01  BUSY  1858  TIMES

CHANNEL  08  IOM  0  NOT BUSY AND DEVICE  02  BUSY   410  TIMES

CHANNEL  12  IOM  0  NOT BUSY AND DEVICE  01  BUSY    10  TIMES

CHANNEL  12  IOM  0  NOT BUSY AND DEVICE  02  BUSY   698  TIMES
```

Figure 8-9. Channel Free and Device Busy Report

two reports. It is an indication of the number of times that an I/O request was queued because the device was busy but there were available channels. Significant values in this report would be an indication that possible relocation of files is required. The Mass Store Monitor can be used with this data tape to determine the files being accessed on this device.

8.5.6.4 Channel Free and Device Free Report. The Channel Free and Device Free Report, figure 8-10, is an indication of the number of I/O requests that were executed by IOS immediately, without any queuing.

8.5.6.5 GEPR Connect Report. The GEPR Connect Report, figure 8-11, shows the number of times that a trace type 7 was processed, without a preceding trace type 22. This unexpected sequence of traces is supposed to occur whenever the system processes a GEPR (I/O error) event. During data collection, the CM captures an I/O status word indicator and the data reduction program checks this status word to verify that a GEPR has actually occurred. This report will indicate how many confirmed GEPRs (i.e., the status word was set) have occurred and how many suspected GEPRs have occurred (i.e., the status word was not set). This report will be of little aid to the analyst and CCTC is still investigating the reason for this trace occurrence, when the status word is not set.

8.5.6.6 Lost Interrupt Report. The Lost Interrupt Report, figure 8-12, indicates that a trace type 7 is being processed for a busy device and/or channel. This is an impossible event and is an indication that a trace type 4 has not been generated. This is usually an indication that the system has generated a lost interrupt. However, if lost data has been generated during data collection, this report may indicate many lost interrupts, which really did not occur.

8.5.6.7 Device ID STIOS Not Connected Report. The Device ID STIOS Not Connected Report, figure 8-13, shows the number of start I/Os that were dropped for each device. A start I/O is dropped when it is found that a STIO trace has occurred for a device and within a user-defined timeframe (5-second default), no connect has been received for the start. The cause of this condition currently is under analysis. This report is also presented by the device ID and must be correlated by using the Device ID Correlation Report.

8.5.6.8 Entries Still in Queue Report. The Entries Still in Queue Report, figure 8-14, shows all entries remaining in the CMDRP queues. This report shows the device number, channel number, IOM number, queue location, and time (in milliseconds) the entry has been in the queue. These entries were active when the monitor terminated.

8.5.7 Idle Monitor Report (File 57). If the Idle Monitor was active when the Channel Monitor was running an Idle Report will be produced next (see figure 8-15). The first few lines will indicate, for each processor, the

```
CHANNEL FREE AND DEVICE FREE REPORT FOR OSCC2 ON 79-12-06

CHANNEL  08  IOM  0  AND DEVICE  01  FREE  15285  TIMES
CHANNEL  08  IOM  0  AND DEVICE  02  FREE   4114  TIMES
CHANNEL  08  IOM  0  AND DEVICE  03  FREE  31540  TIMES
CHANNEL  08  IOM  0  AND DEVICE  04  FREE  13094  TIMES
CHANNEL  08  IOM  0  AND DEVICE  05  FREE   4292  TIMES
CHANNEL  08  IOM  0  AND DEVICE  06  FREE   1007  TIMES
CHANNEL  12  IOM  0  AND DEVICE  01  FREE   3290  TIMES
CHANNEL  12  IOM  0  AND DEVICE  02  FREE   7012  TIMES
CHANNEL  12  IOM  0  AND DEVICE  03  FREE   2724  TIMES
CHANNEL  12  IOM  0  AND DEVICE  04  FREE   3558  TIMES
CHANNEL  12  IOM  0  AND DEVICE  05  FREE   3626  TIMES
CHANNEL  12  IOM  0  AND DEVICE  06  FREE   8781  TIMES
CHANNEL  12  IOM  0  AND DEVICE  07  FREE    890  TIMES
CHANNEL  12  IOM  0  AND DEVICE  08  FREE  15285  TIMES
CHANNEL  12  IOM  0  AND DEVICE  09  FREE   1087  TIMES
CHANNEL  12  IOM  0  AND DEVICE  10  FREE   4147  TIMES
CHANNEL  12  IOM  0  AND DEVICE  11  FREE  18330  TIMES
CHANNEL  16  IOM  0  AND DEVICE  01  FREE    892  TIMES
CHANNEL  16  IOM  0  AND DEVICE  02  FREE    456  TIMES
CHANNEL  16  IOM  0  AND DEVICE  03  FREE    456  TIMES
CHANNEL  16  IOM  0  AND DEVICE  04  FREE    456  TIMES
CHANNEL  16  IOM  0  AND DEVICE  05  FREE    456  TIMES
CHANNEL  16  IOM  0  AND DEVICE  06  FREE    456  TIMES
CHANNEL  16  IOM  0  AND DEVICE  07  FREE    456  TIMES
```

Figure 8-10. Channel Free and Device Free Report

GEPR CONNECT REPORT FOR NMCC2 ON 82-01-19

CHANNEL	08	IOM	0	DEVICE	05	#	OF CONFIRMED GEPRS	0	#	OF SUSPECTED GEPRS	23
CHANNEL	08	IOM	0	DEVICE	07	#	OF CONFIRMED GEPRS	0	#	OF SUSPECTED GEPRS	13
CHANNEL	08	IOM	0	DEVICE	09	#	OF CONFIRMED GEPRS	0	#	OF SUSPECTED GEPRS	9
CHANNEL	08	IOM	0	DEVICE	11	#	OF CONFIRMED GEPRS	0	#	OF SUSPECTED GEPRS	19
CHANNEL	12	IOM	0	DEVICE	15	#	OF CONFIRMED GEPRS	0	#	OF SUSPECTED GEPRS	5
CHANNEL	14	IOM	0	DEVICE	05	#	OF CONFIRMED GEPRS	0	#	OF SUSPECTED GEPRS	26

Figure 8-11. GEPR Connect Report

```
LOST INTERRUPT REPORT FOR NMCC2 ON 80-10-06

IOM  0  PUB  12  LOST INTERRUPT COUNT      1
```


Figure 8-12. Lost Interrupt Report

DEVICE ID STIOS NOT CONNECTED FOR NMCC2 ON 80-10-06

DEVICE ID 8 # OF STARTS DROPPED 1

DEVICE ID 10 # OF STARTS DROPPED 1

Figure 8-13. Device ID STIOS Not Connected

ENTRIES STILL IN QUEUE FOR NMCC2 ON 80-10-06

DEVICE #	CHA #	IOM	QUEUE	TIME IN QUEUE(MS)
20	12	0	189088	33

Figure 8-14. Entries Still in Queue

IDLE REPORT FOR SYSTEM OSCC2 ON 79-07-13

PROCESSOR 0 WENT IDLE 101934 TIMES FOR A % IDLE OF 38

PROCESSOR 1 WENT IDLE 64039 TIMES FOR A % IDLE OF 29
AVERAGE SYSTEM % IDLE WAS 34
% OF IDLE TIME DURING WHICH IO WAS ACTIVE 100
AVERAGE NUMBER OF OUTSTANDING IO's WHEN THE SYSTEM WENT IDLE WAS 7

DEVID	% OF IDLE TIME WITH ACTIVE IO	AVERAGE QUEUE SIZE
1	95	4.66
2	12	0.19
3	21	0.33
4	46	1.16
8	3	0.05
10	1	0.02
11	1	0.02
14	71	0.76
15	4	0.08
16	1	0.02
24	1	0.02

Figure 8-15. Idle Report

number of times that processor went idle and the percent idleness of the that processor. This is followed by an average idle percentage for the entire system. The next line of output indicates for what percentage of system idle time there was active disk I/O; i.e., I/O in progress or I/O request queued. This figure would give a good indication as to whether the CPU is going idle because of a lack of work or rather because work is being delayed by the slower peripheral devices. In figure 8-15, we see that for 100 percent of the idle time, the CPU was actually waiting on disk I/O to be performed. In this case, the CPU could be put to better utilization, if only the speed of the disk subsystem could be increased. The next line of output indicates how many outstanding I/Os were present when the system went idle.

Following this output, a table is generated for every disk device that had any active I/O on it when the CPU went idle. For each such device, the percent of Idle Time during which it had active I/O and the average queue size at that time is given. In figure 8-15, we see that during 95 percent of CPU idle time device ID #1 had outstanding I/O to an average length of 4.66. Device #4 had active I/O for 46 percent of the CPU idle time with an average length of 1.2. By examining this report, we can see that if the queues on these two devices could be reduced there is sufficient CPU power available to handle the additional workload that would be generated. As with earlier reports, the Device ID Correlation Report should be used to convert a unique device ID into an IOM, PUB, device number.

8.5.8 Proportionate Device Utilization Report (File 57). This report shows the proportionate utilization of each device configured on the mass storage subsystem. Figure 8-16 is an example. This histogram identifies each unique device ID (device number zero is an MPC controller) and provides both a count of the number of accesses made to each device (under the column headed INDIV. NUMBER) as well as the percent of all accesses which were to each device (under the column headed INDIV. PRC). The histogram shows the proportionate utilization of each device (i.e., the percent of all accesses which went to each device) in a graphical form. The physical device that each "Device ID" of the histogram represents is shown in the Physical Device ID Correlation Table (see figure 8-6). This report is always generated and cannot be turned off. In this report the user is looking for a device or devices which have significantly more utilization than others in the system. This highly used device would then be a potential bottleneck.

It is desirable, but not always practical, to have equal utilization for each device. The user should be reminded that data in figure 8-16 is cumulative over the monitoring period. The actual accessing pattern could have been periodic with the following form: Many accesses to device 4 followed by many accesses to device 3 followed by many accesses to device 2 followed by many accesses to device 1, etc. Each device could have been a bottleneck for a subperiod of the total monitoring period. This could also

DISTRIBUTION COLLECTED ON SYSTEM OSCC2 AT 13:49:04 ON 81-09-28

DEVICE UTILIZATION BY DEVICE NUMBER

INDIV. NUMBER	CUMUL. NUMBER	CUMUL. PRC	INDIV. PRC	DEVICE ID		PERCENT OF OCCURRENCE 00 05 10 15 20 25 30 35 40 45 50	REPORT 1
						I....I....I....I....I....I....I....I....I....I....I....	
12	12	0.000	0.000	0-	0	I	
18947	18959	9.835	9.828	1-	1	Ixxxxxxxxxx	
8206	27165	14.091	4.257	2-	2	Ixxxx	
20269	47434	24.606	10.514	3-	3	Ixxxxxxxxxxx	
34033	81467	42.260	17.654	4-	4	Ixxxxxxxxxxxxxxxxxx	
238	81705	42.384	0.124	5-	5	I	
25580	107285	55.653	13.269	6-	6	Ixxxxxxxxxxxxx	
.	.	.	.	.	.	.	
.	.	.	.	.	.	.	
.	.	.	.	.	.	.	

192774 ENTRIES 0 OUT OF RANGE

Figure 8-16. Proportionate Device Utilization Report

have been the case if the proportionate utilization of each device was equal. The Channel Monitor can be used to uncover this cyclic type of usage. In addition, the Chronological Device Utilization Report (see subsection 7.5.18) was designed to uncover this type of problem by breaking down device utilization over time, rather than by utilizing a histogram. Nevertheless, when a single or small number of devices has a disproportionately large share of the accesses, they are potential bottlenecks and their usage should be further analyzed.

This report will show all connects that were issued to a given device. This includes all read/write connects, as well as any command type connects issued to a given device.(See subsection 7.3).

8.5.9 Queue Length and Queue Time Histograms (File 57). Figure 8-17 shows the I/O queue length and queue time histograms for the I/O requests to devices as they are processed by IOS. These reports occur in pairs, one pair for each device. These reports are generated only for mass store devices. The first report in the pair is a length report and the second is a time report. The histogram is read in the same manner as the Proportionate Device Utilization Report. In addition to individual percentages, cumulative percentages are also reported. In figure 8-17 (part 2) we see that 99.457% of all requests have a queue time of 12 ms or less while 99.674% of all requests have a queue time of 22 ms or less. The number of entries in these two reports might not be equal. The first report is generated at the time of the request. When an I/O request is made, an entry is made to the histogram indicating the number of requests outstanding or in progress to this device. The second report is generated at the time of the actual connect and indicates the length of time between the request and the actual connect. Since observations have indicated that some STIOS are never connected, the first report may have a higher number of entries than the second. These two reports can be correlated by subtracting the number of STIOS not connected, for a given device, from the number of entries in the queue length histogram. As an example, in figure 8-17 (part 1) we have a total of 921 entries. However, figure 8-17 (part 2) shows only 920 entries. If we now check figure 8-13 for this device (ID #10) we find a total of 1 entry reported. If this figure is subtracted from 921, the two reports correlate.

Figure 8-18 shows the channel queue length and queue time for the I/O queue entries as they are used in the system channel environment. These histograms will be produced for both mass store and tape channels. Once again, the number of entries in these figures will not correlate because of STIOS issued over a given channel, and not over a single device. On the system used for generating the figures in this document, devices 8, 10 and 20 are configured on IOM-0 PUB-12. If we add all references to these devices (see figure 8-13), we have a total of 2. In addition, if we add in references to those devices from figure 8-14 (entries that are still in the queue), the two figures will correlate.

DISTRIBUTION COLLECTED ON SYSTEM NMCC2 AT 11:44:41 ON 80-10-06

I-O QUEUE LENGTH FOR IOM-0,PUB-12,DEVICE-04-- DSS191 DA4

INDIV. NUMBER	CUMUL. NUMBER	CUMUL. PRC	INDIV. PRC	QUEUE LENGTH	PERCENT OF OCCURRENCE 00 10 20 30 40 50 60 70 80 90 100	REPORT 33
					I....I....I....I....I....I....I....I....I....I....I	
914	914	99.240	99.240	0	IXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX	
7	921	100.000	0.760	1- 1	I	

921 ENTRIES TOTAL AVERAGE = 0.00760 VARIANCE = 0.008 STANDARD DEVIATION = 0.087

Figure 8-17. Device Queue Length and Time Histograms (Part 1 of 2)

DISTRIBUTION COLLECTED ON SYSTEM NMCC2 AT 11:44:41 ON 80-10-06

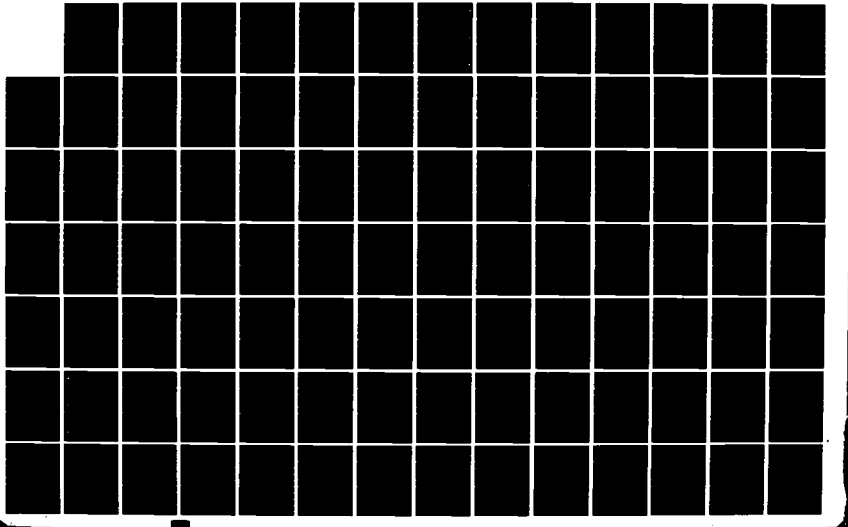
I-O QUEUE TIME-MS FOR IOM-0,PUB-12,DEVICE-04-- DSS191 DA4

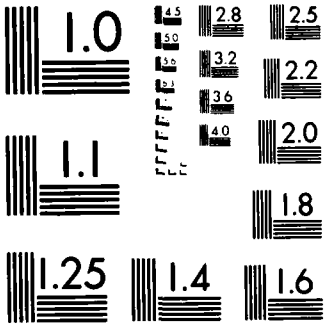
INDIV. NUMBER	CUMUL. NUMBER	CUMUL. PRC	INDIV. PRC	QUEUE TIME	PERCENT OF OCCURRENCE 00 10 20 30 40 50 60 70 80 90 100	REPORT 34
					I....I....I....I....I....I....I....I....I....I....I	
913	913	99.239	99.239	2.000	IXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX	
0	913	99.239	0.	4.000	I	
0	913	99.239	0.	6.000	I	
0	913	99.239	0.	8.000	I	
1	914	99.348	0.109	10.000	I	
1	915	99.457	0.109	12.000	I	
0	915	99.457	0.	14.000	I	
0	915	99.457	0.	16.000	I	
0	915	99.457	0.	18.000	I	
1	916	99.565	0.109	20.000	I	
1	917	99.674	0.109	22.000	I	
0	917	99.674	0.	24.000	I	
0	917	99.674	0.	26.000	I	
0	917	99.674	0.	28.000	I	
1	918	99.783	0.109	30.000	I	
1	919	99.891	0.109	32.000	I	
0	919	99.891	0.	34.000	I	
0	919	99.891	0.	36.000	I	
0	919	99.891	0.	38.000	I	
0	919	99.891	0.	40.000	I	
0	919	99.891	0.	42.000	I	
0	919	99.891	0.	44.000	I	
1	920	100.000	0.109	46.000	I	

920 ENTRIES TOTAL . . AVERAGE = 0.18161 VARIANCE = 5.347 STANDARD DEVIATION = 2.312

Figure 8-17. (Part 2 of 2)

AD-A131 954 GENERALIZED MONITORING FACILITY USERS MANUAL CHANGE 6 5/8
(U) COMMAND AND CONTROL TECHNICAL CENTER WASHINGTON DC
B M WALLACK ET AL. 01 JUL 83 CCTC-CSM-UM-246-82-CHG-6
UNCLASSIFIED DCA/DF-83/001B F/G 9/2 NL





MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

DISTRIBUTION COLLECTED ON SYSTEM NMCC2 AT 11:44:41 ON 80-10-06

I-O QUEUE LENGTH FOR IOM-0,PUB-12

INDIV. NUMBER	CUMUL. NUMBER	CUMUL. PRC	INDIV. PRC	QUEUE LENGTH	PERCENT OF OCCURRENCE 00 10 20 30 40 50 60 70 80 90 100	REPORT 6
9048	9048	100.000	100.000	0- 0	I....I....I....I....I....I....I....I....I....I....I IXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX	

9048 ENTRIES TOTAL AVERAGE = 0. VARIANCE = 0. STANDARD DEVIATION = 0.

Figure 8-18. Channel Queue Length and Queue Time Report (Part 1 of 2)

DISTRIBUTION COLLECTED ON SYSTEM NMCC2 AT 11:44:41 ON 80-10-06

I-O QUEUE TIME-MS FOR IOM-0,PUB-12

INDIV. NUMBER	CUMUL. NUMBER	CUMUL. PRC	INDIV. PRC	QUEUE TIME	PERCENT OF OCCURRENCE	REPORT
					·00 10 20 30 40 50 60 70 80 90 100	7
					I....I....I....I....I....I....I....I....I....I....I	
9045	9045	100.000	100.000	2.000	IXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX	

9045 ENTRIES TOTAL AVERAGE = 0. VARIANCE = 0. STANDARD DEVIATION = 0.

Figure 8-18 (Part 2 of 2)

In determining the length of a channel queue, the following is considered:

- A channel queue of length 0 exists when there is at least one nonbusy, primary/logical channel.
- When all primary/logical channels become busy, the length of the channel queue is calculated by summing the length of all device queues on that channel.
- As an example, assume we had two channels configured with three devices. If device 1 became busy and 4 more requests were made for that device, we would have a device queue of 4 and a channel queue of 0. If during the same time, device 2 received 3 requests, we would generate a queue length 2 for that device but the channel queue would still remain at 0. In the situation described thus far, device contention, not a shortage of channels, is the problem. If we now assume that a connect was issued to device 3, we have now created a channel problem. There is no available channel for the I/O request. At this point, a channel queue of 7 (4 requests for device 1 + 2 requests for device 2 + 1 request for device 3) would be generated. Thus, we see that channel queues may increase in a nonsequential fashion.

8.5.10 Service Time Histograms (File 57). In all of the device histograms, it should be noted that the name of the device is also given. In figure 8-17, queuing statistics were presented for a device with the name DA4. If an exchange took place and the DA4 disk pack was moved from 0-12-04 to 0-12-07, the data reduction program will account for that exchange and any connects that are made to 0-12-07 will be reported on this histogram and not to the 0-12-07 histogram.

In the upper right-hand corner of the report, a report number is indicated. This report number is used only to distinguish one histogram from another and in no way indicates the device to which the report refers. In addition, report numbers may not appear sequentially and this in no way is indicative of a problem. All histogram reports are generated by default and may be turned off by using the LIMITS option (subsection 8.6.13).

Figure 8-19 shows the I/O service time histograms for each I/O channel and device. Each histogram is given in 2ms intervals. The I/O service time is defined as the time (in ms) from connect to the time that IOS processes the terminate interrupt for the I/O request. These histograms are generated for both mass store and tape channels, as well as all mass store devices. On the bottom line of this report, an indication is given as to the percent of total time that this device or channel was busy.

The three device-oriented histograms, just described, have entries placed in them for every connect issued to the device (not just multicommands such as seek-read or seek-write).

DISTRIBUTION COLLECTED ON SYSTEM NMCC2 AT 11:44:41 ON 80-10-06

I-O SERVICE TIME (MS) FOR IOM-0,PUB-14

INDIV. NUMBER	CUMUL. NUMBER	CUMUL. PRC	INDIV. PRC	SERIVCE TIME	PERCENT OF OCCURRENCE 00 05 10 15 20 25 30 35 40 45 50	REPORT 103
					I....I....I....I....I....I....I....I....I....I....I	
10	10	45.455	45.455	2.000	IXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX	
5	15	68.182	22.727	4.000	IXXXXXXXXXXXXXXXXXXXXXXXX	
2	17	77.273	9.091	6.000	IXXXXXXXXX	
0	17	77.273	0.	8.000	I	
0	17	77.273	0.	10.000	I	
0	17	77.273	0.	12.000	I	
0	17	77.273	0.	14.000	I	
0	17	77.273	0.	16.000	I	
0	17	77.273	0.	18.000	I	
0	17	77.273	0.	20.000	I	
.	.	.	.	.	.	
.	.	.	.	.	.	
.	.	.	.	.	.	

22 ENTRIES TOTAL AVERAGE = 9.42141 VARIANCE = 203.826 STANDARD DEVIATION = 14.277

% BUSY = 0.03

Figure 8-19. I/O Service Time Report (Part 1 of 2)

DISTRIBUTION COLLECTED ON SYSTEM NMCC2 AT 11:44:41 ON 80-10-06

I-O SERVICE TIME (MS) FOR [illegible]-0,PUB-12,DEVICE-10-- DSS191 DBO

INDIV. NUMBER	CUMUL. NUMBER	CUMUL. PRC	INDIV. PRC	SERIVCE TIME	PERCENT OF OCCURRENCE 00 05 10 15 20 25 30 35 40 45 50	REPORT 53
					I....I....I....I....I....I....I....I....I....I....I	
3	3	13.043	13.043	2.000	IXXXXXXXXXXXX	
1	4	17.391	4.348	4.000	IXXXX	
2	6	26.087	8.696	6.000	IXXXXXXXXX	
2	8	34.783	8.696	8.000	IXXXXXXXXX	
0	8	34.783	0.	10.000	I	
2	10	43.478	8.696	12.000	IXXXXXXXXX	
2	12	52.174	8.696	14.000	IXXXXXXXXX	
5	17	73.913	21.739	16.000	IXXXXXXXXXXXXXXXXXXXXXX	
.	.	.	.	.	.	
.	.	.	.	.	.	
.	.	.	.	.	.	

23 ENTRIES TOTAL AVERAGE = 14.01183 VARIANCE = 97.224 STANDARD DEVIATION = 9.860

% BUSY = 0.04

Figure 8-19. (Part 2 of 2)

8.5.11 Activity Statistic Report (Files 23 and 24). Figure 8-20 shows the Activity Statistical Report (Parts 1 and 2). For each activity run in the system, Part 1 of this report displays the SNUMB, the activity number, the queue time (in .1 sec) that the activity accumulated for tapes and disks, the connect time (in .1 sec) that the activity accumulated for tapes and disk, and finally, the IDENT and USERID of the activity. The connect time on this report is computed by using as the start time the issuance of trace 7 and as the stop time the issuance of a trace 4. The value output in this report may differ significantly from the value output by SCF. In the opinion of CCTC (C751), the value given by this report more closely approximates the true connect time for an activity. Part 2 of this report gives the number of connects issued to tape and disk by each SNUMB and activity number, and the average queue time for each connect.

By using this report, it is possible to determine those activites which are being queued the most. By examining Mass Store Monitor output for these activities, it could be determined what files and packs were referenced by the activity and possibly some file reorganization could be performed.

System Schedular information for each activity (activity 0) is not reported separately. All activity 0 data is accumulated and reported as a single entry under the SNUMB $GENB. Any job whose SNUMB begins with a $ will not be considered as the System Scheduler, even if it processes as activity zero. In addition, NACE DERAIL jobs and BASHT jobs will also be processed as regular activities, even though they run under an activity number of zero. Finally, there is an input option (ZERO) that allows the user to define other zero activity jobs, that should not be considered as part of the System Scheduler.

From all the reports produced thus far, it is still not possible to determine if particular jobs are in conflict with each other. This information would be extremely useful in being able to move conflicting program files to different disk packs or in scheduling conflicting jobs during different times of the day. There is a CMDRP option that allows the user to obtain a Job Conflict Report for up to four unique devices (see subsection 8.5.12). The user would first determine those devices displaying the largest degree of queuing, or those devices containing the files of the programs receiving the largest queue times and rerun the CMDRP requesting the Job Conflict Report option (see subsection 8.6.1).

8.5.12 Job Conflict Report (Files 31, 32, 33, 34). In figure 8-21, we have a Job Conflict Report for Device ID #10. The first line of the report will give the date, time, and system name on which the data was collected. A separate report will be produced for each device requested by the user with the Device ID number appearing in the upper right hand corner of the report. The first column of the report will list the SNUMB/Activity Number of every job that was queued or caused some other job to be queued on this particular device. The USERID and IDENT for this SNUMB is also reported. Under the column "QUEUED BY" will appear a list of SNUMBs that caused such

a delay. For example, $CALC was delayed by $ PALC 1 time and by $ SYOT 1 time. Under the column "QUEUED" will appear a list of SNUMBs that were delayed by this particular job and the number of times this delay occurred. Therefore, $CALC delayed $ GENB 1 time and itself just one time. Anytime that a given job has been queued by, or queued, 20 different

THIS PAGE LEFT INTENTIONALLY BLANK

ACTIVITY STATISTIC REPORT PART 1 FOR NMCC2 ON 80-10-06

SNUMB ACTIVITY	QUEUE TIME (.1 SEC) TAPES	QUEUE TIME (.1 SEC) DISK	CONNECT TIME (.1 SEC) TAPES	CONNECT TIME (.1 SEC) DISK	IDENT	USERID
7151T- 1	0	0	0	16	1820020/30/6031,BETTY LAFAVERS	DJ3JI32419
TEST - 1	0	0	0	11	1820251/30/2773,C751	DJ8XI70203
7170A- 1	0	0	0	4	XXXNO,SC-PURGE	OPNSUTIL
7170A- 2	12	0	100	65	XXXNO,SC-PURGE	OPNSUTIL
$CALC- 0	0	1	0	49		
$PALC- 1	5	3	1	124		
$SYOT- 0	0	0	0	18		
$RTIN- 0	0	0	0	10		
TSS - 1	0	0	0	167		
$TOLT- 0	0	0	0	15		
$LOGN- 0	0	0	0	2		
NCP - 1	0	0	0	6		
$GENB- 0	0	7	0	1580		
DMSTA- 1	0	0	0	217	ZZZ,OPNSTIL	OPNSUTIL
TEST - 1	0	0	0	10	1820251/30/2773,C751	DJ8XI70203
7170A- 3	0	0	0	2	XXXNO,SC-PURGE	OPNSUTIL
MUM - 1	1	0	235	2	ZZZ,OPNSTIL	OPNSUTIL
HEALS- 1	0	0	0	4	ZZZ,OPNSTIL	HEALS-DATA

Figure 8-20. Activity Statistic Report (Part 1 of 2)

ACTIVITY STATISTIC REPORT PART 2 FOR NMCC2 ON 80-10-06

SNUMB ACTIVITY	NUMBER OF CONNECTS ISSUED TAPE	DISK	AVERAGE QUEUE TIME (MILSEC) PER CONNECT TAPE	DISK
7151T- 1	0	100	0	0
TEST - 1	0	39	0	0
7170A- 1	0	19	0	0
7170A- 2	658	668	1	0
$CALC- 0	0	261	0	0
$PALC- 1	172	1197	2	0
$SYOT- 0	0	101	0	0
$RTIN- 0	0	54	0	0
TSS - 1	0	516	0	0
$TOLT- 0	0	105	0	0
$LOGN- 0	0	3	0	0
NCP - 1	0	29	0	0
$GENB- 0	0	7287	0	0
DMSTA- 1	0	1728	0	0
TEST - 1	0	39	0	0
7170A- 3	0	14	0	0
MUM - 1	166	15	0	0
HEALS- 1	0	106	0	0

Figure 8-20. (Part 2 of 2)

```
JOB DEVICE CONFLICT REPORT COLLECTED ON 06-10-80 AT 11:45 FOR SYSTEM NMCC2 FOR DEVICE ID 10
  SNUMB/ACT   USERID          QUEUED BY    TIMES   QUEUED    TIMES   IDENT

  7151T-  1   DJ3JI32419                           $SYOT-  0      1    1820020/30/6031,BETTY LAFAVERS
                              $SYOT-  0      1
                            **********************************************************************************
  $CALC-  0                                         $GENB-  0      1
                                                    $CALC-  0      1
                              $PALC-  1      1
                              $SYOT-  0      1
                              $CALC-  0      1
                            **********************************************************************************
  $PALC-  1                                         $CALC-  0      1
                            **********************************************************************************
  $SYOT-  0                                         7151T-  1      1
                                                    $CALC-  1      1
                              7151T-  1      1
                            **********************************************************************************
  $GENB-  0                   $CALC-  0      1
                            **********************************************************************************
  7170A-  2   OPNSUTIL                              7170A-  2      1    XXXNO,SC-PURGE
                              7170A-  2      1
                            **********************************************************************************

THE TOTAL NUMBER OF CONNECTS QUEUED ON THIS DEVICE WAS       7
```


Figure 8-21. Job Conflict

SNUMBs, it will be dumped out. This dumping is due to internal array size restrictions. When this dump is performed a message will be printed. No data will be lost and this job will appear again.

This report allows the user to identify exactly which programs are causing the queuing on a given device. The Mass Store Monitor reports can then easily be used to determine the program files on this device that are in conflict.

8.5.13 Special Job Processing Report by Device (File 32). The user has the option of producing two special reports that describe in detail the queue and service times being obtained by Specific Jobs. These reports can be obtained by invoking the "JOB" input option (see subsection 8.6.10). In the Activity Statistic Report (subsection 8.5.11), the total queue time and total service time for every activity processed during the session is reported. In the Special Job Processing Report by Device, these figures are displayed in detail by device number. In this way, it is possible to determine which device/devices are contributing the most to the total queue/service times obtained by a job. Each line of the report (see figure 8-22) describes the service being received by a job (column 1) for a particular device (column 2). The device number can be related to a particular disk drive by using the Device ID Correlation Report (figure 8-6). If a job references more than 25 unique devices, only the first 25 devices referenced will be reported. Column 3 provides the total queue time in milliseconds and column 5 provides the total service time in milliseconds. Column 7 provides the total number of connects issued by this job to this device.

Columns 4 and 6 provide average queue and service times (column 3/7 and column 5/7 respectively). The final column indicates what percentage of the job's total connects were processed on this particular device.

8.5.14 Special Job Processing Report Per 10 Minutes (File 32). The second of the special reports produced by the "JOB" input option is shown in figure 8-23. This report is produced by default every 10 minutes but the user can alter this time restriction by using the "RATE" option (see subsection 8.6.11). Figure 8-23 was produced at 5-minute intervals. This report provides similar information to that described in subsection 8.5.13 except that it is reported for the total job activity, not broken down by individual devices. This report can be very useful in tracking job performance over time. Particular times of day when queue time is substantially higher than normal can be identified. A correlation between the Response Time of TS1, or HI-TRAX, over time (as reported by the Communications Monitor), can be made with this report. If job performance (TS1, FTS, TELNET, HI-TRAX, etc.) is known to be unacceptable at certain times of day, this report can be used to see if the unacceptable performance was due to excessive delay in the I/O subsystem performance.

SPECIAL JOB PROCESSING REPORT BY DEVICE FOR SYSTEM NMCC2 ON 82-01-19

SNUMB	DEVICE #	TOTAL QUEUE	AVERAGE TIME(MS)	TOTAL SERVICE	AVERAGE TIME(MS)	TOTAL CONNECTS	% OF TOTAL CONNECTS
TS1	32	394	3	2122	20	102	0.29
TS1	31	0	0	5	5	1	0.00
TS1	35	0	0	55	18	3	0.01
.	.	.	.	.	.	.	.
.	.	.	.	.	.	.	.
.	.	.	.	.	.	.	.
2460T	1	114047	1	1498991	16	91231	56.35
2460T	21	6244	0	564632	12	43668	26.97
2460T	4	13615	2	114371	25	4558	2.82
2460T	3	9327	1	109815	18	5866	3.62
2460T	2	19352	1	248906	16	14726	9.10
2460T	22	168	0	18256	17	1064	0.66
.	.	.	.	.	.	.	.
.	.	.	.	.	.	.	.
.	.	.	.	.	.	.	.

Figure 8-22. Special Job Processing Report By Device

SPECIAL JOB PROCESSING REPORT PER 5 MINUTE REPORT FOR SYSTEM NMCC2 ON 82-01-19

TIME	SNUMB	TOTAL QUEUE	AVERAGE TIME(MS)	TOTAL SERVICE	AVERAGE TIME(MS)	TOTAL CONNECTS
1300	TS1	36	0	3687	41	89
1300	2460T	4696	0	169300	14	11999
1300	BARRY	NOT CURRENTLY ACTIVE				
1300	NCP	NOT CURRENTLY ACTIVE				
1300	TOTAL	9185	0	304378	17	17223
1305	TS1	2612	2	36967	33	1107
1305	2460T	9548	1	156430	17	8971
1305	BARRY	NOT CURRENTLY ACTIVE				
1305	NCP	77	4	781	41	19
1305	TOTAL	30326	1	568022	23	23718
1310	TS1	3785	3	47657	38	1241
1310	2460T	5018	0	175921	16	10984
1310	BARRY	NOT CURRENTLY ACTIVE				
1310	NCP	1205	21	2694	48	56
1310	TOTAL	26197	1	607949	31	19171

Figure 8-23. Special Job Processing Report Per 5 Minute Report

8.5.15 Special Processing Messages. During the course of processing, several special processing messages may be generated by the CMDRP. Most of these are for information purposes only and can be ignored by the analyst. Following is a list and brief explanation of the most common messages.

- o MONITOR MUM WAS ACTIVE . . . .
  This message is produced for each monitor that was active during the monitoring session.

- o RUN BEING TERMINATED. DATA FOR MONITOR CM . . . .
  The CM was not active and so no reports can be produced.

- o PROCESSOR #N IS (AVAILABLE/RELEASED) AT (TIME)
  This message will indicate the assignment or releasing of processors and provides the processor number involved as well as the time of day the assignment/releasing of processors occurred.

- o FOLLOWING INPUT OPTIONS HAVE BEEN CHOSEN
  An echo print of all nonstandard input options selected will be produced. If any input option is described incorrectly, an error message will be generated indicating the type of error and the card number in error. The user should correct the error and resubmit the job for processing.

- o INCREASE THE VALUE OF NRDEV . . . .
  The number of devices configured on the system exceed the number permitted by the program. See subsection 8.4 for action to be taken.

- o INCREASE THE VALUE OF NRCHAN . . . .
  The number of channels configured on the system exceed the number permitted by the program. See subsection 8.4 for action to be taken.

- o FOR INFORMATION ONLY
  This message will then be followed by several lines of output describing special record types that have been processed, or special processing events that have been executed by the CMDRP. In most cases, the message can be ignored. Those messages that are important, and reveal an error in processing logic, will be described below. All other messages will not be described.

- o JULIAN DATES ARE BAD - RUN TERMINATED
  Every GMF data record is preceded by the current Julian date. The CMDRP has found a Julian date that does not agree with the Julian date found on previous records. This can occur when an old GMF data tape is used without degaussing. Old data is on the tape, and if the new data failed to write an end of file mark on the tape because of a system crash or malfunction, the CMDRP, after

reaching the end of the new data, would attempt to process the old data without realizing that is was old. The check on the Julian date prevents this from happening. The CMDRP will terminate cleanly and all reports will be produced.

- HAVE INCREASING OR BAD SEQUENCE NUMBERS . . . .
  or
  BAD TRACE RECORD . . . .
  A problem has occurred in reading the data tape. If the run is reprocessed, the error may disappear. If it reoccurs, then the tape was generated with an error. In most cases, the CMDRP will recover and processing will not be significantly affected.

- PROCESSING TERMINATED BY NXTRECRD . . . .
  CMDRP has requested that the operator mount a new tape and the operator has responded that he did mount the new tape or is unable to mount the new tape. If he has mounted the new tape, CMDRP is unable to match the initial record or the new tape with the last record on the old tape. User should check the data collection procedure to insure that correct tapes were mounted during the data collection phase. CMDRP will terminate cleanly and all reports will be produced.

- . . . . . CALL CCTC AT . . . . . .
  A series of messages may be produced which indicate a severe processing error. If these occur, the output of the run should be considered suspect until further clarification is obtained from CCTC.

8.6 Default Option Alteration

Most users rely upon the standard CMDRP report formats and their default values as these suit a wide range of needs. A capability to change reports is built into CMDRP.

All inputs are free format with the only requirement being that if any value is to be a zero, the user must type the number 0 on the data card. A zero value may not be inputted as a blank. At least one data card is required; that being the word "END" punched on the data card. The "END" card must be the last data card inputted to the program. It is used to signify the end of input, or the fact that there is no input available. In the Channel Monitor, all reports except the Job Conflict Report (subsection 8.5.12), the Special Job Processing Report by Device (subsection 8.5.13) and the Special Job Processing Report Per 10 Minute Report (subsection 8.5.14) are produced unless explicitly turned off. If the Channel Statistics Reports (subsection 8.5.6) are not desired, they may be turned off using the "OFF STANDARD" format. Individual reports cannot be turned on or off and individual histograms also cannot be turned on or off. If the Activity Statistic Report (subsection 8.5.11) is not desired it may be turned off, using the "OFF STATISTICS" format.

8.6.1 Job Device Conflict Report (Action Code QDEV). This option allows the user to produce the Job Device Conflict Report for up to four (4) different devices. If more than four devices are requested, only the first four devices will be analyzed. The first data card contains the word "QDEV". The second data card contains the number of devices to be analyzed. This value cannot exceed 4. The third data card contains a list of unique device IDs, separated by at least one blank. The unique device IDs can be obtained from the Physical Device ID Correlation Report.

8.6.2 Program Debug Options. There are several debug options available to the user, none of which should be selected unless the user is very familiar with the data reduction program. These options produce large amounts of output, useful only when debugging data reduction program/logic errors.

8.6.2.1 Program Number Debug (Action Code DPRG). This will allow debugging for a given program number. Card 1 contains the word "DPRG", and card 2 contains the program number for which debugging is to occur.

8.6.2.2 Device Debug (Action Code DDEV). This will allow debugging for a given unique device ID. Card 1 contains the word "DDEV", and card 2 contains the unique device ID for which debugging is to occur.

8.6.2.3 Queue Location Debug (Action Code DQUE). This option will allow debugging for a given IO queue location. Card 1 contains the word "DQUE", and card 2 contains the IO queue location for which debugging is to occur. The queue location is inputted as a decimal value.

8.6.2.4 Random Access File Debug (Action Code DEBUG). This option will allow debugging to occur whenever the random file containing excess histograms is read or written. One data card containing the word "DEBUG" is all that is required.

8.6.2.5 Channel Debug (Action Code DCHN). This option will allow debugging to occur for a given channel (ACTCHN). This is not input as an actual channel number, but rather as a relative channel number. Card 1 contains the word "DCHN", and card 2 contains the relative channel number for which debugging is to occur.

8.6.3 Removal of Queue Entries (Action Code DELTA). As explained in subsection 8.5.6.7, it is possible for a T22 trace to be generated and never followed by a corresponding T7-Connect Trace. As long as the T22 trace is active, it will be considered as an active I/O request and may effect the queue reports and histograms. If a connect trace has not occurred within a 5-second (default) time span, this T22 trace will automatically be removed from the queue tables and an entry will be made into the Device ID STIOS Not Connected Report. This option allows the user to alter the 30-second default value. Card 1 contains the word "DELTA", and card 2 contains the default time inputted in milliseconds.

8.6.4 Set a Timespan of Measurement (Action Code TIME). The timespan of data collection can cover many hours, of which only a few may be of interest.

This option allows a user to specify the timespan (or spans) to be displayed in all reports. For example, the user may specify that he wants to collect data from 0500 to 2200 and wants to display data only from 0900 to 1700 in all reports. Up to five timespans may be specified, and they must be serially ordered. The format for this action is given in figure 8-24. All times are expressed as four character fields with no intervening blanks. Time is based on a 24-hour clock. If the user wants to request the time 4:07 he must input a "0407." If a start time but no stop time is desired, no characters should be entered after the minutes of the start time. If a stop time is requested, there must be a start time corresponding to it. Individual reports may not be timespanned. The timespans will affect the entire set of output.

8.6.5 Turn a Report On/Off (Action Code ON/OFF). These options allow the user to turn certain report types on or off. The report types effected by these options are the Channel Statistics Reports (subsection 8.5.6) and the Activity Statistics Reports (subsection 8.5.11). These reports are on by default and, therefore, the user should need only to use the OFF option if these reports are not desired. The ON option is supplied only for purposes of completeness. The format for this option consists of the word ON or OFF on card 1 and the word STANDARD or STATISTICS on card 2. The word STANDARD is used to turn off the Channel Statistics Reports and the word STATISTICS is used to turn off the Activity Statistics Reports.

8.6.6 Continue Data Reduction After an Input Option Error (Action Code ERROR). This option allows data reduction to continue when an error has been detected and reported in an input option request. The default value reports the error and aborts the data reduction procedures. The format for this option is the word ERROR on the data card.

8.6.7 W6.4/2H Data Reduction (Action Code RN). This option allows a user to process a GMF data tape (W6.4/2H or W7.2/4JX) under a W6.4/2H software release. It consists of the word RN on a data card.

8.6.8 Record Limitation by Connects (Action Code NCONN). This option allows a user to process only a specific number of connects. This option is especially useful if the tape contains an error on it and cannot be completely processed. Using this option, the user can process the tape or tapes up to the point where the tape error exists. This option requires two data cards. The first data card contains the word NCONN with the second card containing the number of connects to be processed.

8.6.9 Record Limitation by Records (Action Code NREC). This option allows a user to process only a specific number of tape records. This option is especially useful if the tape contains an error on it and cannot be

```
Card  1  A
Card  2  N M
Card  3  B C D E F ...
```

where

A = the word TIME
N = the word TOTAL
M = Number of different times appearing on card 3
B,C,D,E,F = Start and stop times used to define timespans.
Times must be separated by one or more blanks.

Figure 8-24. Input Option TIME Format

completely processed. Using this option, the user can process the tape or tapes up to the point where the tape error exists. This option requires two data cards. The first data card contains the word NREC with the second card containing the number of tape records to be processed.

8.6.10 Special Job Processing Reports (Action Code JOB). The Special Job Processing Report described in subsections 8.5.13 and 8.5.14 can be obtained with this option. These reports will not be produced unless this option is invoked. The format consists of three data cards. Card 1 contains the word JOB, card 2 contains the number of special jobs to be reported (not to exceed 8), and card 3 contains the actual SNUMBs, separated by at least one blank column.

8.6.11 Change the Time Quantum Value for the Special Job Processing Report Per 10 Minutes (Action Code RATE). The user can change the time quantum value used to produce the Special Job Processing Report Per 10 Minute by inputting a new time quantum in seconds. Two cards are required. The first card contains the word RATE and the second card contains the new quantum in minutes. The default value is 10 minutes.

8.6.12 END Card (Action Code END). This card must be present at all times and must be the last data card supplied. It consists of the word END typed on the card.

8.6.13 Limit the Processing and Output (Action Code LIMITS). This option will allow the user to control the amount of output produced and the amount of record processing performed. Card number 1 contains the word LIMITS and card number 2 contains either the word ONLYSP, the word NOHIST, or the word SUMARY. If the word ONLYSP is used then the Channel Monitor program will process only those data records that are generated by the SNUMBs requested under the JOB input option (see subsection 8.6.10). All other data will be ignored. The user must take care when examining the histograms and reports that are produced. The user must remember that only a limited amount of data has been processed. If the word NOHIST is used then no histograms will be produced. This option can be used in conjunction with the ONLYSP option (must have two LIMITS input cards) or can be used by itself. In the latter case, all data will be analyzed, but no histograms will be produced. Finally, if the word SUMARY is used then the user will not receive any histograms, but he will receive a single line of print for each device which provides the same information that occurs on the summary line of each histogram (last two lines of a histogram). The SUMARY option can be used in combination with either of the other two options (i.e., the user can turn off histograms and only receive the summary, or he can receive both the summary and the histograms). A sample of the summary reports is not provided in this document.

8.6.14 Zero Activity Processing (Action Code ZERO). This option allows the user to define up to 10 jobs which the user desires to see handled as normal activities, even though they process with an activity number of

zero. Under normal conditions, any activity processing with an activity number of zero will be considered a System Scheduler Job (see subsection 8.5.11). To use this option, the first data card should have the word ZERO. The second data card contains the number of jobs following on the third card. This number may not exceed 10. The third card contains the list of SNUMBs, separated by at least one blank column.

8.7 JCL

The data reduction procedures consist of a single FORTRAN program having a main level and multiple subroutines. A description of the more important JCL cards is presented below (see figure 8-25).

The $ LIMITS card should be changed to meet the user needs. The run time (99) and output limit (30K) may both need to be altered as required by the duration of the monitoring run. The CMDRP requires 48K of memory in order to execute plus an additional 2K for SSA space. During the initial loading process, CMDRP will actually require 60K of memory, but 10K will be released immediately upon loading.

The statement

```
$ DATA I*
```

is used to identify the data cards that follow as described in subsection 8.6. At least one data card is required, that being an "END" request.

8.8 Multireel Processing

If more than a single reel of data has been collected, a series of messages will be outputted to the console informing the operator that a new data reel is required. The following are the messages produced.

a. DISMOUNT REEL #XXXXX THEN MOUNT REEL NUMBER YYYYY ON ZZZZZ

   In this case, XXXXX is the old reel number, YYYYY is the new reel number, and ZZZZZ is the tape drive ID.

   If the operator fails to mount the new tape, the above message will be repeated three times, after which the program will terminate, and all reports produced.

b. IS TAPE XXXXX MOUNTED ON DRIVE ID YYYYY (Y/N)

   In this case, XXXXX is the tape number being requested for mounting and YYYYY is the tape drive ID.

   This message occurs when the data reduction program finds the wrong tape has been mounted (by comparing internally generated

```
COL   1        8          16

      $        IDENT      1820251/30/3044
      $        SELECT     B29IDPX0/OBJECT/CM
      $        TAPE       01,X1D,,12345
      $        LIMITS     99,50K,-4K,30K
      $        DATA       I*
      $        Data Cards - at least an "END" card must be present
      $        ENDJOB
```

Figure 8-25. CMDRP JCL

tape labels). If the operator answers N, the message in (c) below is produced. If the operator answers Y, the data reduction program will terminate and all reports will be produced. In this case, the data reduction program is unable to process the tape. Even though the operator is mounting the correct tape, the internal label on the new tape does not match that being requested by the old tape. The user should check the data collection session to insure that the operator did not respond with an incorrect tape number during multireel change.

After entering the Y or N, the operator will need to hit the EOM key twice in order for the response to be transmitted.

c. WRONG REEL JUST MOUNTED, DISMOUNT AND MOUNT REEL XXXXX ON ZZZZZ

In this case, XXXXX is the new reel number, and ZZZZZ is the tape drive ID.

d. CAN TAPE XXXXX BE MOUNTED ON DRIVE YYYYY (Y/N)

In this case, XXXXX is the new desired reel and YYYYY is the tape drive ID.

If the operator fails to answer this message it will be repeated until he responds with a "Y" for YES or "N" for NO. If he types in "Y", then message (a) will be repeated. If he types in "N", then the program will be terminated and all reports will be produced.

8.9 Tape Error Aborts

During the course of processing it is possible that the operator will be required to abort the data reduction program due to an irrecoverable tape error. If such a condition occurs, the operator should abort the job with a "U" abort. This will allow the data reduction program to enter its wrap-up code processing and produce all reports generated prior to the tape error.

SECTION 9. COMMUNICATIONS ANALYSIS MONITOR DATA REDUCTION PROGRAM (CAMDRP).

9.1 Introduction

The Communications Analysis Monitor Data Reduction Program is a FORTRAN program that sequentially processes data the Communications Analysis Monitor collected and wrote on tape. CAMDRP produces a number of reports depicting the usage of terminals, the response being received by terminals and the various DAC subsystems, and a special analysis report on Time Sharing Response. Report descriptions are presented in subsection 9.5.

There are two inputs to the CAMDRP. The first is the data tape produced by the CAM in the General Monitor Collector. The second is a set of report option control cards used to alter the reports in some way other than the standard default. The various user input options and their formats are described in subsection 9.6. The actual reports produced by CMDRP are described in subsection 9.5.

9.2 Data Collection Methodology

The CAM in the General Monitor Collector processes a GMF generated trace type 14 and collects information to monitor the usage of the entire terminal and DAC subsystems. The information collected on the occurrence of the above trace enables the CAMDRP to identify the DAC Subsystem Activity, response time being received by both DAC subsystems and terminals, and the extent to which any terminal is being utilized. The method used for generating the CAM traces is described in subsection 5.2.6 and the formats for the CAM generated records used by the CAMDRP are described in subsection 5.4.7.

9.3 Analytical Methodology

All communications between the H6000 and an online user is controlled by the GCOS module .MDNET (.DNWW in W6.4). This module contains a series of buffers, called mailboxes, that are used to store data passing between the datanet and the H6000. Whenever either machine wants to communicate with the other, information is placed in a mailbox and an interrupt is generated. The Communications Analysis Monitor (CAM) is designed to examine the mailboxes each time they are changed and to generate a GMF trace type 14. The trace type 14 is used by CAMDRP to provide data transfer rates, machine response times and user think times. The data transfer rates are derived from the number of words transferred for each interaction. Machine response time can have multiple definitions. One definition is the amount of time from the transfer of the first character of data by the user (carriage return) to the first response back to the user from the system. This definition is not precise in that TSS transmits a CLEAR SCREEN ACK back to the terminal prior to actual data transmission. Several seconds may pass before the user receives any usable data at his screen.

A second definition of response time is the same start time (carriage return) but the stop time is when the user has received his last piece of data before being required to give another response. In the case of multiscreen output, response ends when the user is requested to input a screen clear. This definition also is not precise in that the system response is not complete until possibly a full screen of data has been transmitted. This definition also lumps GCOS and subsystem (TSS, TRAX) response together. However, it is felt that this method is a more realistic method of response time calculation, and is the method used by CAMDRP.

NOTE: Users monitoring a job execution will be credited with a large response time, usually equal to the job execution time since the status monitoring subsystem requires no input until the job is finished. These users will be reported in the out-of-range response time average.

User think time is defined by CAM to be the amount of time from the start of data transmission to the user until the receipt of the first character of user response. This includes any datanet delay time (monitored by the datanet monitor) and any user wasted time (coffee break, phone call, etc.). However, this is the best definition available with the type of data captured by the CAM. Figure 9-1 presents a pictorial representation of these definitions.

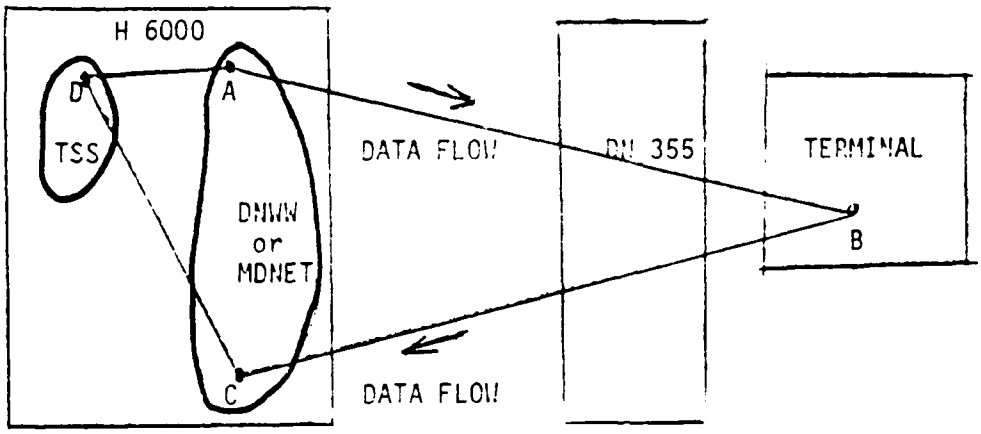
9.4 Data Reduction Methods

The CAMDRP reads only the trace type 14 records and any special records. It ignores lost data records, which can cause loss of some logons and disconnects. The CAMDRP logs a user onto a subsystem only when "connect to slave" command is captured. This command gives the actual subsystem the user is connecting to (TSS, TRAX, etc.). If, when CAMDRP first begins processing, a user is found to already be logged onto the system and no "connect to slave" command has been found, the user is logged onto an "UNKNWN" subsystem. This is because the "connect to slave" command is the only time the actual subsystem name is known. However, for every terminal logged on to TSS, each time a response is generated, the USERID of the terminal user is collected. CAMDRP uses this information to log the user on to TSS.

If, during processing, a datanet is found to have crashed, all users connected to that datanet are disconnected by the CAMDRP and processed as an end terminal session. If a reduction time frame ends, all users are disconnected as if their terminal session ended, and all reports are printed.

9.5 CAMDRP Output

The CAMDRP produces a header page and either a 355 Mailbox Report or Statistical Summary Reports, Terminal Session Reports, and requested histograms. The following subsections will describe all the reports produced by the CAMDRP.



MACHINE RESPONSE TIME
TIME FROM C TO D TO A

USER THINK TIME
TIME FROM A TO B TO C

Figure 9-1. Communications Analysis Concept

9.5.1 Header Page. The header page contains the total number of records read, trace records read, communications analysis (type 14) trace entries read, and the types and number of communication interface commands encountered. The date and time of data analyzed are printed, along with all the data tape numbers processed. An example of the header page is shown in figure 9-2. The version number printed on this report should say 01-82.

9.5.2 Trace Dumps. These reports are produced only when specified by the LIST, ALL, or MAIL input options. These dumps include a 355 Mailbox Trace Dump (Total - option LIST, or Selective - option ALL) and a Terminal Mailbox Dump (option MAIL). These options are described in subsections 9.6.4 and 9.6.8.

9.5.2.1 355 Mailbox Report-Trace Dump. This report is produced only if the user requested input option LIST or ALL. The Mailbox Report provides a formatted listing of the individual communication mailboxes collected by the data acquisition program (CAM). A sample of this output is given in figure 9-3. The report contains a single mailbox per line, separated into 13 fields with column headings for each field appearing at the top of each new page. Column headings and their meanings are as follows:

a. ID - Specifies the unique terminal ID.

b. CMND - Specifies the category of command from one machine to the other. This heading is further explained in the paragraph following this list.

c. OP - Specifies the detailed OP code within a command category requested of the other machine.

d. P - Gives the DN355 processor number. Numbering starts with 0.

e. LINE - Gives the logical line number which is used to reference line tables in the DN355 and H6000.

f. TYPE - Shows the terminal type.

g. TLY - Tally word in H6000 word units. On an output data transfer, it specifies the number of words transferred from the H6000 to the DN355 (four characters per word in DAC mode and six characters per word in remote batch mode). On an input data transfer, it specifies the maximum number of words the H6000 allows the DN355 to transfer.

h. ITLY - Input tally. This specifies the number of words actually input into the H6000 on an input data transfer. It will always be less than or equal to the TLY.

COMMUNICATIONS MONITOR VERSION 01-82 FOR NMCC2 ON 122181

- SUMMARY OF TRACE ENTRIES READ -

# OF RECORDS READ:		293
TOTAL TRACE ENTRIES READ:		119232
TOTAL TRACE 14 ENTRIES READ:		119231
COMMANDS:		
	RCD	49323
	RTX	20412
	WCD	978
	WTX	48427
	ACU	0
	RJM	0
	RSVP	91
	ECHO	0

MONITOR OVERHEAD - MONITOR	TIME(SEC)	% OVERHEAD
EXEC	600	2
CAM	300	1
TOTAL		3

START TIME(HH:MM:SS.SSS)
11:54: 0.000
STOP TIME(HH:MM:SS.SSS)
16:03: 34.749
FOR A TOTAL OF 4.15 HOURS

MONITOR TAPE USED - C010

Figure 9-2. CAMDRP Header Page

COMMUNICATIONS MONITOR TRACE DUMP FOR NMCC2 ON 122181

ID	CMND	OP	P	LINE	TYPE	TLY	ITLY	SI	CMI	ICM	TOD	RTOD	
CY	RCD	ACPDIP	U	00032	VP786W	0	0	00	0	000000	045737764021	22: 6:	0.608
CY	RTX	INPACC	0	00032	VP786W	519	13	00	0	000000	045737765114	22: 6:	0.617
CY	WTX	DACOUT	0	00032	VP786W	5	0	00	0	000000	045737770142	22: 6:	0.641
CY	RCD	SNDOTP	0	00032	VP786W	0	0	00	0	000000	045740060445	22: 6:	1.092
KP	WTX	DACOUT	1	00024	VP7705	21	0	00	0	000000	045740212305	22: 6:	1.811
CY	WTX	DACOUT	0	00032	VP786W	14	0	00	0	000000	045740220154	22: 6:	1.857
UC	RCD	SNDOTP	0	00025	TTY	0	0	00	0	000000	045740411615	22: 6:	2.830
KS	RCD	SNDOTP	1	00036	VP786W	0	0	00	0	000000	045740723005	22: 6:	4.440
CS	RCD	SNDOTP	0	00026	VP7705	0	0	00	0	000000	045740723571	22: 6:	4.445
CS	WTX	DACOUT	0	00026	VP7705	440	0	00	0	000000	045741726256	22: 6:	8.562
KS	WTX	DACOUT	1	00036	VP786W	440	0	00	0	000000	045741726372	22: 6:	8.563
CY	RCD	ACPDIP	0	00032	VP786W	0	0	00	0	000000	045742114324	22: 6:	9.507
CY	RTX	INPACC	0	00032	VP786W	519	25	00	0	000000	045742115447	22: 6:	9.516
RA	RCD	BREAKC	0	00023	RCT	0	0	00	0	000000	045742532522	22: 6:	11.669

Figure 9-3. Mailbox Report

i. ST - Status of the operation.

j. CNT - Count of the number of meaningful characters included in the next field (ICM); a decimal integer.

k. ICM - Intercomputer message field. This is a word in the mailbox for storing extra data to be passed to the other machine. The only information the Communications Monitor looks at in this word is the name of the slave (e.g., TSS) to which a terminal is attempting to connect on a "connect to slave" operation code.

l. TOD - Elapsed time since the system was booted in a 12-digit octal format with 1/64 millisecond resolution.

m. RTOD - Real time-of-day; the wall clock time in the format HH:MM:SS, with one millisecond resolution. It is computed by adding the wall clock starting time (MME GETIME) of the data acquisition program to the elapsed time from the first mailbox to the current mailbox.

The command field specifies the category of request between the two machines. The four most important commands are as follows:

- RCD - READ command data
- WCD - WRITE command data
- RTX - READ text
- WTX - WRITE text.

RCD specifies a request from the DN355 to the H6000, whereas WCD is a request from the H6000 to the DN355. TRX and WTX request actual data transfers with RTX for input and WTX for output.

If a DAC program wants to communicate with a terminal connected to it, the program issues either an "output" request to the system or an "output wait for input" request. When the "output" request (MME GEROUT 03) is issued, the program supplies the data to be sent to the terminal and the program expects notification when the operation is completed. First, a "WTX DACOUT" mailbox will appear at this interface which causes the output message to be transferred from the H6000 to the DN355. After the DN355 transmits the data to the terminal, an "RCD SNDOTP" mailbox appears which informs the H6000 that the output is completed and more output may now be sent. If this program sends more outputs, this mailbox sequence will appear repeatedly.

The mailbox sequence is different when an "output wait for input" (MME GEROUT 04) is issued by the program. The program sends output and waits for the user to enter his input. The program may be swapped out while waiting for the input. First, a "WTX DACOUT" mailbox appears with the same significance as before. No mailbox appears after the output has been completely transmitted to the terminal. Only after the user has entered his line and has depressed the carriage return does an "RCD ACPDIP" mailbox appear at the interface. This action initiates a request for the H6000 to "accept DAC input" from the DN355. The DN355 does not directly transfer the data since it is not known to which address the data should be transferred or even whether the program is in H6000 memory. Once GCOS determines the address to which the data should be transferred (swapping the program in, if necessary), a third mailbox, "RTX INPACC," appears at the interface informing the DN355 that the H6000 will "accept input" and where in the H6000 to move it. This three mailbox sequence is characteristic of the "output wait for input" operation.

In remote batch mode, a continuous stream of either input or output is being transferred. If it is output, the sequence which will be repeated will be "WTX ACCOUT" followed by "RCD SNDOTP." The first mailbox informs the DN355 to transfer batch output from the H6000. The second is the DN355 requesting more output.

On remote batch input, the repetitive sequence is "RCD ACPINP" followed by "RTX INPACC." The first is the DN355 request to the H6000 to accept batch input, and the second is the H6000 accepting it and that the transfer from DN355 to H6000 should be made.

The formatted device names used by both the formatted Mailbox Report and the Statistical Summary Reports correspond to a two-digit octal field appearing in the mailbox. The octal codes and the corresponding device names used in the program are shown in table 9-1.

9.5.2.2 Terminal Mailbox Dump. This report is produced only if the MAIL option is requested along with either the LIST or ALL option. Also, the CAM must have run with the complete communications data option (monitoring of specific terminals). This report gives a dump, in ASCII, Octal, BCD, or all three of all communication traffic between the H6000 and the DN355 for the specific terminals that were monitored. This data will include all SOH (start of header), CR (carriage return), LF (line feeds), etc. For terminals with upper/lower case data, upper case data will be followed by a period (i.e., T.H.E is THE and THE is the). This option allows a user to track exactly what a terminal user is doing. CAUTION - This data may be classified since passwords may be printed. A sample figure is not presented in this document.

Table 9-1. Octal Codes and Device Names for the Mailbox and Statistical Summary Reports (Part 1 of 2)

16	355	Direct DN355 Link
17	XXXX17	Reserved for the System
20	2741	2741 Teletypewriter
21	XXXX21	Reserved for the System
22	XXXX22	Reserved for the System
23	RLP300	Remote Line Printer
24	XXXX24	Reserved
25	MSLINK	Mass Store Link
26	XXXX26	Reserved for the System
27	XXXX27	Reserved for the System
30	MRS200	MRS200 Document Handler
31	DRD200	DRD200 Document Handler
32	DRDDHU	DRD236 or DHJ1604/8/12/16 Document Handler
33-47		Reserved for the System
50-67		Reserved for Test and Diagnostics
61-77		Reserved for Customer Use
01	XXXX01	Reserved for the System
02	XXXX02	Reserved for the System
03	RCT	Remote Computer Terminal
04	TTY	Teleprinter
05	DNET4	DATANET 760 - Screen Size = 4 Lines X 46 Characters

Table 9-1. (Part 2 of 2)

06	DNET8	DATANET 760 - Screen Size = 8 Lines X 46 Characters
07	DNET16	DATANET 760 - Screen Size = 16 Lines X 46 Characters
10	DNET26	DATANET 760 - Screen Size = 26 Lines X 46 Characters
11	VIP775	VIP Series 755/775 - Screen Size = 22 Lines X 92 Characters
	VP786W	VIP Series 786 (WWMCLS) - Screen Size = 22 Lines X 92 Characters
13	VIP 12	VIP Series 7700 - Screen Size = 12 Lines X 80 Characters
14	VIP22	VIP Series 7700 - Screen Size = 22 Lines X 46 Characters
15	VP7705	VIP Series 7700 - Screen Size = 24 Lines X 80 Characters

9.5.3 Statistical Summary Reports. These reports are produced unless the user specifically requests the 355 Mailbox Report. The Statistical Summary Reports include DAC Device and Subsystem Summaries, Remote Batch Device Summary, and Terminal ID Summary.

9.5.3.1 DAC Devices Summary Report. The DAC Device Summary Report (shown in figure 9-4) indicates the activity for each of the different DAC terminal types utilized during the data collection session. The distinctive device types include TTYs, IBM 2741, and several categories of displays (e.g., VIP786W and 7705). Also represented in these reports are devices which use DAC protocols, such as the DN355-DN355 link implemented between the CCTC and ANMCC.

For each DAC device, the following values are reported in terms of mean and standard deviation (where appropriate):

- Number of sessions collected - Number of terminal sessions accumulated in this category.
- Session length (sec) - Time from log-on to log-off
- Input length (char) - Number of characters in an input
- Output length (char) - Number of characters per output
- Number of outputs/output group - Number of distinct outputs in consecutive order (total output length = output length times the number of outputs/output group)
- User think time (sec) - Time from start of last output transfer in output group to end of input transfer
- Machine response time (sec) - Time from end of input transfer to request for next input transfer (see definition in section 9.3)
- Inter-output time (sec) - Time from start of previous output transfer to start of succeeding output transfer
- Character rate (char/sec) - Total number of characters transferred, divided by the terminal session length.
- Number of inputs - Total number of user responses during the time period.
- Average number of inputs - Average number of user responses per session (number of inputs/number of sessions collected)
- Number outputs - Total number of system responses during the time period

DAC DEVICE SUMMARY FOR NMCC2 ON 122181

	VP7705		TTY	
	MEAN	STD DEV	MEAN	STD DEV
# SESSIONS COMPLETED	39.		2.	
SESSION LENGTH (SEC)	1044.9	1415.4	2460.0	2037.9
INPUT LENGTH (CHAR)	22.6	53.4	4.4	1.5
OUTPUT LENGTH (CHAR)	245.7	241.7	44.6	13.0
# OUTPUTS PER OUTPUT GROUP	1.9	38.8	1.0	0.2
USER THINK TIME (SEC)	17.8	35.4	6.4	7.5
MACHINE RESPONSE TIME (SEC)	3.2	17.6	0.3	1.1
INTER-OUTPUT TIME (SEC)	1.7	1.4	1.0	0.
CHARACTER RATE (CHAR/SEC)	38.6	39.1	0.5	0.7
# INPUTS	975.0		25.0	
AVG # INPUTS	25.0		12.5	
# OUTPUTS	5392.0		50.0	
+ AVG # OUTPUTS	138.3		25.0	
FLAG	0		0	

FLAGS:

0 - NORMAL LOG-ON, DATA COLLECTION, TERMINATION.
2 - DUPLICATE ID ENCOUNTERED DURING LOG-ON. OLD ID LOGGED-OFF.
5 - TERMINAL LOGGED-ON VIA SUBROUTINE INSERT. NONSTANDARD PROCEDURE.
THIS FLAG HAS BEEN DELETED FOR NON-BATCH TERMINALS.

Figure 9-4. DAC Device Summary Report

- Average number of outputs - Average number of system responses per session (number of outputs/number of sessions collected

- Flag - Logical flags indicating any unusual conditions in this category. Flag explanations appear on the printout.

The user should note that summaries for the same batch device with different flags are not mixed. Thus, for example, there may be several summaries for RLP 3000 with the majority of normal sessions reflected in one summary and exceptions in the others. The exceptions imply conditions such as GMC starting its collection in the middle of some terminal session for which the session length cannot be determined.

9.5.3.2 DAC Subsystem Summary Report. The DAC Subsystem Summary Report (figure 9-5) summarizes the characteristics of users of DAC subsystems such as TSS and TPS. The heading of each report gives the subsystem utilized. The categories summarized are the same as those categories discussed in subsection 9.5.3.1.

Invariably, a number of the subsystems summarized are bogus due to user typing errors. For example, the following misspellings of TSS may appear: 'YSS', and 'TS'. These reports do not imply that these subsystems exist, only that some user attempted to log-on to a system with that name. Users logged onto a subsystem other than TSS and prior to the CAM starting will be considered as logged onto subsystem UNKNWN (see subsection 9.4).

9.5.3.3 Remote Batch Device Summary Report. The Remote Batch Device Summary Report profiles the devices using remote batch mode communication protocols such as remote line printers (RLP300) and remote computers (RCT) (figure 9-6).

For each device, the following values are reported:

- Number of jobs collected - Total number of distinct jobs reflected in this report.

- Number of input jobs - That part of the total number of jobs which are input jobs.

- Number of output jobs - That part of the total number of jobs which are output jobs. Certain RCT (H716) reports contain jobs that may be counted as both input and output; more detailed examination of the raw data is required to verify this circumstance.

- Job length (sec) - Time from the first to the last data transfer for that job.

DAC SUBSYSTEM SUMMARY FOR NMCC2 ON 122181

	TPE		TSS		UNKNWN	
	MEAN	STD DEV	MEAN	STD DEV	MEAN	STD DEV
# SESSIONS COLLECTED	1.		27.		4.	
SESSION LENGTH (SEC)	104.0	0.	1345.2	1493.2	1003.8	1612.4
INPUT LENGTH (CHAR)	8.0	7.3	23.3	54.5	0.	0.
OUTPUT LENGTH (CHAR)	35.5	62.1	89.5	222.0	438.1	69.6
# OUTPUTS PER OUTPUT GROUP	1.0	0.2	1.0	0.1	610.8	979.0
USER THINK TIME (SEC)	4.6	4.0	18.3	36.0	0.	0.
MACHINE RESPONSE TIME (SEC)	4.0	9.3	3.3	17.9	0.	0.
OUTPUT TIME (SEC)	0.	0.	0.0	0.2	1.7	1.3
CHARACTER RATE (CHAR/SEC)	16.0	0.	11.8	12.0	281.5	325.0
# INPUTS	12.0		934.0		0.	
AVG # INPUTS	12.0		32.6		0.	
# OUTPUTS	42.0		2819.0		2443.0	
AVG # OUTPUTS	42.0		104.4		610.8	

Figure 9-5. DAC Subsystem Summary Report

REMOTE BATCH DEVICE SUMMARY

	RLP300		RLP300		RCT	
	MEAN	STD DEV	MEAN	STD DEV	MEAN	STD DEV
# JOBS COLLECTED	1.		16.		2.	
# INPUT JOBS	0.		0.		2.	
# OUTPUT JOBS	1.		16.		2.	
JOB LENGTH (SEC)	16.0	0.	250.9	281.1	6210.0	2.8
INPUT LENGTH (CHAR)	0.	0.	0.	0.	8.9	58.7
OUTPUT LENGTH (CHAR)	747.4	637.5	1097.5	747.4	11.3	2.7
INPUT CHARACTER RATE (CHAR/SEC)	0.	0.	0.	0.	15.5	6.4
OUTPUT CHARACTER RATE (CHAR/SEC)	327.0	0.	218.4	67.2	0.	0.
FLAG*1*	5		0		5	

1 FLAGS:

0 - NORMAL LOG-ON, DATA COLLECTION, TERMINATION.
2 - DUPLICATE TO ENCOUNTERED DURING LOG-ON. OLD ID LOGGED-OFF.
5 - TERMINAL LOGGED-ON VIA SUBROUTINE INSERT. NONSTANDARD PROCEDURE.

Figure 9-6. Remote Batch Device Summary Report

- Input length (char) - Number of input characters in distinct input.

- Output length (char) - Number of output characters in distinct output.

- Input character rate (char/sec) - Total number of input characters divided by the job length.

- Output character rate (char/sec) - Total number of output characters divided by the job length.

- Flag - Logical flags denoting special conditions. Flag values are explained on the report printout. Reports for the same device, with differing flag values, are maintained separately.

9.5.3.4 Terminal ID Summary Report. The Terminal ID Report (figure 9-7) summarizes the activity on a particular terminal over the entire data collection session. Terminal IDs are unique and unchanging for each terminal configured on a given system. These terminal ID assignments are made by the computer operations personnel and become effective at system boot time. At some sites with multiprocessors, the terminal ID assignments may be different, depending on whether the multiple systems are split. Terminal IDs are two alphanumeric characters which correspond to the 12-bit identification, specified in calls to the system, used to identify the terminal in which communications are directed. Each terminal ID report is headed by this 2-character ID and the device type of that terminal. Items reported are as follows:

- Number of sessions collected - Total number of terminal sessions on that terminal.

- Session (job) length (sec) - Length of terminal session or job.

- Input length (char) - Length of distinct inputs. Also appended is the conversion factor, from words to characters, obtained from the mailbox giving the length in units of H6000 words. DAC uses ASCII characters at four characters per word and remote batch uses BCI characters at six characters per word.

- Output length (char) - Length of distinct outputs.

- Batch input character rate (char/sec) - Total number of input characters in job divided by job length.

- Batch output character rate (char/sec) - Total number of output characters in job divided by job length.

TERMINAL ID SUMMARY FOR MMCC2 ON 122181

	RC	RCT	D7	VP7705	D1	VP7705
	MEAN	STD DEV	MEAN	STD DEV	MEAN	STD DEV
SESSIONS OR JOBS COLLECTED	3.		4.		4.	
SESSION (JOB) LENGTH (SEC)	94.7	65.7	979.8	1478.5	1070.3	1556.2
INPUT LENGTH (CHAR)*1*	93.9	279.7	24.4	39.4	42.9	78.4
OUTPUT LENGTH (CHAR)*1*	3.9	7.6	50.1	80.7	59.0	88.9
BATCH INPUT CHAR RATE (CHAR/SEC)	8.7	11.6	0.	0.	0.	0.
BATCH OUTPUT CHAR RATE (CHAR/SEC)	0.	0.	0.	0.	0.	0.
DAC CHARACTER RATE (CHAR/SEC)	0.	0.	15.5	17.9	13.3	15.3
USER THINK TIME (SEC)	9.0	20.5	22.1	28.3	31.4	52.5
MACHINE RESPONSE TIME (SEC)	0.	0.	2.5	4.4	1.7	2.4
INTER-OUTPUT TIME (SEC)	0.	0.	0.	0.	0.	0.
% TERMINAL USAGE	5.582		77.027		84.1	
# INPUTS	20.0		157.0		128.0	
AVG # INPUTS	6.7		39.3		32.0	
# OUTPUTS	23.0		499.0		378.0	
AVG # OUTPUTS	7.7		124.8		94.5	
FLAG	0		0		0	

1 WORD-TO-CHARACTER CONVERSION DIFFERS FOR DAC AND REMOTE BATCH-
DAC : CHARACTERS = # WORDS * 4 - 2
BATCH CHARACTERS = # WORDS * 6

2 FLAGS:
0 - NORMAL LOG-ON, DATA COLLECTION, TERMINATION.
2 - DUPLICATE ID ENCOUNTERED DURING LOG-ON. OLD ID LOGGED-OFF.
5 - TERMINAL LOGGED-ON VIA SUBROUTINE INSERT. NON-STANDARD PROCEDURE.
THIS FLAG HAS BEEN DELETED FOR NON-BATCH TERMINALS.

Figure 9-7. Terminal ID Summary Report

- DAC character rate (char/sec) - Total number of DAC characters divided by terminal session length.
- User think time (sec) - Time from start of last output transfer in output group to end of input transfer.
- Machine response time (sec) - Time from end of input transfer to request for next input transfer. See definition in section 9.3.
- Inter-output time (sec) - Time from start of previous output transfer to start of succeeding output transfer.
- % terminal usage - The percent of the monitored time this terminal ID was active. Calculated by summing up the total terminal connect time and dividing by the total session length.
- Number inputs - Total number of user responses during the time period
- Average number of inputs - Average number of user responses per session (number of inputs/number of sessions collected)
- Number of outputs - Total number of system responses during the time period
- Average number of outputs - Average number of system responses per session (number of outputs/number sessions collected
- Flag - Logical flags explained in the printout

9.5.4 Delta Time Period Summary Report. This report (figure 9-8) is used to monitor overall communication activity on the system as a function of time. It shows the total number of characters input and output, in DAC and remote batch modes, during consecutive time periods. This report is produced when the time period to be used is specified using the DELTA input option (subsection 9.6.2).

Up to four DN355s can be configured. In multi-DN355 environments, this report can show whether the loads on the two systems are balanced. These reports also show the number of different terminals that were active during each time period.

The times, shown in the left hand column, are 24-hour wall clock times. Interspersed in these reports are messages reflecting the aborting or rebooting of either DN355.

9.5.5 Histogram Reports. These reports are produced only if the user requests them with the HISTG input option (subsection 9.6.3). Three

DELTA TIME PERIOD SUMMARY FOR NMCC2 ON 122181

	DATANET-355 C					
	DIRECT ACCESS			REMOTE BATCH		
TIME	#ACT TRMS	INPUT CHARS	OUTPUT CHARS	#ACT TRMS	INPUT CHARS	OUTPUT CHARS
12:27:48	2	58	9070	0	0	0
12:28:18	2	26	9278	0	0	0
12:28:48	2	50	9782	0	0	0
12:29:18	2	14	8786	0	0	0
12:29:48	2	10	8788	0	0	0
12:30:18	1	0	7022	0	0	0
12:30:48	1	0	8788	0	0	0
12:31:18	2	26	10766	0	0	0
12:31:48	2	0	9398	0	0	0
12:32:18	1	0	8788	0	0	0
12:32:48	2	2	7026	0	0	0
12:33:18	2	18	9130	0	0	0
12:33:48	2	86	9106	0	0	0
12:34:18	2	14	10310	0	0	0
12:34:48	2	10	10002	0	0	0

Figure 9-8. Delta Time Period Summary Report

DISTRIBUTION COLLECTED ON SYSTEM NMCC2 AT 11:54:15 ON 81-12-21

MACHINE RESPONSE TIME FOR SUBSYSTEM TSS

INDIV. NUMBER	CUMUL. NUMBER	CUMUL. PRC	INDIV. PRC	SECOND		PERCENT OF OCCURRENCE 00 05 10 15 20 25 30 35 40 45 50	REPORT 1
						I....I....I....I....I....I....I....I....I....I....I....I	
2421	2421	31.00	31.00	0-	0	Ixxxxxxxxxxxxxxxxxxxxxxxxxxxxxx	
3531	5952	76.20	45.20	0-	3	Ixxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx	
945	6897	88.30	12.10	4-	6	Ixxxxxxxxxxxx	
331	7228	92.50	4.20	7-	9	Ixxxx	
156	7384	94.50	2.00	10-	12	Ixx	
85	7469	95.60	1.10	13-	15	Ix	
69	7536	96.50	0.90	16-	18	Ix	
47	7585	97.10	0.60	19-	21	Ix	
46	7631	97.70	0.60	22-	24	Ix	
27	7658	98.00	0.30	25-	27	I	
.	.	.	.	.	.	.	
.	.	.	.	.	.	.	
.	.	.	.	.	.	.	

7810 ENTRIES TOTAL AVERAGE = 4.12228 VARIANCE = 335.431 STANDARD DEVIATION = 18.315
190 (2%) OUT OF RANGE AVERAGE FOR THESE = 286.95 IN RANGE AVERAGE = 3.43

Figure 9-9. Machine Response Time Report

DISTRIBUTION COLLECTED ON SYSTEM NMCC2 AT 11:54:15 ON 81-12-21

SESSION LENGTHS FOR SUBSYSTEM TSS

INDIV. NUMBER	CUMUL. NUMBER	CUMUL. PRC	INDIV. PRC	MINUTE	PERCENT OF OCCURRENCE 00 05 10 15 20 25 30 35 40 45 50	REPORT 3
					I....I....I....I....I....I....I....I....I....I....I	
15	15	10.60	10.60	0- 0	Ixxxxxxxxxx	
43	58	40.90	30.30	0- 5	Ixxxxxxxxxxxxxxxxxxxxxxxxxxxx	
14	72	50.80	9.90	6- 10	Ixxxxxxxxx	
16	88	62.10	11.30	11- 15	Ixxxxxxxxxx	
8	96	67.70	5.60	16- 20	Ixxxxx	
14	110	77.60	9.90	21- 25	Ixxxxxxxxx	
.	.	.	.	. .	.	
.	.	.	.	. .	.	
.	.	.	.	. .	.	

142 ENTRIES TOTAL AVERAGE = 17.33099 VARIANCE = 433.799 STANDARD DEVIATION = 20.828

Figure 9-11. Session Length Report

experiences a response time greater than or equal to the requested limit. The information printed includes Terminal ID, subsystem name, response time in seconds, and time of day. Refer to figure 9-12.

9.5.7 User Think Time Limit Report. This report is produced only if the user requests it with the THINK input option (subsection 9.6.7). This report will print out a line of information every time a terminal experiences a response time greater than or equal to the requested limit. The information printed includes Terminal ID, subsystem name, think time in seconds, and time of day. Refer to figure 9-12.

9.5.8 Terminal Session and High Terminal Usage Reports. The Terminal Session Report (figure 9-13) is produced whenever the Statistical Summary Reports are requested. The report gives an account of every session that occurs during the monitoring session. Every time a user logs off or is logged off due to a DN355 abort, TCALL, or monitor termination, an entry into this report is produced. The report gives the Log On Time, Log Off Time, Terminal ID, USERID, Subsystem, Subsystem Size (MIN MAX), Session Length, Response Time, # Inputs, # Outputs. If a terminal was logged onto a subsystem when CAM started, then there is no immediate way for CAM to determine the subsystem being used by the terminal. In this case, the CAM data record will be checked to see if a USERID exists for the terminal. If one does exist, it means the terminal is logged on to TSS and the log on name is changed from UNKNWN to TSS. If there is no USERID for this terminal in the record, the subsystem name is set to UNKNWN. This occurs for the WIN lines and any user logged on to VIDEO when the monitor is started. If a user JDACS to a new subsystem, CAMDRP will disconnect the current line, calculate all statistics and reconnect the line to the new subsystem. Both the USERID and Subsystem Sizes are given only for TSS users. It has been noticed that some WIN pseudo terminals appear logged onto TSS, but have no USERID and no input/outputs. The cause for this discrepancy is under investigation. The Session Length is given in seconds. The Response Time is given in seconds and is the average response time over the session. The # Inputs is the number of input requests sent by the user. The # Outputs is the number of output response groups sent to the user. This report can help pinpoint excessive response times. It can also be used to determine if a terminal is logged onto the system and is not being used (low inputs, high or low outputs, long session length).

The High Terminal Usage Report is included as part of the Terminal Session Report and provides a list of terminals that have been logged on for a specified percent (default 75%) of the session. This will list terminals by ID and type, give the percent of time the terminal was logged on, the number of sessions during this time, the number of inputs and the number of outputs from the terminal. (See figure 9-14.)

9.5.9 Opcode Count Report. This report (figure 9-15) is produced whenever the System Summary Reports are produced. This report gives a listing of all the opcodes that were transmitted between the H6000 and the DN355, and a count of how many of each opcode there were. This report is of interest mainly when the following opcodes appear:

EXCESS THINK/RESPONSE TIME REPORT FOR NMCC2 ON 122181

TERMINAL C6 ON SUBSYSTEM TSS HAD A THINK TIME OF 208 SECS AT 16:35:55

TERMINAL CI ON SUBSYSTEM TSS HAD A THINK TIME OF 120 SECS AT 16:48:21

TERMINAL BT ON SUBSYSTEM TSS HAD A THINK TIME OF 223 SECS AT 16:48:57

TERMINAL BD ON SUBSYSTEM TSS HAD A THINK TIME OF 197 SECS AT 16:49:11

TERMINAL BT ON SUBSYSTEM TSS HAD A RESPONSE TIME OF 50 SECS AT 16:43:54

Figure 9-12. Response Time/User Think Time Limit Report

TERMINAL SESSION REPORT FOR NMCC2 ON 102182 AT 14:18: 0.000

LOG ON TIME	LOG OFF TIME	TERMINAL ID	USERID	SUBSYSTEM	SUBSYSTEM SIZE MIN	SUBSYSTEM SIZE MAX	SESSION LENGTH	RESPONSE TIME	NUMBER INPUTS	NUMBER OUTPUTS
18:53:17	18:54:37	CN	DJ3JI32411	TS1	1K	3K	80 SEC	4.500 SEC	4	8
15:01:25	18:54:37	4C		SENDC			13993 SEC	0. SEC	0	0
18:48:31	18:54:37	AA	DJ8XC344JG	TS1	1K	26K	367 SEC	3.324 SEC	39	125
18:03:07	18:54:37	CV	DJ3NI33203	TSS	0K	29K	3090 SEC	10.120 SEC	103	499
14:18:58	18:54:37	A8		UNKNWN			16539 SEC	0. SEC	0	4618
18:42:07	18:54:37	AP	DJ3JI32401	TELNET			751 SEC	8.844 SEC	32	69

Figure 9-13. Terminal Session Report

HIGH TERMINAL USAGE REPORT FOR NMCC ON 030382 AT 12:33: 0.000

TERMINAL ID	TERMINAL TYPE	PERCENT UTILIZATION	NUMBER SESSIONS	NUMBER INPUTS	NUMBER OUTPUTS
E3	VP7705	89.1	6.0	131	1448
D0	RLP300	99.1	2.0	0	198
01	XXXX24	96.9	12.0	51688	51692
02	XXXX24	96.8	12.0	0	28207
B7	VP7705	94.2	3.0	282	1156
BP	RCT	91.6	6.0	8687	8693
DN	VP7705	100.0	1.0	0	3250

Figure 9-14. High Terminal Usage Report for NMCC

OPCODES COUNT FOR NMCC2 ON 122181 AT 16:03: 34.749

TERMINAL ACCEPTED	164	DISCONNECT THIS LINE	117
DISCONNECT ALL LINES	2	ACCEPT CALLS	4
INPUT ACCEPTED	20412	ACCEPT OUTPUT	242
ACCEPT FINAL OUTPUT	12	OUTPUT NOT AVAILABLE	1
ACCEPT DIRECT OUTPUT	27614	ACCEPT DIRECT OUTPUT-INPUT	20501
REJECT REQUEST-TEMP	49	TERMINAL REJECTED	392
DISCONNECT ACK	26	6000 INITIALIZED	4
BREAK ACKNOWLEDGED	219	TERMINAL CHARACTERISTICS CHANGE	58
ACCEPT NEW TERMINAL	556	DISCONNECT LINE	26
SEND INITIAL OUTPUT	2	SEND OUTPUT	27841
CONNECT TO SLAVE	160	ACCEPT DIRECT INPUT	20461
BREAK	219	TERMINAL CHARACTERISTICS CHANGE ACK	58

***** END REACHED *****

Figure 9-15. Opcode Count Report

OPCODE	DESCRIPTION
11	Output not available
16	Reject request (temporary)
17	Reject request (permanent)
20	Terminal rejected
110	Backspace output

These opcodes indicate a delay in data transmission or a communications problem. If these opcodes show up consistently, and in significant numbers, a detailed analysis should be conducted.

9.5.10 Response Time Report. This report is produced whenever the user sets the interval time using the input option RATE (subsection 9.6.11) or SNUMB (subsection 9.6.12). The report shows, for each interval, the time of day, the response time for the requested subsystems, and the number of opcode rejects, RSVPs and RJMs. (See figure 9-16). The column headings are as follows:

TOD - Time Of Day

RESP - average response time over the time period

I/R - average response time of those responses considered acceptable (see sections 9.5.6 and 9.6.6)

#I - number of responses in the acceptable range

#O - number of responses in the unacceptable range. This number is important in validating the figure in the RESP columns. One extremely bad response can cause a skewed average response.

TRM - average number of terminals on this subsystem during the time frame

OPREJT - number of Opcode Reject Temporary commands received during the period

OPREJP - number of Opcode Reject Permanent commands received during the period

RJM - number of Reject Message commands received during the period

RSVP - number of Resend requests received during the period

NOTE: If TSS is one of the SNUMBs requested, all TS SNUMBs (TS1-TS4) will be represented under TSS.

9.5.11 Error Messages. The CAMDRP can produce multiple error messages relating to the data type. Most of these messages are actually warning messages, which the CAMDRP will try to recover from and will continue to process.

The most prevalant error message is the warning message "TERMINAL ID NOT FOUND." This message usually occurs when a terminal has been logged onto the system prior to the CAM starting to collect its data. When the CAMDRP tries to find a particular user who is receiving or transmitting data in its tables, that user will not be found since the CAMDRP did not find any log-on record for him. The user is logged onto the system and the CAMDRP continues processing.

The main reason for the CAMDRP to abnormally stop processing is the error message "NO MORE ROOM IN TERMID ARRAY." This means that an internal array has been filled. This is usually the terminal ID array. The parameter MAX must be increased to enlarge the required arrays. The current size of MAX is 50. This can be exceeded if there are a large number of users on the system when the CAM is started. To increase this value, the user should log onto TS1. Enter EDIT 0 B29IDPXO/SOURCE/CAM. Then enter:

```
CASE
VERI
RS:/MAXb=b50/;*:/MAXb=bxx/
```

where xx is the new maximum number of terminals to be allowed on the system. The CAM data reduction program will be increased by 175 words for each terminal above 50 allowed. All other messages produced will be self-explanatory. If they do not indicate a severe error, the words "For Information only" will appear with the message.

9.5.12 H6000-DN355 Reject Report. This report displays all the terminal IDs that have had some type of error opcode from or to the DN355. These opcodes are RJM, RSVP, ECHO, OPCRJT, OPCRJP. This report also lists all the terminals seen on all the datanets (see figure 9-16.1).

9.5.13 Abort/Disconnect Report. This report indicates each time the DN355 aborts and is reinitialized. It also presents each time specified line IDs (option DISC) disconnect. These disconnects indicate the line ID, the subsystem it was connected to, when it disconnected, how long it had been connected and the number of inputs/outputs associated with the line ID during the session. This report can be used to monitor how often the WIN lines lose their connections (see figure 9-16.2).

9.5.14 TS1 Initial Parameter Report. This report indicates the initial preset values for TS1. These values are SIZE parameters, LIMIT parameters, SWAP FILES parameters and a list of all ALLOCATED DEVICES per file code. This report is produced once, so if any parameters are changed during the run (such as TS1 max size), the change is not reported. See figure 9-16.3 for a sample of this report.

9.5.15 Mailbox Busy Report. This report is printed out each time a Response Time Report line is printed. This report indicates the time of day, the datanet, the total number of busy mailboxes during the time frame, the total number of times the mailboxes were looked at (number of responses), the total number of mailboxes available during that time frame (number of responses * seven mailboxes), the percent of the mailboxes busy, the number of mailbox denials, if any, the percent of the mailboxes available that were denied, the maximum mailboxes busy during the time frame, and the maximum number of denials seen at any one time (see figure 9-16.4).

RESPONSE TIME REPORT FOR NMCC ON 030382

	TSS						FTS						TLCF								
TOD	RESP	I/R	# I	# O	USER	RESP	I/R	# I	# O	TRM	RESP	I/R	# I	. . .	OPREJT	OPREJP	RJM	RRSVP
12:38:01	0	0	659	7	13 I	1	1	8	0	1 I								
12:43:01	0	0	835	10	15 I	2	2	2	0	2 I	0	0	3	. . .	1	0	0	0
12:48:01	0	0	849	6	18 I	2	1	43	1	4 I	6	6	1	. . .				
12:53:01	2	0	604	6	18 I	9	2	32	2	4 I				. . .				
12:58:02	2	0	469	7	18 I	19	3	14	1	4 I				. . .				
13:03:02	2	0	247	4	18 I	39	1	5	2	4 I				. . .	1	0	0	0
13:13:03	1	0	976	6	17 I	7	1	33	1	4 I				. . .	7	0	4	0
TOTAL	6	2	4639	46	17	10	1	137	7	4		3	4					

Figure 9-16. Response Time Report

H6000 - DN355 REJECT REPORT FOR NMCC ON 021683

TERMINAL ID	RJM	RSVP	ECHO	OPCRJT	OPCRJP COMMANDS
A4	0	35	0	3	0
01	0	0	0	3	0
A4	0	0	0	1	0

DN/TERMINAL REPORT FOR NMCC2 ON 021683

TERMINALS ON DATANET 0

A6 A8 AQ A4 AA 02 01 AB AT AD A0 AL AJ AF AR A2 AP A7 AN AM A5 AS

TERMINALS ON DATANET 1

FJ CN CA F6 C6 C9 CK FM CT F3 F2 C8 CL CC FU FE FV CE CM CV CI CP F7 CJ C7 CO F4 FI FN C2 CR CD

Figure 9-16.1. H6000-DN355 Reject Report

```
                    ABORT REPORT FOR NMCC  ON 030382

TERMINAL ID BP ON SUBSYSTEM UNKNWN DISCONNECTED AT 12:57:12 AFTER   246 SECS WITH   876 INPUTS AND   876 OUTPUTS

TERMINAL ID BP ON SUBSYSTEM NACE-U DISCONNECTED AT 12:57:29 AFTER     3 SECS WITH     6 INPUTS AND     7 OUTPUTS

TERMINAL ID B4 ON SUBSYSTEM UNKNWN DISCONNECTED AT 12:57:49 AFTER   283 SECS WITH   876 INPUTS AND   876 OUTPUTS

TERMINAL ID 02 ON SUBSYSTEM UNKNWN DISCONNECTED AT 13:04:39 AFTER   692 SECS WITH     0 INPUTS AND   280 OUTPUTS

TERMINAL ID 01 ON SUBSYSTEM UNKNWN DISCONNECTED AT 13:04:39 AFTER   692 SECS WITH   599 INPUTS AND   599 OUTPUTS

                    ***** ABORT 355-0 AT 13:28:35 AT RCDCNT     144
```


Figure 9-16.2. Abort Report

TS1 INITIAL PARAMETER REPORT

SIZE PARAMETERS

INITIAL TS1 MAX SIZE	180K	SIZE GROWTH FACTOR	20K
MINIMUM TS1 SIZE	60K	SIZE REDUCTION FACTOR	20K
MAX TIME ALLOWED FOR SIZE CHANGES	60SECS	MAX TIME ALLOWED BETWEEN GMORE REQUESTS	10SECS
MEMORY SIZE REDUCTION TIME INTERVAL	120SEC		

LIMIT PARAMETERS

MAX NUMBER OF TERMINALS	90	MAX NUMBER CONCURRENT DRL TASKS	4
LARGE SUBSYSTEM SIZE FENCE	20K	LARGE SUBSYSTEM WAIT TIME PENALTY	8
NUMBER OF 32MS TIME QUANTUMS	8	FREQUENCY OF PRIORITY B DISPATCHING	1
RECONNECTION TIME LIMIT	300SEC		

SWAP FILES

NUMBER OF SWAP FILES IN USE	2	SWAP FILES IN USE: #S #T	
MIN SWAP FILE SIZE	1200LLINKS	MIN SWAP FILE GROWTH SIZE	300LLINKS

ALLOCATED DEVICES

FILE	DEVICE/FILE NAME
#D	DQ1
#P	
#Q	
.D	.D
SS	
#S	
#T	

Figure 9-16.3. TS1 Initial Parameter Report

MAILBOX BUSY REPORT FOR NMCC2 ON 062582

TIME OF DAY	DN	BUSY MAILBOXES	# MBX CHECKS	POSSIBLE MAILBOXES	% BUSY	DENIALS	% DENIED	MAX BUSY	MAX DENIED
17:44:52	0	3606	12273	85911	4			3	
17:44:52	1	10110	12273	85911	12	1	0.001	7	1
17:54:52	0	3883	12653	88571	4			3	
17:54:52	1	10153	12653	88571	11	10	0.011	5	2
18:04:55	0	3090	7736	54152	6			2	
18:04:55	1	5304	7736	54152	10	1	0.002	6	1

Figure 9-16.4. Mailbox Busy Report

9.6 Default Option Alteration

The Communication Analysis Monitor Data Reduction Program (CAMDRP) uses the trace tape generated by the GMC and user data cards as input. The user has several optional inputs. These options include Timeframe Reduction, Delta Timeframe, Histogram, Trace Dump, User Think Time Limit, Machine Response Time Limit, Record Count Limitation, Terminal Busy Limit, Response Time Report time frame, and Terminal Mailbox Dump. These options are evoked by specifying an input option (action code) and any other required inputs specified in the following subsections. Default options are given in subsection 9.6.13.

All inputs are free format with the only requirement being that if any value is to be a zero, the user must type the number 0 on the data card. A zero value may not be inputted as a blank. In addition, at least one data card with the word END specified is required.

9.6.1 Timeframe Reduction Report (Action Code TIME). This option allows the user to specify up to five different consecutive processing timeframes. For example, the user may collect data from 0500 to 2200 hours and specify that he wants to process only 1900 to 2000 hours and 2100 to 2200 hours. The timespans set through this option apply to all report processing.

Up to five timespans may be specified and they must be chronologically ordered. In specifying time spans, a start time and a stop time must be supplied. If a start time is specified, but no stop time is specified, there should be no data on the card after the start time is specified. All times are expressed as S1M1 where S1 is an hour and M1 is a minute. All times are considered to be on a 24-hour clock, and must be expressed as 4-digit fields with no intervening blanks. If the user wants to request the time 4:07 he must input a "0407." Figure 9-17 has the card formats and an example.

9.6.2 Delta Timeframe Report (Action Code DELTA). This option allows the user to produce a report (subsection 9.5.4) for each DATANET-355 giving the number of active terminals, input characters, and output characters every X seconds where X is a user-specified value. The first data card contains the option name "DELTA". The second data card contains the length of the delta timeframe in seconds followed by the word COMPRESS or NO. This word should be separated from the time by at least one intervening blank. The word COMPRESS is included if the user wants to suppress the printing of lines with all zeros for data (i.e., no activity during this timeframe), or the word NO is included if the user wants blanks or zero lines printed.

```
CARD   1   TIME
CARD   2   N
CARD   3   H1M1     H2M2     H3M3     ....H10M10
```

Where

N = Number of different times appearing on card 3
H1M1 = Time used to define a timespan. Individual times must be separated from each other by at least 1 blank column. (see section 9.6.1)

Figure 9-17. Card Format - Input Option TIME

9.6.3 Histogram Report (Action Code HISTG). This option allows the user to get up to 20 groups of histograms (subsection 9.5.5). Each group includes a histogram for machine response time in seconds, user think time in seconds, and session length time in minutes.

These histograms can be produced by subsystem, device types, or terminal ID. The user is allowed any combination of these three categories, with a maximum of 20 total items. The input option name HISTG is used for these histograms. The second data card contains the number of subsystems wanted, number of device types wanted, and number of terminal IDs wanted.

The next card(s) contains the subsystem names, if requested. The following card(s) contains the device types, if requested. (See table 9-2 for a list of acceptable device type names.) The next card contains up to 20 two-character terminal IDs. The format of these cards is shown in figure 9-18.

9.6.4 Trace Dump Report (Action Code LIST or ALL). These options allow the user to get a formatted dump of all CAMDRP traces. The input option name for the Trace Dump Report (subsection 9.5.2) is LIST. The input option name for both Trace and Summary Report data (subsection 9.5.2) is ALL. If the user wishes to limit the amount of output, an option exists to print out only certain subsystems, terminal IDs, or terminal types. This is accomplished by using the HISTG input option with the LIST input option. In this case, no histograms will be produced and the trace dump will be limited to the requested terminals, subsystems, and terminal types. This option consists of a single data card containing the word LIST or ALL.

9.6.5 Record Count Limitation (Action Code NREC). This option permits the user to process only a certain number of tape records. The input option name for the Record Count Limitation is NREC. The second data card contains the number of tape records to be read.

9.6.6 Response Time Limit (Action Code RESP). This option permits the user to specify an acceptable machine response time in seconds. If any terminal exceeds or equals this limit, the terminal ID and response time are printed out (subsection 9.5.6). The format of this option is the word RESP on the first data card and the response time limit on the second data card.

9.6.7 Think Time Limit (Action Code THINK). This option permits the user to specify an acceptable User Think Time in seconds. If any user exceeds or equals this limit, the terminal ID and think time are printed out (subsection 9.5.7). The format of this option is the word THINK on the first data card and the think time limit on the second data card.

Table 9-2. Acceptable Device Type Names

DEVICE TYPE	CAMDPP INPUT
Remote Computer	RCT
Teletype	TTY
VIP 775	VIP775
VIP 786W	VP786W
VIP 12	VIP12
VIP 22	VIP22
VIP 7705	VP7705
2741	2741
RLP 300	RLP300
Mass Store Link	MSLINK
MRS 200 Document Handler	MRS200
DRD 200 Document Handler	DRD200

```
CARD   1   HISTG
CARD   2   B   C   D
CARD   3   E   F   G...
CARD   4   H   I   J...
CARD   5   K   L   M...
```

Where

B = Number of subsystems wanted
C = Number of device types wanted
D = Number of terminal IDs wanted
E,F,G = Up to 10 subsystem *names* separated by at least one blank (may go to more than 1 card)

H,I,J = Up to 10 device types separated by at least one blank (may go to more than 1 card)
K,L,M = Up to 20 terminal IDs separated by at least one blank

Figure 9-18. Histogram Reports, Input Option HISTG

9.6.8 Terminal Mailbox Dump (Action Code MAIL). This option allows the user to get a dump of the terminal traffic collected for the specified terminal IDs. (Reference subsection 9.5.2.2) This output can be in ASCII, BCD, OCTAL, or all three. See figure 9-19 for the format of this option. NOTE - this option will turn on the LIST option.

9.6.9 Terminal Busy Limit (Action Code BUSY). This option allows the user to change the threshold for the High Terminal Usage Report (subsection 9.5.8). Whenever a terminal is connected to the system for greater than the desired limit, the terminal ID will be printed. Two cards are required for this option. The first card has the word BUSY on it and the second card contains a % busy limit value.

9.6.10 W6.4/2H Data Reduction (Action Code RN). This option allows the user to process a GMF data tape (W6.4/2H or W7.2/4Jx) under a W6.4/2H software release. It consists of the letters RN on a data card.

9.6.11 Response Time Report Time Frame (Action Code RATE). This option allows the user to produce a report (subsection 9.5.10) giving the average response time over a time interval for both TimeSharing and all subsystems combined. This option requires two data cards. The first card contains the word RATE and the second card contains the number of minutes between response time printouts.

9.6.12 Response Time Report SNUMB (Action Code SNUMB). This option allows the user to produce response times for up to three different DAC subsystems (subsection 9.5.10), giving the average response over a time frame for the system and each requested SNUMB. This option requires three data cards. The first card contains the word SNUMB. The second data card contains the number of SNUMBs to be used. The last data card contains these SNUMBs.

NOTE: If one of these SNUMBs is TSS, all TSS logons (TSS, TS1, TS2, TS3, TS4) will be presented under the heading TSS in this report. The user can also request a separate column for any of the TS SNUMBs.

9.6.13 Disconnect Report Line IDs (Action Code DISC). This option allows a user to define up to 10 different line IDs to be reported on the Abort/Disconnect Report every time they disconnect (see section 9.5.13). This option consists of three data cards. The first card contains DISC. The second card contains the number of line IDs to be monitored. The third card contains the specified line IDs, separated by a blank.

9.6.14 Terminate Input Options (Action Code END). This option is required as the last input option data card. It may be the only data card if standard default options are selected. It consists of the word END on a data card.

9.6.15 Default Options. The default options for the variable input are as follows:

ACTION CODE	Option	Default Value
TIME	Timeframes	None, total tape processed.
DELTA	Delta Timeframes	None, this report is not processed. If Delta time is given but not the word COMPRESS, all data are printed.
HISTG	Histograms	None, no histograms produced.
LIST	Trace	None, report data reduction is done, not trace dump.

THIS PAGE INTENTIONALLY LEFT BLANK

```
CARD   1   MAIL
CARD   2   N
CARD   3   B   C   D
```

Where

N = Number of terminals to be dumped (1 or 2)
B = The letters BCD for a BCD dump
    The letters ASCII for an ASCII dump
    The letters OCTAL for an OCTAL dump
    The letters ALL for all three dumps
C = Terminal ID #1
D = Terminal ID #2 (present only if N=2)

Figure 9-19. Terminal Mailbox Dump, Input Option MAIL

NREC	Record Count Limit	None, total tape processed
RESP	Response Time Limit	9999 seconds
THINK	User Think Time Limit	9999 seconds
MAIL	Terminal Mailbox Dump	None, no mailbox dump is produced. Only valid with specially collected data tapes.
BUSY	Terminal Busy	75%
RN	Process on NW6.4	Assumed CAMDRP will be executed under a WW7.2/4JS system
RATE	Response Report Time Frame	9999 minutes (no reports printed)

9.7. JCL

Following is a sample of the JCL required to run CAMDRP. The only change required involves the use of input data. CAMDRP requires 50K of memory to process, but will use 65K during the load procedure. Immediately upon loading, the additional 15K will be released.

```
$ IDENT       User Accounting Information

$ USERID      User Identification/Password/Classification

$ SELECT      B29IDPX0/OBJECT/CAM

$ LIMITS      999,50K,-4K,15K

$ TAPE        01,X1D,,Tape #

$ DATA        I*

   DATA cards

   (at least an "END" card must be present)
```

9.8. Multireel Processing

If more than a single reel of data has been collected, a series of messages will be outputted to the console

informing the operator that a new data reel is required. The following are the messages produced.

a. DISMOUNT REEL #XXXXX THEN MOUNT REEL NUMBER YYYYY ON ZZZZZ

In this case, XXXXX is the old reel number, YYYYY is the new reel number, and ZZZZZ is the tape drive ID.

If the operator fails to mount the new tape, the above message will be repeated three times, after which the program will terminate, and all reports produced.

b. IS TAPE XXXXX MOUNTED ON DRIVE ID YYYYY (Y/N)

In this case, XXXXX is the tape number being requested for mounting and YYYYY is the tape drive ID.

This message occurs when the data reduction program finds the wrong tape has been mounted (by comparing internally generated tape labels). If the operator answers N, the message in (c) below is produced. If the operator answers Y, the data reduction program will terminate and all reports will be produced. In this case, the data reduction program is unable to process the tape. Even though the operator is mounting the correct tape, the internal label on the new tape does not match that being requested by the old tape. The user should check the data collection session to insure that the operator did not respond with an incorrect tape number during multireel change.

After entering the Y or N, the operator will need to hit the EOM key twice in order for the response to be transmitted.

c. WRONG REEL JUST MOUNTED, DISMOUNT AND MOUNT REEL XXXXX ON ZZZZZ

In this case, XXXXX is the new reel number, and ZZZZZ is the tape drive ID.

d. CAN TAPE XXXXX BE MOUNTED ON DRIVE YYYYY (Y/N)

In this case, XXXXX is the new desired reel and YYYYY is the tape drive ID.

If the operator fails to answer this message it will be repeated until he responds with a "Y" for YES or "N" for NO. If he types in "Y", then message (a) will be repeated. If the types in "N", then the program will be terminated and all reports will be produced.

9.9 Tape Error Aborts

During the course of processing it is possible that the operator will be required to abort the data reduction program due to an irrecoverable tape error. If such a condition occurs, the operator should abort the job with a "U" abort. This will allow the data reduction program to enter its wrap-up code processing and produce all reports generated prior to the tape error.

SECTION 10. DATANET-355 DATA REDUCTION PROGRAM (DDRP).

10.1 Introduction

The DATANET-355 Data Reduction Program (DDRP) is a FORTRAN program that sequentially processes data the DATANET-355 GRTS Monitor collected and wrote on tape. DDRP produces a number of reports describing the performance of the DATANET-355 front ends. CPU and resource usage, subchannel usage, and user response times are but some of the many reports that can be obtained. Report descriptions are presented in subsection 10.5.

There are two inputs to the DDRP. The first is the data tape produced by the GRTM in the General Monitor Collector. The second is a set of report option control cards used to alter the reports in some way other than the standard default. The various user input options and their formats are described in subsection 10.6. The actual reports produced by DDRP are described in subsection 10.5.

10.2 Data Collection Methodology

The GRTM in the General Monitor Collector processes a GMF generated trace (type 62) and collects information to monitor the performance of the DATANET-355 front ends. The performance information is collected within the DATANET-355 and transmitted from the DATANET to the GMF which is resident in the H6000. Under normal operating conditions, the DATANET-355 does not generate performance data. In order to generate the desired performance data, it is necessary to alter the standard DATANET-355 GRT software. The method for altering the GRT software is described in subsection 5.2.7.3. The alters which must be applied to the GRT software not only allow the DATANET-355 to generate performance data, but also provide the mechanism by which the DATANET-355 communicates with the GMF. This communication process is described in subsections 5.2.7.4 and 5.2.7.5. The format for the GRTM generated records used by the DDRP are described in subsection 5.4.8.

10.3 Analytical Methodology

Within the GRTS-II Monitor, a series of 10 different events are captured and the time of their occurrence is recorded. The 10 different events are as follows:

1. Connect time of a terminal.

2. DN355 requesting output from the H6000.

3. H6000 telling the DN355 to accept direct output.

4. H6000 telling the DN355 to accept direct output and return input.

5. DN355 telling the H6000 to accept input.

6. H6000 telling DN355 that it accepted the input.

7. Disconnect time of a terminal.

8. DN355 has received output from the H6000.

9. DN355 has begun to transmit data to a terminal.

10. DN355 has completed transmission of data to a terminal.

Using the above events, seven different response times are defined and reported at the DATANET level and for specified terminals. These seven different response times are as follows:

1. H6000 Data Acceptance Time From the DATANET (Event 6 - Event 5)

2. DATANET Acceptance Time From the H6000 (Event 8 - Event 3 or 4)

3. DATANET Processing Time (Event 9 - Event 8)

4. DATANET Transmission Time (Event 10 - Event 9)

5. User Processing Time + DATANET Input Time (Event 5 - Event 10)

6. Machine Response Time (Event 3 or 4 - Event 6 or 2)

7. System Response Time (Event 4 - Event 6)

Machine Response Time is calculated as the sum of the individual H6000 responses. As an example, when a user lists a large file on a VIP, the H6000 will transmit several lines at a time to the DATANET-355. When the screen is full, the H6000 requests input from the user (i.e., clear screen) so the H6000 knows when to continue transmitting data. The sum of the transmission times of each of the segments, up to the request for input, is considered machine response time. A sample trace sequence by event number would be:

a. Event #5 (DATANET-355 asking H6000 to list file)

b. Event #6 (H6000 has received request)

c. Event #3 (H6000 sending four or five lines back to DATANET)

d. Event #8 (DATANET-355 has received output)

e. Event #9 (Transmission to terminal begun)

f. Event #10 (Transmission to terminal completed)

g. Event #2 (DATANET requesting more output from H6000)

h. Event #3 (H6000 sending additional lines of the file)

i. Event #8 (as before)

j. Event #9 (as before)

k. Event #10 (as before)

l. Event #2 (DATANET requesting more output from H6000)

m. Event #4 (H6000 has now completed sending one screen full of data and is requesting input from user; i.e., clear screen, break, etc.)

$$\text{Machine Response} = (c-b) + (h-g) + (m-l)$$

System response is calculated as the total elapsed time from user input to request for more user input. In the above example, this would be (m-b).

10.4 Data Reduction Methods

The DATANET-355 Data Reduction Program outputs consist of histograms, serial plots, and tabular reports. These outputs have a number of parameters which can be modified by card input at run time. Most reports may be turned on or off, and a set of timespans may be specified for measurement. These options are completed described in subsection 10.6. The method for altering histograms and plots require an understanding of the internal parameters that control these outputs. These parameters will be described in this subsection.

The user has the option of obtaining a set of eight histograms describing the response being obtained during various phases of processing. These histograms will be described in subsection 10.5.3. These histograms are normally not produced and must be requested for specific DATANETs and/or terminals. The method for requesting histograms will be described in subsection 10.6.1 and 10.6.20. A set of eight histograms will be produced for each requested DATANET and/or terminal. The definitions of these response times are given in subsections 5.2.7.7.3, 5.4.8, and 10.3 of this document.

There are two characteristics for each histogram directly available to the user. The first of these characteristics is the lowest value, stored in the LOW array. Any input received by this histogram that is

less than this value will be entered into the first histogram cell along with all entries that fall within the range of that cell. For all eight histograms, this array defaults to zero. There is an input action which allows the user to modify this array for any of the eight histograms (see subsection 10.6.20).

The second characteristic is the range or width of each cell. It is determined by the INTEV array. Once again, there is a user action which allows the user to modify this array for any of the eight histograms. The default values are (5,2,2,1,10,2,2,1) respectively.

The size of the histogram is determined from the internal parameter MXTBSZ which specifies how many cells or "buckets" are contained in the histogram. It indirectly specifies the upper value which could fall within the display range of the histogram. If an entry is made above this range, it is placed in an out-of-range cell of the histogram. The number of the out-of-range occurrences is printed with the histogram, along with their average value, so that the histogram can be altered to handle the variability of the range properly. If an increase beyond this is desired, MXTBSZ must be increased causing an increase in the memory size of the data reduction program. This must be done via an edit of the source code, ensuring that all values of MXTBSZ are changed. The current value of MXTBSZ is 50.

There are three characteristics directly available to the user for each individual plot used. The actual plot formats and information that is plotted will be described in subsection 10.5.2. The first characteristic, MAXNUM, is the maximum absolute length of the abcissa (print lines) to be plotted. The second and third characteristics, YMAX and YMIN, define the upper and lower limits of the horizontal ordinate. These two values, along with the width of the print page, determine the value of one increment, along the horizontal axis (referred to as DELTA on the plot). The actual formulas used are:

```
DELTA = INTEGER ((YMAX - YMIN)/126 + .5)
YMAX  = DELTA* 126 + YMIN
```

Figure 10-1 shows the default values for the plots. Subsection 10.6.2 describes how plot parameters may be altered.

10.5 DDRP Output

The DDRP outputs consist of histograms, serial plots, and tabular reports. Table 10-1 lists all reports available to the user by default or by input request. The input requests will be described in subsection 10.6 and the actual reports will be described in this subsection. The reports of the DATANET Reduction Program can be used to study the following for each active DATANET:

Plot #	Max Lines In Plot	Y-Axis Max	Y-Axis Min
1	No limit	500	0
2	No limit	3000	0
3	No limit	30000	500
4	No limit	50000	15000
5	No limit	100	0
6	No limit	100	0
7	No limit	10000	0
8	No limit	150	0
9	No limit	100	0

Figure 10-1. Default Values For Plots

Table 10-1. DATANET-355 Reports (Part 1 of 3)

Tabular Reports Generated by Default

Report ID/ File Code	Description
RESPD/21	Average Response Time For All Users by DATANET
BUFF/21	HSLA/Subchannel Transmission Report by DATANET
RESPL/28	Average Response Time For Specifically Designated Line IDs
HSLA/21	HSLA Subchannels Being Monitored
CARD/21	Card Images from I* of the GMF Monitor (always generated)
-/23	List of Active Line IDs (always generated)
THRESH/20	HSLA Threshold Report (default = 9999999)
-/42	Response Interval Unmatched Pairs Verification (always generated when DATANET reports are produced)
-/42	Response Reject Messages

Tabular Reports Generated By Request

MATCH/20	Response Interval Matched Pairs Verification
TRACE/22	Annotated List of the DATANET Trace

Table 10-1. (Part 2 of 3)

Histogram Reports Generated by Request
For a DATANET or Special Line ID

Report ID/ File Code	Description
HISTG1/21	H6000 Data Acceptance Time From Requested Device
HISTG2/21	Data Acceptance Time For Requested Device From H6000
HISTG3/21	Processing Time For Requested Device
HISTG4/21	Transmission Time For Requested Device
HISTG5/21	User Processing + Data Input Time For Requested Device
HISTG6/21	H6000 Machine Response Time For Requested Device
HISTG7/21	System Response/User Request For Requested Device
HISTG8/21	H6000 Responses/User Request For Requested Device
	Above set of eight histograms is repeated for each requested DATANET or special line ID.

Table 10-1. (Part 3 of 3)

Plot Reports Generated by Default

Report ID/ File Code	Description
PLOT1/31	Transactions Sent to Host (cumulative every 5 minutes)
PLOT2/32	Transactions Received From Host (cumulative every 5 minutes)
PLOT3/33	36-bit Words Sent to Host (cumulative every 5 minutes)
PLOT4/34	36-bit Words Received From Host (cumulative every 5 minutes)
PLOT5/35	% Idle Time Over a 5-Minute Period
PLOT6/36	% Buffer Requests Denied Over a 5-Minute Period
PLOT7/37	18-bits Average Number of Words Available For Buffers Over a 5-Minute Period
PLOT8/38	Average Number of Users Logged On Over a 5-Minute Period
PLOT9/39	Host RSVPs Received by DATANET (cumulative every 5 minutes)

- Amount of activity by DATANET - includes buffer denials, number of transactions sent to and received from the host, number of 36-bit words sent to and received from the host, host RSVPs received, the amount of time spent idle, and a count of the number of users logged onto the DATANET;

- Response time by DATANET and by specified line ID - includes analysis of a series of response time delays;

- HSLA subchannel activity by DATANET - includes number of transmits and receives.

Immediately prior to the reports listing, a summary of processing information is printed. Included in this summary are the following:

- list of all input options selected by the user;

- indication of sequential input tapes being mounted;

- error messages (self-explanatory);

- if the timespan option was used, an indication of when the various timeframes were reached.

In the following sections an example of each report is displayed and a simple explanation of what data was used in the creation of the report is given.

10.5.1 DATANET-355 Tabular Reports. The following tabular reports can be obtained in a data reduction run.

10.5.1.1 Average Response Time for All Users by DATANET (Report RESPD). The report shows the average response times, for each response time category by DATANET. The number of requests and accumulated response times are also displayed. Each response time catagory is defined in subsection 10.3. The report also displays the tapes from which the data was obtained and the date and time to which the report corresponds. The times for which each DATANET was active are given at the bottom of the report. This report is printed as file code 21 at closure of a timespan and end of processing. At each printing, measurements begin anew for each DATANET affected. Refer to figure 10-2 for an example of this report. In generating this report, the individual response times for all terminals on a DATANET are averaged together to determine the response time of the DATANET. In the current version, the default value for the number of terminals that can be processed is 50. If more then 50 terminals can be active, on a given DATANET, the parameter MAXID must be changed via an edit, and the program recompiled. In addition, the size of the program will

AVERAGE RESPONSE TIMES FOR ALL USERS BY DATANET FOR SYSTEM NMCC ON 81-08-27
USING TAPES D0001 D0002 D0003

DATA NET	H6000 DATA ACCEPTANCE TIME NUMBER OF REQUESTS	TOTAL WAIT (SECONDS)	AVERAGE (MSEC)	DN 355 DATA ACCEPTANCE TIME NUMBER OF REQUESTS	TOTAL WAIT (SECONDS)	AVERAGE (MSEC)
0	192	0.960	5.	384	1.920	5.
1	144	0.720	5.	282	1.440	5.

DATA NET	DN 355 PROCESSING TIME NUMBER OF REQUESTS	TOTAL WAIT (SECONDS)	AVERAGE (MSEC)	TRANSMISSION TIME NUMBER OF REQUESTS	TOTAL WAIT (SECONDS)	AVERAGE (MSEC)
0	384	1.920	5.	384	1.920	5.
1	268	1.440	5.	288	1.440	5.

Figure 10-2. Average Response Time for All Uses by DATANET (Report RESPD) (Part 1 of 2)

USER THINK TIME

DATA NET	NUMBER OF REQUESTS	TOTAL WAIT (SECONDS)	AVERAGE (MSEC)
0	384	768.000	2000.
1	268	567.000	2000.

H6000 MACHINE RESPONSE TIME

NUMBER OF REQUESTS	TOTAL WAIT (SECONDS)	AVERAGE (MSEC)
372	564.000	1516.
279	423.000	1516.

SYSTEM RESPONSE PER USER REQUEST

DATA NET	NUMBER OF REQUESTS	TOTAL WAIT (SECONDS)	AVERAGE (MSEC)
0	384	727.000	1893
1	268	658.000	2455

OF H6000 RESPONSES PER USER REQUEST

NUMBER OF REQUESTS	TOTAL RESPONSES	AVERAGE
384	825	2.15
268	652	2.43

DATANET 0 SEEN FROM 06:06:00 TO 08:22:00
DATANET 1 SEEN FROM 08:10:00 TO 08:22:00

Figure 10-2. (Part 2 of 2)

increase by 15 words for each integer increase of MAXID. This report will be generated only if the response time portion of the DATANET-355 GMF Monitor was active during data collection (see subsection 5.5.9).

10.5.1.2 HSLA/Subchannel Transmission Report by DATANET (Report BUFF). This report displays the HSLA subchannel usage by DATANET. For each subchannel being monitored, the number of transmits and receives are printed. A totals line for each DATANET is also displayed. The data for this report is obtained from the terminal monitoring entries (type 3 records) of the host buffer. This report, one for each affected net, is output on file code 21 at closure of a timespan and end of processing. At each printing, measurements begin anew. This report will be generated only if the HSLA monitoring portion of the DATANET-355 GMF Monitor was active during data collection (see subsection 5.5.9). Refer to figure 10-3.

10.5.1.3 Average Response Time for Specifically Designated Line IDs (Report RESPL). This report presents the average response times, for each response time category, for each line ID designated for special analysis. Each response time category is defined in subsection 10.3. The response times are tallyed on a terminal session basis. A "connect to slave" record (subtype 1 of the response time record entries) initiates the data collection for a specific line ID, and a "line disconnect" record (subtype 7 of the response time record entries) terminates collection. The line disconnect sends a line-of-print to the report, showing the connect and disconnect TOD as well as the average response times for that line ID. A summary paragraph is listed whenever a net disconnect, closure of a timespan or end of processing occurs. The summary shows the average response times for each active line on the net affected, as well as the TODs during which the data was collected. The period during which the data was collected for the summary depends upon what causes the previous summary - a net disconnect resets average response time counters only for line IDs attached to that net. Any other cause resets the average time counters for all line IDs. File code 28 is dedicated to this report. An example is shown as figure 10-4. This report will be generated only if the response time portion of the DATANET-355 GMF Monitor was active during data collection (see subsection 5.5.9).

10.5.1.4 HSLA Subchannels Being Monitored (Report HSLA). This report, derived from terminal monitoring entries (type 3 records), lists each HSLA number and subchannel monitored for data collection (see subsection 5.9.9) in the GRTS-II monitor run. The report is written once for each DATANET eligible for processing, when the net's first T62 trace is reduced. Figure 10-5 shows an example of this report. It is written to file code 21.

START 08:02:00 81/09/27
STOP 09:00:00

TAPE NOS. 000001 000002

HSLA/SUBCHANNEL TRANSMISSION REPORT FOR DATANET 0

HSLA	HSLA S/C	TRANSMITS	RECEIVES
1	1	13000	13000
1	2	13000	13000
2	1	13000	13000
DATANET 0 TOTALS		39000	39000

Figure 10-3. HSLA/Subchannel Transmission Report (Report BUFF)

START 8:08:00 ON 81/10/27 FOR SYSTEM NMCC FIRST TAPE WAS D0001

AVERAGE RESPONSE TIME IN MILLISECONDS FOR LINE-ID's DG DN DM

LINE-ID	CONNECT TOD	DISCONN TOD	DAC NAME	H6000 ACCEPT TIME	DN ACCEPT TIME	DN PROC TIME	TRANS TIME	USER TIME	MACH TIME	SYST TIME	# RESP
DG	08:10:00	8:12:00	SENDA	5.	5.	5.	5.	2000.	1516	1700	8
SUMMARY FOR THE CHOSEN LINE-ID'S AT 08:13:00 27/10/81 TRIGGERED BY DATANET DISCONNECT											
DN	08:10:05	8:13:00	*UNKN*	5.	5.	5.	5.	2000.	1516	1700	8
SUMMARY FOR THE CHOSEN LINE-ID'S AT 08:13:00 27/10/81 TRIGGERED BY ATANET DISCONNECT											
DM	08:06:08	8:13:00	*UNKN*	5.	5.	5.	5.	2000.	1516	1700	8
DG	08:16:00	8:18:00	SENDA	5.	5.	5.	5.	2000.	1516	1700	8
DG	08:20:00	8:22:00	SENDA	5.	5.	5.	5.	2000.	1516	1700	8
SUMMARY FOR THE CHOSEN LINE-ID'S AT 08:26:00 27/10/81 TRIGGERED BY TIME-SPAN ENDED											
DM	08:16:00	8:26:00	SENDA	5.	5.	5.	5.	2000.	1516	1700	8
DG	08:16:00	8:26:00	SENDA	5.	5.	5.	5.	2000.	1516	1700	8

Figure 10-4. Average Response Time for Designated Lines (Report RESPL)

```
START 8:08:00 ON 81/10/27 FOR SYSTEM NMCC

TAPE NOS. D0001 D0002 D0003

          HSLA SUBCHANNELS BEING MONITORED BY DATANET
DATANET 1
               MONITORING HSLA 1 SUBCHANNEL 0
               MONITORING HSLA 1 SUBCHANNEL 1
               MONITORING HSLA 1 SUBCHANNEL 2
               MONITORING HSLA 1 SUBCHANNEL 3
               MONITORING HSLA 1 SUBCHANNEL 4
               MONITORING HSLA 1 SUBCHANNEL 5
               MONITORING HSLA 1 SUBCHANNEL 6
               MONITORING HSLA 1 SUBCHANNEL 7
               MONITORING HSLA 1 SUBCHANNEL 8

DATANET 3
               MONITORING HSLA 1 SUBCHANNEL 0
               MONITORING HSLA 1 SUBCHANNEL 1
               MONITORING HSLA 1 SUBCHANNEL 2
               MONITORING HSLA 1 SUBCHANNEL 3
               MONITORING HSLA 1 SUBCHANNEL 4
               MONITORING HSLA 1 SUBCHANNEL 5
               MONITORING HSLA 1 SUBCHANNEL 6
               MONITORING HSLA 1 SUBCHANNEL 7
               MONITORING HSLA 1 SUBCHANNEL 8
               MONITORING HSLA 1 SUBCHANNEL 9
               MONITORING HSLA 1 SUBCHANNEL 10
               MONITORING HSLA 1 SUBCHANNEL 11
               MONITORING HSLA 1 SUBCHANNEL 12
               MONITORING HSLA 1 SUBCHANNEL 13
               MONITORING HSLA 1 SUBCHANNEL 14
               MONITORING HSLA 1 SUBCHANNEL 15
```

Figure 10-5. HSLA Subchannels Being Monitored (Report HSLA)

10.5.1.5 Card Images from I* of the GRTS-II Data Collection (Report CARD). These card images were originally control data from the GMC. The first card was interpreted by GMC as specifying which of the ten monitors were to remain inactive during the data collection run. If monitor #6, coded "M6," was not listed, then the GRTS-II monitor should have been active and T62 traces expected. Succeeding card images informed GMC of the datanet monitoring options selected by the user (see subsection 5.5). They are repeated here to provide environment data. Figure 10-6 gives an example of this report. It is printed once on file code 21.

10.5.1.6 List of Active Line IDs. Displayed on this report is the identity of each line ID active during the reduction (to a maximum of 50). Only nets selected for analysis (not deselected) contribute. The data is taken from every response time entry (type 2) carrying a line ID tag. The report is listed on file code 23 whenever closure of a timespan, disconnect, or end of processing occurs. Figure 10-7 is an example. This report will be generated only if the response time portion of the DATANET-355 GMF Monitor was active during data collection (see subsection 5.5.9).

10.5.1.7 HSLA Threshold Report (Report THRESH). This report, derived from terminal monitoring entries (type 3 records), indicate whenever the number of transmits or receives, corresponding to any monitored HSLA subchannel (see subsection 5.5.9), exceeds a preset threshold value. The default threshold values are set high so that this report is not normally produced. If the user wants to obtain this report he must activate user input option THRESH (see subsection 10.6.17). The threshold values are compared to the number of transmits or receives monitored across an HSLA subchannel between any two consecutive DATANET traces. The message produced indicates the threshold surpassed (transmits=XMITS, receives=RCVS), the time of day, the HSLA subchannel, and the amount of time that has passed since the previous trace. Figure 10-8 is an example of this report. This report will be written to file code 20.

10.5.1.8 Response Interval Unmatched Pairs Verification. Unmatched pairs are said to occur for a specific line ID when successive data response time records for that line cannot be affirmed as matched pairs (see subsections 10.6.21, 10.3 and 10.5.1.9).

Unmatched pairs are noted within the DATANET-355 Reduction Program, and are printed when discovered. The report appears on file code 42. A summary count is printed whenever closure of a timespan or end of processing occurs. A typical line of output is shown below.

Example of unmatched pairs verification:

UNMATCHED ACCEPT DIRECT OUTPUT RECORD FOR DNET 1 LNID DG SUBSYSTEM TSS TYPE VIP 785 AT 08:09:33

```
CARD INPUT IMAGES FROM I* FILE OF THE GMF MONITOR FOR SYSTEM NMCC AT 08:13:00 ON 81/10/27 FIRST TAPE WAS D0001

         1         2         3         4         5         6         7         8
12345678901234567890123456789012345678901234567890123456789012345678901234567890

INPUT CARD # 1    M0 M1 M2 M3 M4 M7 M8 M9
INPUT CARD # 2    D1;HSLA1;SCHO-31;
INPUT CARD # 3    D3;HSLA1;SCHO-31
```


Figure 10-6. Card Images from I* FILE of the GMF Monitor (Report CARD)

```
ACTIVE LINE ID REPORT FOR SYSTEM NMCC AT 08:13:00 ON 81/10/27
USING TAPES D0001 D0002 D0003


          DATANET 1 SEEN FROM  08:02:00 TO 08:12:00  REPORT TRIGGERED BY DATANET DISCONNECT


                                        LINES SEEN ON SELECTED NET
                                            1-DG    1-DA    1-DM
**********************************************************************************************************************


          DATANET 0 SEEN FROM  08:02:00 TO 08:12:00  REPORT TRIGGERED BY DATANET DISCONNECT


                                        LINES SEEN ON SELECTED NET
                                            0-AD    0-AB
**********************************************************************************************************************


          DATANET 1 SEEN FROM  08:16:00 TO 08:26:00  REPORT TRIGGERED BY DATANET DISCONNECT


                                        LINES SEEN ON SELECTED NET
                                            1-DG    1-DA    1-DM
**********************************************************************************************************************
```


Figure 10-7. List of Active Line IDs

```
HSLA THRESHOLD REPORT FOR SYSTEM NMCC AT 08:15:00 ON 81-10-27      FIRST TAPE D0001

HSLA THRESHOLD ** FOR DNET 1 AT 08:25:00 HSLA 1 HSALSC 2 XMITS INCREASED BY 3000--TIME SINCE PREV RCD THIS NET = 240.00
                                                                                                                SECONDS

HSLA THRESHOLD ** FOR DNET 0 AT 09:15:20 HSLA 2 HSLASC 5 RCVS INCREASED BY 2700--TIME SINCE PREV RCD THIS NET = 240.00
                                                                                                                SECONDS
```


Figure 10-8. HSLA Threshold Report (Report Thresh)

This message indicates that the data is not being received in the order expected and should be considered as an error condition. CCTC (C751) should be contacted when this condition is found to exist. A similar problem exists if the word "UNMATCHED" is replaced by the word "DUPLICATE." Both these cases indicate data not being received in the order expected. Table 10-2 displays all such occurrences of these messages and reasons for the problem. This report will be generated only if the response time portion of the DATANET-355 GMF Monitor was active during data collection (see subsection 5.5.9).

10.5.1.9 Response Interval Matched Pairs Verification (Report MATCH). Data transmission between host and front end is governed by the series of events described in section 10.3. A matched pair is affirmed on detecting any sequence that generates a valid response time measurement as described in section 10.3, for the same line ID within a single T62 trace. This report will be generated only if the response time portion of the DATANET-355 GMF Monitor was active during data collection (see subsection 5.5.9).

Matched pairs are noted within the DATANET-355 Data Reduction Program, and are printed when affirmed, and when both the following two conditions are true: (1) the option has been specifically selected (see subsection 10.6.15) and (2) the response time threshold (separately specified for each response time category) has been exceeded. The report appears on file code 20. A typical line of output is shown below.

Example of matched pairs verification:

DN-355 DATA ACCEPTANCE TIME FOR 75 SECS ON DNET 1 FOR LNID DG TYPE = VIP 7855 FOR SUBSYSTEM VIDEO AT 08:08:12

This report may be used to capture the occurrence of response or processing times that are considered to be excessive.

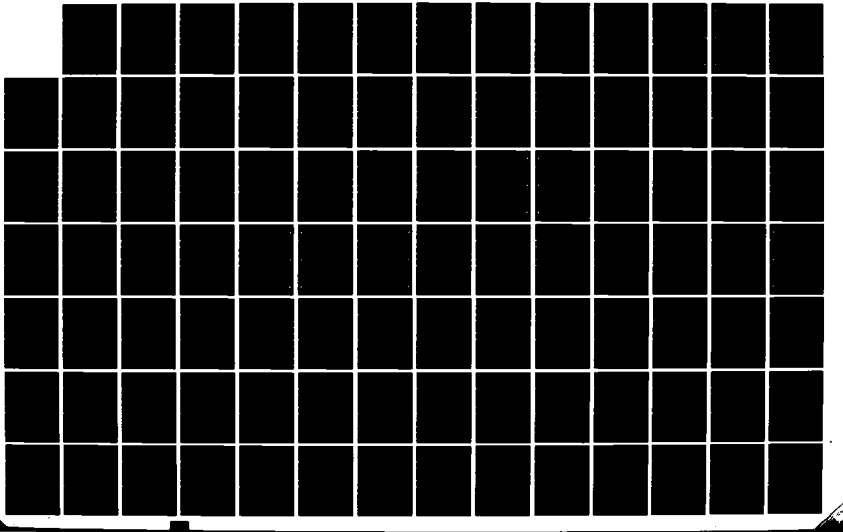
10.5.1.10 Annotated List of Datanet Traces (Report TRACE). This report can only be produced by use of the special input option TRACE (see subsection 10.6.12). It will produce voluminous output and should only be used for debug purposes. The report will produce a formatted dump of the various DATANET 355 GMF data records and is only meaningful to someone familiar with the correct record formats. No sample of this report is given in this document.

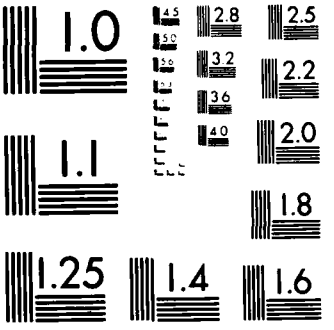
10.5.2 Plots. The plots described in the following subsections can be obtained from a data reduction run. Table 10-1 lists the data collections that can be plotted. Each plot contains the information for all known DATANETs. Plot listings are found on file code 25. Since the plots can produce a tremendous amount of output, it is suggested that the plot time interval not be made too small. The

Table 10-2. Response Time Error Messages

ERROR MESSAGE	REASON
Duplicate Send Output	A duplicate event 2 has been received before the occurrence of an event 3 or 4.
Duplicate Accept Direct Output or Duplicate Accept Direct Output/Input	A duplicate event 3 or 4 has been received before the occurrence of an event 8.
Unmatched Accept Direct Output or Unmatched Accept Direct Output/Input	An event 3 or 4 has occurred before an event 2 or 6.
Unmatched Accept Direct Input	An event 5 has occurred before an event 10.
Unmatched Input Accepted	An event 6 has occurred before an event 5.
Duplicate Input Accepted	A duplicate event 6 has been received before the occurrence of an event 3 or 4.
Duplicate Total System Response Time	A duplicate event 6 has occurred before an event 4.
Unmatched Output Received	An event 8 has occurred before an event 3 or 4.
Duplicate Output Received	A duplicate event 8 has occurred before an event 9.
Unmatched Output Started	An event 9 has occurred before an event 8.
Duplicate Output Started	A duplicate event 9 has occurred before an event 10.
Duplicate Output Completed	A duplicate event 10 has occurred before the occurrence of an event 5.
H6000 Rejected Accept Direct Input (THIS IS IS NOT AN ERROR CONDITION)	H6000 too busy to respond to DN 355, see subsection 10.6.22.

AD-A131 954 GENERALIZED MONITORING FACILITY USERS MANUAL CHANGE 6 6/8
(U) COMMAND AND CONTROL TECHNICAL CENTER WASHINGTON DC
B M WALLACK ET AL. 01 JUL 83 CCTC-CSM-UM-246-82-CHG-6
UNCLASSIFIED DCA/DF-83/001B F/G 9/2 NL





MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS 1963 A

default plot interval is five minutes and can be changed with a user input option (see subsection 10.6.14).

One horizontal line is output at the specified, or default, plot time interval. Every tenth line displays the current time of day. If a curve ordinate is on or beyond the axis limit, it will be positioned at the axis limit. If two or more curves share the same ordinate simultaneously, the position of coincidence will be marked with the cardinal number of the curves sharing it. The cell "height" (value of one increment along the horizontal axis) is specified by a DELTA value given in the heading. The vertical axis is the time axis with nominal increments of one plot time interval, but rarely will the increments be equal. The parameters which determine the plot delta intervals (horizontal increments) are described in subsection 10.4 and methods for altering the parameters are described in subsection 10.6.2. Each plot contains a heading line giving the date and the starting time of the collection, a title describing the collection, the identifying mark of each curve, and the DELTA value. When the plot time interval is exceeded for any given DATANET, all DATANETs will be plotted. Therefore, it is probable that in a multiple DATANET environment a given plot entry will reflect different amounts of time for each of the DATANETs. In order to make the plots more readable, it is suggested that only a single DATANET be plotted at one time. This will require multiple reductions to be run, turning off the plot option for different DATANETs, but output will be easier to read. The method for doing this is described in subsection 10.6.13.

10.5.2.1 Number of Transactions Sent to Host (Plot ID PLOT1). This report plots the cumulative number of transactions sent to the host, over a plot interval of time, by DATANET, as obtained from the host buffer for CPU utilization (record type 1, subtype 4). Figure 10-9 shows an example. In this plot we see that DATANET 0 has accumulated 3560 transactions between 0800, when the monitor started, and 0822, when the first plot entry was made. In this plot, each dash (delta value) is equivalent to 40 transactions, and therefore, the "A" is plotted at the 89th dash mark. During that same interval of time, DATANET 1 had accumulated 4400 transactions. From 0822 until the next plot interval of time (approximately 15 minutes, depending on when the data buffer was received), DATANET 0 again accumulated 3560 transactions, but DATANET 1 accumulated only 2520 transactions. Therefore, the B is written out at the 63rd mark. The plots need to be read as a continuous line of output. Therefore, the user will need to connect the same lettered points in order to track the change that has occurred. During the next time interval, DATANET 0 decreased to 2520 transactions, while DATANET 1 increased to 3560 transactions.

10.5.2.2 Number of Transactions Received From Host (Plot ID PLOT2). This report plots the cumulative number of transactions received from

DISTRIBUTION COLLECTED ON SYSTEM NMCC AT 8:00:00 ON 04-10-81

TRANSACTIONS SENT TO HOST (CUMULATIVE/15 MIN) PLOT-1

DELTA= 4.0000E 01 A=DNET 0 B=DNET 1 C=DNET 2 D=DNET 3

TIME OF DAY SCALED COUNTS

0 1240 2520 3760 4400

```
y---------y---------y---------y---------y---------y---------y---------y---------y---------y---------y
X  08:22;:00                                                                    A                   B
2                                                                  B            A
2                                                                  A            B
```

CURVE	MINIMUM	MAXIMUM	AVERAGE
DNET 0	0.2520E 04	0.3560E 04	0.3213E 04
DNET 1	0.2520E 04	0.4400E 04	0.3473E 04
DNET 2	0.	0.	0.
DNET 3	0.	0.	0.

Figure 10-9. Number of Transactions Sent to Host (Plot 1)

the host, over a plot interval of time, by DATANET, as obtained from the host buffer for CPU utilization (record type 1, subtype 5). This plot should be interpreted in the same manner as figure 10-9 and therefore no example is included.

10.5.2.3 Number of 36-Bit Words Sent to Host (Plot ID PLOT3). This report plots the cumulative number of 36-bit words sent to the host, during a plot interval of time, by DATANET, as obtained from the host buffer for CPU utilization (record type 1, subtype 6). This plot should be interpreted in the same manner as figure 10-9 and therefore no example is included.

10.5.2.4 Number of 36-Bit Words Received From Host (Plot ID PLOT4). This report plots the cumulative number of 36-bit words received from the host, during a plot interval of time, by DATANET, as obtained from the host buffer for CPU utilization (record type 1, subtype 7). Figure 10-10 provides a sample of this type plot and should be interpreted in the same manner as figure 10-9. Note the occurrence of the 2's as both Datanet A and B have reported identical values during the second and third time increments.

10.5.2.5 Percent Idle Time Over a Plot Interval (Plot ID PLOT5). This report plots the percent of time, over a plot interval, that each DATANET was idle. If a DATANET crashes during this interval, the amount of time spent down, due to the crash, is disregarded (i.e., the length of the interval for the DATANET that crashed will be reduced by the amount of time the DATANET was down). This information is obtained from the host buffer for CPU Utilization (record type 1, subtype 9). Figure 10-11 provides an example of this plot. In examining this plot, it can be seen that during the first time interval, DATANET 0 was 32 percent idle, while DATANET 1 was 30 percent idle. During the second time interval, DATANET 0 became 27 percent idle, while DATANET 1 remained 30 percent idle. During the third time interval, DATANET 0 had a 29 percent idle time, while DATANET 1 had a 27 percent idle time.

10.5.2.6 Percent Buffer Requests Denied (Plot ID PLOT6). This report plots the percent of buffer requests that were denied during a single plot interval of time, by DATANET, as obtained from the host buffer for CPU utilization (record type 1, subtype 1 and 10). As can be seen in the figure, two DATANETs are being monitored for three time intervals. Both DATANETs are reporting 60 percent of their buffer requests being denied. Since two DATANETs are reporting the same value, the number "2" appears on the plot, indicating the intersection of two values. Refer to figure 10-12 for a sample of this type plot.

10.5.2.7 Number of 18-Bit Words Available For Buffers (Plot ID PLOT7). This report plots the average number of 18-bit words vailable for buffers, over a plot interval of time, by DATANET, as

DISTRIBUTION COLLECTED ON SYSTEM NMCC AT 8:00:00 ON 04-10-81

36-BIT WORDS RECEIVED FROM HOST (CUMULATIVE/15 MIN) PLOT-4

DELTA= 4.0000E 01 A=DNET 0 B=DNET 1 C=DNET 2 D=DNET 3

TIME OF DAY SCALED COUNTS

0 1240 2520 3760 4400



X 08:22:00 A B
2 2
2 2

CURVE	MINIMUM	MAXIMUM	AVERAGE
DNET 0	0.3760E 04	0.3760E 04	0.3760E 04
DNET 1	0.3760E 04	0.4360E 04	0.3950E 04
DNET 2	0.	0.	0.
DNET 3	0.	0.	0.

Figure 10-10. Number of 36 Bit Words Received From Host (Plot 4)

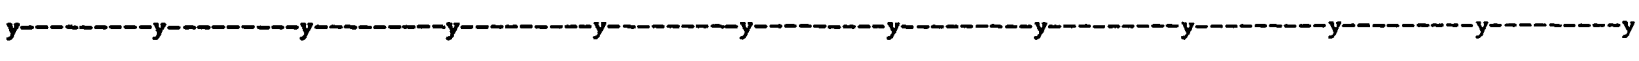
DISTRIBUTION COLLECTED ON SYSTEM NMCC AT 8:00:00 ON 04-10-81

% IDLE TIME/15 MINUTES PLOT-5

DELTA= 1.0000E 00 A=DNET 0 B=DNET 1 C=DNET 2 D=DNET 3

TIME OF DAY SCALED COUNTS

0 31 63 94 110



X 08:22:00 B A
2 A B
2 B A

CURVE	MINIMUM	MAXIMUM	AVERAGE
DNET 0	0.2700E 02	0.3200E 02	0.2933E 02
DNET 1	0.2700E 02	0.3000E 02	0.2900E 02
DNET 2	0.	0.	0.
DNET 3	0.	0.	0.

Figure 10-11. Percent Idle Time (Plot 5)

DISTRIBUTION COLLECTED ON SYSTEM NMCC AT 8:00:00 ON 04-10-81

% BUFFER REQUESTED DENIED/15 MIN PLOT-6

DELTA= 1.0000E 00 A=DNET 0 B=DNET 1 C=DNET 2 D=DNET 3

TIME OF DAY SCALED COUNTS

0 31 63 94 110

y---------y---------y---------y---------y---------y---------y---------y---------y---------y---------y

X 08:22:00 2
2 2
2 2

CURVE	MINIMUM	MAXIMUM	AVERAGE
DNET 0	0.6000E 02	0.60000E 02	0.60000E 02
DNET 1	0.6000E 02	0.60000E 02	0.60000E 02
DNET 2	0.	0.	0.
DNET 3	0.	0.	0.

Figure 10-12. Percent Buffer Requests Denied (Plot 6)

obtained from the host buffer for CPU utilization (record type 1, subtype 2). This plot is interpreted in the same manner as other plots and therefore, an example is not provided.

10.5.2.8 Number of Users Logged On The System (Plot ID PLOT8). This report plots the average number of users logged on each DATANET, during a plot interval of time, as obtained from the host buffer for CPU utilization (record type 1, subtype 3). This plot is interpreted in the same manner as other plots and therefore, an example is not provided.

10.5.2.9 Number of Host RSVPs Received (Plot ID PLOT9). This report plots the cumulative number of host RSVPs received, during a plot interval of time, by DATANET, as obtained from the host buffer for CPU utilization (record type 1, subtype 8). This plot is interpreted in the same manner as other plots and therefore, an example is not provided.

10.5.3 Histograms. Histogram reports are produced for each response time category described in section 10.3. Histograms can be produced for all known DATANETs (up to four) and special line IDs (up to 10). Included on each listing are the time and date the data were collected, the system name on which they were collected, response category, and the report ID. The histogram proper includes cells count and frequency, both individual and cumulative, as well as the range and scale of each cell. A bar chart comparing the cell frequencies, and a statistical summary of in-range and out-of-range events completes the chart. Figures 10-13 and 10-14 show sample histograms as produced by a DATANET and a special line ID. Histograms will be outputted at the end of a timeframe or at the end of processing. A closing line on each histogram shows the TODs between which the histogram data are valid, and why the histogram was printed. Successive histograms for the same application are never cumulative.

Figure 10-13 will be used to describe the procedure that should be followed in interpreting a histogram. Entries in the column headed SECOND give the range of times which form each histogram bucket. The first bucket is used to report 0 measured values, if any, with all other buckets used to represent a range of values. The user should refer to subsection 10.4 for a discussion of range values within histograms and subsection 10.6.20 for the method of altering the standard ranges. The entries in the column headed INDIV. NUMBER give the number of reponses that were measured that fell within the time range specified by the SECOND column. In figure 10-13, we observe that 144 responses measured between 2 and 3 seconds.

DISTRIBUTION COLLECTED ON SYSTEM NMCC AT 8:08:00 ON 04-11-81

H6000 MACHINE RESPONSE TIME FOR DATANET 1

INDIV. NUMBER	CUMUL. NUMBER	CUMUL. PRC.	INDIV. PRC.	SECOND	PERCENT OF OCCURRENCE 00 10 20 30 40 50 60 70 80 90 100
					I....I....I....I....I....I....I....I....I....I....I
135	135	0.484	0.484	1- 1	Ixxxxxxxxxxxxxxxxxxxxxxx
144	279	1.000	0.516	2- 3	Ixxxxxxxxxxxxxxxxxxxxxxxxxx

279 ENTRIES TOTAL AVERAGE = 1.51613 VARIANCE = 0.250 STANDARD DEVIATION = 0.500

TRIGGERED BY TIME - SPAN ENDED AT 08:26:00
DNET 1 SEEN FROM 08:10:00 TO 08:22:08

Figure 10-13. Histogram of DATANET Response

DISTRIBUTION COLLECTED ON SYSTEM NMCC AT 00:00 ON 04-11-81

H6000 MACHINE RESPONSE TIME FOR LNID DC

INDIV. NUMBER	CUMUL. NUMBER	CUMUL. PRC.	INDIV. PRC.	SECOND	PERCENT OF OCCURRENCE 00 10 20 30 40 50 60 70 80 90 100
					I....I....I....I....I....I....I....I....I....I....I
195	195	0.484	0.485	1- 1	Ixxxxxxxxxxxxxxxxxxxxxxx
208	403	1.000	0.516	2- 3	Ixxxxxxxxxxxxxxxxxxxxxxx

403 ENTRIES TOTAL AVERAGE = 1.51313 VARIANCE = 0.250 STANDARD DEVIATION = 0.500

TRIGGERED BY END-OF-FILE AT 09:00:00
DNET 0 SEEN FROM 08:02:00 TO 09:00:00

Figure 10-14. Histogram Special Line ID Response

Similarly, the columns headed INDIV. PRC and CUMUL. PRC give the individual and cumulative percentages of all responses which were made within each time range. In figure 10-13, we observe that 48% of all responses were 1 second in duration, while 51% were between 2 and 3 seconds in duration. At the same time, it can be observed (CUMUL. PRC) that 100% of all responses were measured at 3 seconds or less.

10.6 Default Option Alteration

The DDRP uses the trace tape generated by the GMC and user data cards as input. The user has several optional inputs. These options are evoked by specifying an input option (action code) and any other required inputs specified in the following subsections. The general format for an option request is as follows:

The first card is an action code describing the action to be taken. Subsequent cards modify report parameters for some of the action codes. All inputs are free format with the only requirement being that all zeros must be typed on the data card. At least one data card with the word END specified is required as the last data card.

Available action codes (and default implications) (and character code) are:

1. Turn histogram on; (no histograms) (HISTG)

2. Modify a plot; (standard plots) (PLOT)

3. Turn a specific report on; (all reports on except ID TRACE, MATCH, and histograms) (ON)

4. Turn a specific report off; (all reports on except ID TRACE MATCH, and histograms) (OFF)

5. Set a timespan for measurement; (time not a criterion for measurement) (TIME)

6. Process Data Reduction Program on a WW6.4/2H system (WW7.2/4JS processing) (RN)

7. Turn all reports off except those specified; (all reports on except ID TRACE, MATCH, and histograms) (ALLOFF)

8. Turn all reports on except those specified; (all reports on except ID TRACE, MATCH, and histograms) (ALLON)

9. Do not stop on an option request error; (stop on an input error) (ERROR)

10. Number of physical tape records to process before stopping; (number of records not a constraint) (NREC)

11. All reports off except plots; (all reports on except ID TRACE, MATCH, and histograms) (REPORT)

12. Dump formatted record types listed; (no dump) (TRACE)

13. DATANETs not to analyze for plots; (all are plotted) (NOPLOT)

14. Set minimum plot interval; (five minute plot interval) (INTERV)

15. Verify response time records (type 2) matched (at end of processing of each T62 trace) (no verification) (MATCH)

16. Accept line IDs for special analysis; (none get special analysis) (SPECL)

17. Modify threshold values (all values 9999999) (THRESH)

18. DATANETs not to analyze for reports other than plots; (all nets are analyzed) (NORPT)

19. Debug (no debug) (DEBUG)

20. Modify parameters for histograms (default values described in section 10.4) (CHANGE)

21. Error Report (all unmatched responses produced) (DUPLIC)

22. Reject Report (reject messages are produced but summary is not) (REJECT)

There is no order required for the options, and multiple entries in each are permissible. If several inputs refer to the same report, the last one encountered will have precedence. Also, if a report is off by default and is modified, it will be turned on through its request for modification. All input cards are free format unless otherwise described. If multiple parameters occur on a given card they must be separated by at least 1 blank column.

Three tabular reports depart from the normal scheme (see table 10-1): the reports titled "List of Active Line IDs" and "Response Interval Unmatched Pairs Verification" are generated if a T62 trace is processed and cannot be turned off. The report titled "Response Reject Messages" is also produced by default and can only be altered with the REJECT user option.

10.6.1 Histogram Desired (Action Code HISTG). If the user desires histograms, he must use this action code. There is no other method for obtaining histograms. The user may obtain histograms for either DATANETs or special terminals. This option would be used for obtaining DATANET histograms. Terminal histograms can be requested under action code "SPECL". For every requested DATANET and/or terminal, a set of eight histograms will be produced. Currently, the data reduction program can produce a maximum of 40 histograms (five sets of eight histograms each). If more than this are desired, the user will need to edit and recompile the source code. In performing the edit, the user should change the value of "RPTCNT" from the current setting of 57 to the required new value. To calculate the new value of "RPTCNT," the following formula should be used:

RPTCNT = 17 + 8* (number of sets of histograms desired)

It should be noted that for every integer increase in the value of RPTCNT, the amount of memory required to run this program increases by 80 words. The first card contains the word HISTG and the second card contains the DATANET number (0-3) for which the histograms are desired.

10.6.2 Plot Alteration (Action Code PLOT). For any given plot the user can alter the maximum number of lines to be printed on the plot, the minimum value of the plot and the maximum value of the plot (see subsection 10.4). By specifying the minimum and maximum values of the plot, the user has in effect specified the interval size of the plot. If a value to be plotted is smaller then the minimum value specified, the element will be plotted so as to intersect the left axis of the plot. If a value to be plotted is larger then the maximum value specified, the element will be plotted as an asterisk on the right hand axis of the plot. If the user wants to insure that the entire plot is printed, he should input a value of -1 for the number of lines to be printed. The minimum and maximum values must be specified as real numbers, i.e. decimal points must be punched on the data card. Table 10-1 lists those measurements currently being plotted and figure 10-1 lists the default values for all plots. Figure 10-15 shows the format required for this action code.

10.6.3 Turn a Report On (Action Code ON). This action allows a user to turn a single report on (that is off by default) for this run. The user specifies the report ID (see table 10-1). No change to the default parameters will be made, and a report already on will remain on. All reports are on by default except the histograms, and report ID's TRACE and MATCH. None of these can be turned on by using this option. To turn histograms on, the user must use the "HISTG" or "SPECL" options, and to turn on the other two reports, the user must use the "TRACE" and "MATCH" options respectively. This option is only included for completeness and will never need to be used. The format

Card #1 A
Card #2 B C D E

Where

A = The letters PLOT
B = Plot ID to be altered (see TABLE 10-1)
C = Number of lines to be printed or a -1 if entire plot is desired
D = Maximum value for plot inputted as a real number, i.e. with decimal point
E = Minimum value for plot inputted as a real number, i.e. with decimal point

Figure 10-15. Plot Alteration (Action Code-PLOT)

consists of two cards. The first contains the word ON and the second card contains the report ID from table 10-1 to be turned on.

10.6.4 Turn a Report Off (Action Code OFF). This action permits a user to turn a single report off (that is on by default) for this run. The user specifies the report ID (see table 10-1). No change of default parameters is made, and a report may be turned off multiple times. All reports and plots that are currently on by default can be turned off with this option except for the report titled "Response Reject Messages." This report can be altered only with the use of the input option REJECT. The format for this option is the same as the ON option except the word OFF replaces the word ON.

10.6.5 Set a Timespan of Measurement (Action Code TIME). The system allows a user to specify the timespan (or spans) in which the collected data is to be processed and reports generated. For example, the user may collect data from 0500 to 2200 hours and specify that he wants to process only 0900 to 1700 hours on all reports. The timespans set through this option apply to all report processing.

Up to five timespans may be specified and they must be chronologically ordered. In specifying timespans a start time and a stop time must be supplied. If a start time is specified, but no stop time is specified, there should be no data on the card after the start time is specified. All times are expressed as S1M1 where S1 is the start hour and M1 is the start minute. All times must be expressed as 4 character fields with no intervening blanks. If the user wants to request the time 4:07 he must input a "0407". Times are considered to be on a 24 hours clock.

Individual timespans may be set for the entire data reduction run (report ID=TOTAL), or for report ID's THRESH and TRACE, or for any plot. Report ID's for plots are indicated in TABLE 10-1. No other individual timespans may be set. The format for this card is shown in figure 10-16.

10.6.6 Process Reduction on a WW6.4/2H System (Action Code RN). The data reduction program is designed to be processed on a WW7.2/4JS GCOS release. However, the program can be executed under a WW6.4/2H GCOS system with the use of this option. The option consists of an RN typed on the data card.

10.6.7 Turn All Reports Off Except Those Specified (Action Code ALLOFF). All reports except those explicitly identified here are to be turned off. The card format is shown in figure 10-17.

10.6.8 Turn all Reports On Except Those Specified (Action Code ALLON). All reports except those explicitly identified here are to be turned on. The card format is shown in figure 10-18. This action

```
Card 1      A
Card 2      N     M
Card 3      B     C     D     E     F .....
```

Where

A = The letters TIME
N = Report ID to be timespanned. (see subsection 10.6.5)
M = Number of different times appearing on card 3
B, C, D, E, F = Start and stop times used to define the timespan. (see subsection 10.6.5)

Figure 10-16. Set Timespan(s) For Processing (Action Code TIME)

```
Card 1      A
Card 2      N
Card 3      I    J    K . . .
```

Where

A = The letters ALLOFF
N = The number of reports to be specified on card 3
I, J, K = Report ID's (from TABLE 10-1)
not to be turned off. N of these numbers
must appear on this card.

Figure 10-17. Action Codes For Turning Reports Off

```
Card 1      A
Card 2      N
Card 3      I    J    K . . . .
```

Where

A = The letters ALLON
N = The number of reports to be specified on card 3
I, J, K = Report ID's (from TABLE 10-1) not to be turned on. N of these numbers must appear on this card.

Figure 10-18. Action Codes For Turning Reports ON

cannot be used to turn on histograms. Action codes "HISTG" and "SPECL" must be used for histograms. It also should not be used to turn on report ID's MATCH and TRACE. These two reports are turned via special input options. For these reasons the user should never use this option and it is provided only for a matter of completeness.

10.6.9 Continue Data Reduction After an Input Option Error (Action Code ERROR). This code allows data reduction to continue when an error has been detected in an input option request. The default value will abort data reduction and report the error. No cards are required after the action card which contains the word ERROR.

10.6.10 Stop Processing (Action Code NREC). This option permits a user to specify the number of tape records to process before stopping. This is useful if a hard tape error aborts the run and causes reports to be lost. In order to obtain these reports, it is necessary to stop processing prior to the tape error. The first card contains the word NREC and the second card contains the number of tape records to be processed.

10.6.11 Plots Only (Action Code REPORT). This action code permits the plots to be generated and curtails all other reports. No cards are required after the action card which contains the word REPORT.

10.6.12 Print formatted DATANET-355 Record Types (Action Code TRACE). The GRTS-II monitor of the GMC distinguishes three record types, designated 1, 2, and 3. Generally, type 1 records provide gross net performance data; type 2 records capture events pertinent to each line ID; type 3 records show gross performance of each HSLA subchannel. This Action code permits the user to specify which type record to format and print. An additional option allows selection of a header showing various counters and clocks.

The first card contains the word TRACE. The second card would indicate which of the possible options the user is going to describe. This card can contain the word TERM, or LIMIT, or RECTYP, or END.

If the word TERM appears on the second card, then the third card would contain the number of terminals (maximum of four) for which the type two records are to be formatted. The fourth card would contain the list of terminal IDs for which type two record formatting is desired. The IDs should be separated by at least one blank.

The next card would once again contain the word LIMIT, or RECTYP, or END.

If the word LIMIT appears on the card then the following card would contain the number of GRTS records that should be formatted before the formatting option will be deactivated. The card following would once again contain the word RECTYP, or END.

If the word RECTYP appears on the card then the following card would contain the number of record types desired and the DATANET for which they are desired. A zero for a DATANET # must be typed. A DATANET number of -1 indicates that all DATANETs should be processed in the requested manner. The following card would list the actual record types desired. The card following should contain the word END.

If the word END appears on the card processing of the dump option will terminate. Any combination of the above options may follow a TRACE action card, but the END card must be the last card of the group. The TRACE report is off by default and will be activated by the processing of this option.

10.6.13 DATANETs Not to Plot (Action Code NOPLOT). This option permits the user to "turn off" plot data for specified DATANETs. The effect on an active plot is to zero the Y-coordinate for the unwanted DATANET. The DATANET identifiers are 00 through 03. This option should be used when reducing a tape containing data from multiple DATANETs. Under the default conditions, data from every DATANET will appear on the same plot, as different curves. While this allows the user to easily compare the activity on multiple DATANETs it does present several problems. One problem is that the plot tends to be difficult to read because of the large amount of data contained in a small area. The other problem is one of plot limits. If the activity on the various DATANETs is significantly different, it becomes difficult to set meaningful minimum and maximum plot boundaries.

By turning off plots for specified DATANETs, the user can see each DATANET plotted separately, or else only see certain DATANETs plotted together. If the user selects this option, multiple reductions will need to be run in order to see the plots for all DATANETs. Figure 10-19 shows the format for this action.

10.6.14 Set a Plot Interval (Action Code INTERV). This option allows the user to set the minimum plot interval. The interval criterion is applied to each DATANET individually and if the test approves the plot, the latest data from all DATANETs are plotted. The default interval is 5 minutes (300 seconds). The first card contains the word INTERV and the second card contains the new plot interval in seconds.

10.6.15 Response Time Records (Type 2) Matched (Action Code MATCH). This option causes a message to be printed whenever a response interval is completed (matched). This option can produce voluminous output and should be used with caution if the data tape contains more than a few records. The "matched" option can be used to track the origin of unusual response times reported in DATANET and special-analysis line ID histograms. A series of seven threshold

Card 1 A
Card 2 N
Card 3 M L O....

Where

A = The letters NOPLOT
N = Number of DATANETs appearing on the third card, whose plots will be turned off
M, L, O = DATANET numbers whose plots are to be turned off. N such numbers must appear on card.
If the DATANET number is 0, the 0 must be punched on the card.

Figure 10-19. Turn Off Plots (Action Code NOPLOT)

values can be specified to force the tracking of times greater than, or equal to, the threshold value. The definition of the various response times are given in sections 5.2.7.7.3, 5.4.8, and 10.3 of this document. The first card contains the word MATCH. The second card would contain the seven new threshold values that are desired. Seven numbers must appear on this card. The threshold must be expressed in milliseconds. This report is not produced under default conditions and will be turned on by the occurrence of this option.

10.6.16 Accept Line IDs for Special Analysis (Action Code SPECL). Up to 10 line IDs can be designated for special analysis on a terminal session basis, i.e., from line connect to line disconnect. Each session is summarized on the tabular report ID RESPL.

It is also possible to obtain histograms for these special lines. In this case, both tabular reports and histograms would be produced. If the histogram option is selected, the user must insure that the value of "RPTCNT" is set correctly (see subsection 10.6.1). If the user wants more than 10 line IDs investigated, the internal parameter "MXSPEC" can be altered, via an edit, and the program recompiled. All occurrences of "MXSPEC" must be changed. For each integer increase in "MXSPEC", and additional 30 words of memory is required. The first card contains the word SPECL. The second card would contain the number of special terminals desired and the third card would contain a list of terminal IDs separated by at least one blank.

The terminal IDs should be expressed as a 2 character field followed immediately by an H or N. The "H" signifies that histograms are desired, while the "N" signifies that no histograms are desired.

10.6.17 Modify Threshold Values (Action Code - THRESH). This option causes a warning message to be issued whenever the event count corresponding to any HSLA subchannel exceeds the specified threshold value since the previous trace for that DATANET. Two threshold values can be specified. They correspond to the following event counts on the HSLA subchannel: transmits and receives. The default values for these counts are initially set high enough that the warning messages should not be issued. In addition, the user may request that a message be produced every time the number of buffer denials exceeds a certain value or the number of available words of buffer space falls below a certain value. This information is provided only when plots are being produced. Figure 10-20 shows the cards required for this action.

10.6.18 Suppress Nonplot Reports for Selected Nets (Action Code - NORPT). This option allows the turning off of all reports (other than plots) for a specified DATANET. Figure 10-21 shows the format for this action.

```
Card 1      A
Card 2      B
Card 3      C     D
```

Where

A = The letters THRESH
B = The letters AVAIL if we are setting buffer space availability; the letters DENIAL if we are setting buffer denial threshold; or the letters HSLA if we are setting HSLA activity thresholds.
C, D = Actual threshold values. D will be blank if card 2 was an "AVAIL" or a "DENIAL" card.

Figure 10-20. Threshold Values (Action Code THRESH)

```
Card 1      A
Card 2      N
Card 3      M    L    O ...
```

Where

A = The letters NORPT
N = Number of DATANETs appearing on the Third card, whose reports will not be produced
M, L, O = DATANET numbers whose reports are to be turned off. N such numbers must appear on card. If the DATANET number is 0, the 0 must be punched on the card.

Figure 10-21. Turn off Reports For Selected DATANETs (Action Code NORPT)

10.6.19 Debug (Action Code DEBUG). Extensive debug facilities are provided within the program. By default all debug is turned off. This option should be used only by somebody very familiar with the logic of the program. Figure 10-22 shows the format for this action.

10.6.20 Alter Histogram Default Parameters (Action Code CHANGE). The user should reference section 10.4 for a discussion of Histogram Parameters. Figure 10-23 provides the format for this option.

10.6.21 Unmatched/Duplicate Error Report (Action Code DUPLIC). In processing response time data, it has been discovered that responses may not be received in an order considered correct by the data reduction program. This can occur because of a lost data occurrence or because of an anomoly in the functioning of the DATANET-355. The correct order and an explanation of the responses being captured are presented in sections 5.2.7.7.3, 5.4.8, and 10.3 of this document. Normally, when this error occurs an error message will be produced. Table 10-2 shows the various error messages that may be produced and describes when they would occur. This option allows the user to turn off these error messages, or else produce the error messages only for those terminals selected for special analysis (action code SPECL). Figure 10-24 provides the format for this option.

10.6.22 Special Reject Report (Action Code REJECT). Instances occur when the H6000 will reject a DATANET-355 request for service. This rejection normally occurs when the H6000 is to busy to immediately respond to the 355 interrupt. In these cases the DATANET will be forced to retransmit its request. When this event is detected, the data reduction program will produce a reject error message. In addition, a summary of the number of such rejects may also be produced. Under default conditions, this summary report is not produced. This option allows the user to turn off the error messages or else to request the summary report. Figure 10-25 provides the format for this option.

10.7 JCL

The following paragraphs describe the JCL usage necessary for reducing a GRTS-II monitor data tape. The usual $ SNUMB, $ IDENT, and $ USERID cards are included as required for the installation. The program is controlled by the following JCL cards:

```
$ LOWLOAD
$ OPTION FORTRAN
$ SELECT B29IDPXO/OJBECT/GRT
$ LIMITS 20,48K,-2K,10000
```

```
Card 1      A
Card 2      B    C
```

Where

A = The letters DEBUG
B = The letters TYPE1 if type 1 and type 3 records are to be debugged; or the letters TYPE2 if type 2 records are to be debugged; or the letters BOTH if all records are to be debugged.
C = DATANET number to be debugged

Figure 10-22. Debug Option (Action Code DEBUG)

```
Card 1      A
Card 2      B     C     D
```

Where

A = The letters CHANGE
B = The letters LOW if the low value is modified or the letters INTEV if the interval size is modified.
C = Histogram ID to be modified (see table 10-1)
D = New parameter

Figure 10-23. Alter Histogram Parameters (Action Code CHANGE)

```
Card 1      A
Card 2      B
```

Where

A = The letters DUPLIC
B = The letters Off if the error messages are not desired at all, or the letters TERM if the error messages are desired only for special terminals

Figure 10-24. Unmatched/Duplicate Error Report (Action Code DUPLIC)

Card 1 A
Card 2 B

Where

A = The letters REJECT
B = The letters OFF will turn off the reject error message but produce the summary reports. The letters BOTH will cause the error messages to be produced as well as the summary report.

Figure 10-25. Reject Report
(Action Code REJECT)

The run time, 0.2 hour, and SYSOUT limits, 10000 lines, may be altered as required by the scope of the monitored run. During the loading process, the program will reach a memory size of 56K, but the extra 8K will be released immediately upon loading.

The GRTS-II monitor data tape specifications are provided by the following card:

```
$ TAPE 01,X1D,,tape.no,,NO-RING
```

The TAPE card specifies the tape type (7- or 9-track) and the tape number.

The input options file is identified by a

```
$ DATA I*
```

card, followed by the option cards as described in section 10.6. At least one data card containing an "END" must be present.

10.8 Multireel Processing

If more than a single reel of data has been collected a series of messages will be printed at the computer console informing the operator that a new data reel is required. The following are the messages produced.

a. DISMOUNT REEL #XXXXX THEN MOUNT REEL NUMBER YYYYY ON ZZZZZ

   In this case, XXXXX is the old reel number, YYYYY is the new reel number, and ZZZZZ is the tape drive ID.

   If the operator fails to mount the new tape, the above message will be repeated three times, after which the program will terminate, and all reports produced.

b. IS TAPE XXXXX MOUNTED ON DRIVE ID YYYYY (Y/N)

   In this case, XXXXX is the tape number being requested for mounting and YYYYY is the tape drive ID.

   This message occurs when the data reduction program finds the wrong tape has been mounted (by comparing internally generated tape labels). If the operator answers N, the message in (c) below is produced. If the operator answers Y, the data reduction program will terminate and all reports will be produced. In this case, the data reduction program is unable to process the tape. Even though the operator is mounting the

correct tape, the internal label on the new tape does not match that being requested by the old tape. The user should check the data collection session to insure that the operator did not respond with an incorrect tape number during multireel change.

After entering the Y or N, the operator will need to hit the EOM key twice in order for the response to be transmitted.

c. WRONG REEL JUST MOUNTED, DISMOUNT AND MOUNT REEL XXXXX ON ZZZZZ

In this case, XXXXX is the new reel number, and ZZZZZ is the tape drive ID.

d. CAN TAPE XXXXX BE MOUNTED ON DRIVE YYYYY (Y/N)

In this case, XXXXX is the new desired reel and YYYYY is the tape drive ID.

If the operator fails to answer this message it will be repeated until he responds with a "Y" for YES or "N" for NO. If he types in "Y", then message (a) will be repeated. If the types in "N", then the program will be terminated and all reports will be produced.

10.9 Tape Error Aborts

During the course of processing it is possible that the operator will be required to abort the data reduction program due to an irrecoverable tape error. If such a condition occurs, the operator should abort the job with a "U" abort. This will allow the data reduction program to enter its wrap-up code processing and produce all reports generated prior to the tape error.

SECTION 11. CENTRAL PROCESSING UNIT AND TAPE REDUCTION PROGRAM (CPUTRP)

11.1 Introduction

The Central Processing Unit and Tape Reduction Program (CPUTRP) is a FORTRAN program that sequentially processes the data recorded on tape by the CPU Monitor and the Tape Monitor of the General Monitor Collector (GMC). It produces a number of reports depicting the usage of the CPUs or the usage of the tape subsystem during the monitoring period. A list of these reports is found in table 11-1 and report descriptions are presented in subsection 11.5.

There are two inputs to the CPUTRP. The first is the data tape produced by the CPU Monitor and/or Tape Monitor in the General Monitor Collector. The second input is a set of report option control cards used to alter the reports in some way other than standard default. The various user input options and their formats are described in subsection 11.6. The actual reports produced by the CPUTRP are described in subsection 11.5.

11.2 Data Collection Methodology

The CPUM in the General Monitor Collector will create its own direct transfer traces (63 and 70) in order to collect data sufficient to analyze the utilization of the Central Processing Units. The method for generating these direct transfer traces is described in subsection 5.2.3 and the formats for the CPUM generated records used by the CPUTRP are described in subsection 5.4.4.

The Tape Monitor in the General Monitor Collector processes GCOS trace types 50, 51 and 52, and collects information to monitor the usage of the tape subsystem. The information collected on the occurrence of the above traces enables the CPUTRP to identify the jobs using tapes, which drives are used, how long a job is delayed due to nonavailability of tape drives, and the length of time a job is allocated to a given drive.

11.3 Analytical Methodology

There is no special analytical procedures used in this program. The program merely reports on the usage of the CPU and tape subsystem as it is reported by the General Monitor Collector.

11.4 Data Reduction Methodology

The CPUTRP is the only reduction program that actually produces reports from the data gathered by two different monitors (CPU and Tape Monitors). Therefore, a procedure needed to be developed whereby the user could specify which set of reports he desired, or if desired, allow him to produce the set of reports from each monitor. This capability could result in a given set of tapes needing to be analyzed as many as two times

within a given run. In order to manage this dual monitor reporting capability, a more complex user interface was required, then previously used by the other data reduction programs. This interface will be described in subsection 11.6.

11.5 CPUTRP Output

Output is divided, conceptually, into three parts - Execution and Error Reports, CPU Reduction Reports, and Tape Reduction Reports. These categories are shown in table 11-1, and further described below. An example of each report is displayed, and a simple explanation of how it was derived from the data is given.

11.5.1 Execution and Error Reports (Files 6, 7 and 8). These reports are written to three files: codes 6, 7 and 8. On file code 8, the user will find a printout containing processing or execution information - what the CPU and Tape Reduction Program found on the data tapes. Included in this information is the following:

- o A list of the monitors in execution during the GMC data collection.
- o An echo of the user input options, and the program's interpretation of them.
- o The time, date, tape number and RCSR clock at the beginning of collection.
- o If the timeframe option is used, a report of when the various timeframes were reached.
- o A count of the traces reduced within each timeframe option.
- o If a reduction period exceeded 9+ hours, an indication that a 35-bit internal clock overflowed (no error implied).
- o The time, date, tape number and RSCR clock at end of data collection, but only if a termination trace is processed.
- o When the NEW option is user selected, the items above are repeated on a new page.
- o Version number - The software that is compatible with this documentation should indicate VERSION 7.2-11.74, 24 SEP 1982.

A sample of this report is shown in figure 11-1.

Errors, except for tape handler errors, are listed on file code 6. Any such occurrence is of concern to C751 and should be reported for

Table 11-1. Central Processing Unit and Tape Reduction Reports

CPU Reports:

a. A periodic tabular report which shows the cumulative CPU utilization and queuing by various program categories (default period is 10 minutes) (ID #0).

b. An integer valued histogram of the number of activities in the CPU queue (ID #3).

c. A real valued histogram of the CPU burst length for user activities (ID #1).

d. A real valued histogram of the CPU burst length for all activities (ID #2).

e. An integer valued histogram displaying the percentage of time a given activity spent waiting in the dispatcher's queue (ID #8).

f. A periodic plot of the CPU idle time (default period is 10 minutes) (ID #4).

g. A periodic tabular report of WIN CPU usage (default period is 10 minutes) (ID #5).

h. A report showing the CPU usage and queuing statistics for each activity processed (ID #6).

Tape Reports:

i. An integer valued histogram of the number of tape drives in use (time corrected) (ID #7).

j. A tabular report which describes, for each activity that used tapes, the device and channel utilization and the delay time waiting drives (no report ID # assigned).

Execution Reports:

k. A report which provides an overview of the data reduction - it describes the initial and final data tapes, the card input options and user selected time frames as they occur (no report ID # assigned).

l. A report which shows errors detected by the tape handler, and, if selected, debug output (no report ID # assigned).

m. A report which shows errors detected by program modules other than the tape handler (no report ID # assigned).

```
CPU AND TAPE REDUCTION, VERSION 7.2 - 11.74, 24 Sep 1982

     ACTIVE MONITORS:  ****  ****  CPUM  TM    CM    ****  ****  ****  ****  ****  ****  ****

     DATA COLLECTION BEGAN WITH TAPE D0020 AT 11:44:41 ON 82-12-15,
     RSCRA/Q WERE  000000000000 731542120046

CARD INPUT IS INTERPRETED AS FOLLOWS:

 1: HPRINT(1) CPRINT(300)
     HISTOGRAM REPETITION PERIOD IS 1          SECONDS MINIMUM.
     CPU REPORT REPETITION PERIOD IS NOMINALLY 300         SECONDS.

 2: NEW(D0020)
     TAPE D0020  WILL LEAD OFF A FOLLOWING REDUCTION.

      REDUCED          275  T70   TRACES.

     DATA COLLECTION ENDED WITH TAPE D0020 AT 11:56:48 ON 82-12-15
     RSCRA/Q WERE  000000000000 737066771640
```

Figure 11-1. Sample Execution Report

interpretation and corrective action (202-695-0856). The occurrence of errors will, as a rule, invalidate the run.

All tape handler messages, including tape error messages, are found on file code 7. The following messages are of informative value and are generally self explanatory.

- "MOUNTING ANOTHER REEL #..."
- "END REACHED IN NXTREC AT RCDCNT ..."
  Explanation: A tape or operator error forced the tape handler to treat the condition as an end-of-file. This message may be preceded by the following
- "ERROR ON THE TAPE READ"
- "TAPE MOUNTED APPEARS INCORRECT BUT IS NEW ... PROCESSING WILL CONTINUE ..."
  Explanation: A newly mounted continuation reel does not conform to continuation conventions, but is also not positively incorrect. The run is probably invalid.
- "INCORRECT TAPE MOUNTED 3 TIMES OR NEXT TAPE CANNOT BE READ. RUN ENDED".
  Explanation: A newly mounted continuation reel is positively incorrect. This will be preceded by
- "WRONG TAPE MOUNTED. WANTED ..."
- "JULIAN DATES DO NOT AGREE. RUN ENDED"
  Explanation: A new physical record bears an incorrect Julian date. This event may result from re-using an old GMC data tape in another collection run in which GMC termination was improper or incomplete.
- "GMFEXC FOUND END OF TAPE ON FIRST READ"
- "GMFEXC ERROR: FIRST RECORD IS IN ERROR OR WRONG TAPE MOUNTED. PHYSICAL RECORD DUMPED".
  Explanation: The occurrence of this message is of concern to C751 only if there was no tape mounting or tape number confusion. This may be preceded by
- "GMFEXC ERROR: ASKED FOR TAPE #... GOT ..."
  Explanation: The NEW option was used, and the user specified tape number does not agree with the number found on the NEW tape.

All other tape handler messages are of concern to C751 and should be reported for correction.

Debug output, when selected, shares file code 7 with the tape handler.

11.5.2 Central Processing Unit Monitor Reports. The CPU title page is printed to file code 14, immediately ahead of any histograms. The title page contains a summary of the systems configuration, the time reduction effectively started and stopped, as well as identifying the system which was monitored and the tape numbers containing the data set.

Following this information is a list of all dispatcher options that were active. This information is obtained from word 0 of the dispatcher module (.MDISP). If Priority B processing was enabled, then the SNUMBs being granted Priority B processing will be listed.

The user can alter the dispatcher algorithm to satisfy particular installation requirements. The three most commonly used algorithms are Urgency Thruput, I/O Priority and Priority B processing. If none of these options are selected, job urgency and channel time will be ignored and the standard rules for job selection will be observed. In effect, the standard rules imply a First-In/First-Out mode of operation.

While the dispatching algorithm involves some rather complex decisions, the primary driving force behind the algorithm is the urgency of a job. When the I/O Priority option is set, a job's urgency code is computed at 192 millisecond intervals of processor time. If the job is an ordinary slave program and its urgency code is not 0, the code is reduced by 1. Each time an urgency code is reduced to 0, a processor-bound job is dispatched at the next dispatch. If a job is a priveleged slave program, or if its urgency code is 0, the urgency code is recomputed as follows:

Urgency Code = L+T

where T = 1 if total channel time is at least equal to total processor time; otherwise T = 0.

L = 0,1,2,3,4 depending on the ratio of local channel time to local processor time.
set to 0 if ratio $<$ 1:1
set to 1 if ratio $\geq$ 1:1 $<$ 4:1
set to 2 if ratio $\geq$ 4:1 $<$ 16:1
set to 3 if ratio $\geq$ 16:1 $<$ 64:1
set to 4 if ratio $\geq$ 64:1

When the Urgency Thruput option is set, a job's dispatcher urgency is computed whenever the job's dispatcher urgency reaches 0. The dispatcher urgency is reduced by one each time the job is placed in the queue. A job's dispatcher urgency is computed as (U+2)/8 where U is a job's processing urgency from .CRSN1. This formula explains why it is very important to ensure that system programs and priority jobs are given processing urgency levels (U value) significantly higher than the ordinary slave job. Assume that $PALC has a processing urgency level of 54 and a

user job has a processing urgency level of 40. In this example, $PALC has a dispatcher urgency of (54+2)/8 = 7 while the slave job has an urgency level of (40+2)/8 = 5. Therefore, $PALC is given priority for only two dispatches until it will be considered equal to the slave job. On the other hand, if the slave job has a processing urgency of 6, then it would have a dispatching urgency equal to (6+2)/8 = 1. In this case, $PALC would be given priority for seven out of eight dispatches.

It is possible for a site to have both I/O Thruput and Urgency Thruput set. In this case, a combination of the above algorithms is implemented.

Finally, the Class B level of priority is a dispatching option which can be used to ensure faster processing for a set of one to three jobs. The order in which the jobs are given affects priority; that is, each job is checked for dispatch in the order in which it occurs. The Frequency Count is the number of dispatches to other jobs after two consecutive dispatches to this priority job. Thus, a Frequency Count of one allows the priority job two out of three dispatches. The Time Quantum is the number of 32-millisecond intervals allowed per time quantum. If both the Frequency and Time Quantum values are 0, priority for this job is delayed until a change priority request is issued.

The next several lines of output describe the overhead of all GMF monitors that were active during data collection. The monitor name is given, its CPU time in seconds, and its overhead as a function of total processor power. The GMF executive overhead is separated from the actual monitors and is listed as "EXEC". The monitor "NAME" is an area of code within the Mass Store Monitor and even though listed separately, it is also included under the monitor "MSM". The monitor "FMS" is also an area of code within the Mass Store Monitor, but in this case it has not been included under the monitor "MSM". Monitor "CM" describes the processor overhead of subroutine T4 (terminate processing) and subroutine T22 (start I/O processing). Monitor "MSM" describes the processor overhead of subroutine T7 (connect processing). Therefore, if the Channel Monitor was active, but the Mass Store Monitor was not, this report will list both "CM", "MSM" and "FMS" monitors. If both the Channel and Mass Store Monitors were active, then the combined overhead of both monitors can be found by the same sum above. For purposes of this report, percent overhead is computed as (CPU time used by monitor)/(total CPU time) x (number of processors).

If a termination record is not reduced for any reason, the lines describing monitor overhead will not be printed. Figure 11-2 provides a sample report.

11.5.2.1 CPU Queue Length Distribution (File 14). Figure 11-3 shows a sample histogram of how busy the CPU queues were. It indicates how many activities were waiting for a processor at the time the samples were taken. This histogram will be produced only if the dispatcher queuing option of the CPU Monitor has not been disabled. It should be noted that

the dispatcher queue is not analyzed over a continuous time frame, but rather is sampled once every 1500 dispatches.

All histograms produced by the CPUTRP are interpreted in the following manner:

- In the INDIV NUMBER column, the histogram displays the number of occurrences of a particular event. The particular event being evaluated is represented by the figures in column 5. Therefore, in figure 11-3 we see that 231 times the CPU queue had a length of 0, while 44 times the queue length was 1.

- In the INDIV PROB. column, the histogram displays the probability that a given event will occur. Therefore, there is an 84% probability that the queue length was 0; i.e., 84% of the time there was no queue, while 16% of the time the queue length was 1.

- The CUMUL NUMBER and CUMUL PROB. columns are merely the cumulative probability distribution of the histogram.

CPU AND TAPE REDUCTION, VERSION 7.2 - 11.74, 24 SEP 1982

```
**********  THE CPU AND TAPE MONITOR  **********

        REDUCTION STARTED AT 1730:20.349 WED 82-12-27
AND COMPLETED AT 1909:10.372 WED 82-12-27 A TOTAL TIME OF 1.65 HOURS
        ON SYSTEM NMCC2  RUNNING 4J2 ON TAPE D0020

        THE SYSTEMS CONFIGURATION CONSISTED OF:

                2 - 6680  CENTRAL PROCESSORS

                    DISPATCHER OPTIONS
                IN CORE PUSH AREA
                DYNAMIC BUFFERING OF SSA MODS
                CLASS B PRIORITY ON

                    PRIORITY B JOBS
            SNUMB       FREQ       TIME QUANTUM
            TS1           2             7

MONITOR OVERHEAD -  MONITOR     TIME(SEC)     % OVERHEAD
                    -------     ---------     ----------
                    EXEC            47          3.26
                    CPU             47          3.26
                    -------     ---------     ----------
                    TOTAL           94          6.52
```

Figure 11-2. CPU Title Page

CPU AND TAPE REDUCTION, VERSION 7.2 - 11.74, 24 SEP 1982

DISTRIBUTION COLLECTED ON SYSTEM NMCC2 AT 1730:20.349 WED 82-12-27

CPU QUEUE LENGTH DISTRIBUTION

REPORT 1

INDIV. NUMBER	CUMUL. NUMBER	CUMUL. PRC	INDIV. PRC	QUEUE LENGTH	PERCENT OF PROBABILITY OF OCCURRENCE 00 10 20 30 40 50 60 70 80 90 100
					I....I....I....I....I....I....I....I....I....I....I
231	231	84.00	84.00	0- 0	IXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
44	275	100.00	16.00	1- 1	IXXXXXXXX

275 ENTRIES TOTAL AVERAGE = 0.16000 VARIANCE = 0.134 STANDARD DEVIATION = 0.367

Figure 11-3. CPU Queue Length Distribution

In figure 11-4, we find that 40% of all CPU bursts have a duration of 6ms or less while 96.5% of all CPU bursts have a duration of 12ms or less.

- o The last line of the histogram provides some standard statistical output such as average, variance and standard deviations.

11.5.2.2 CPU Burst Distribution for User Activities/All Activities (File 14). Figures 11-4 and 11-5 show sample histograms of the CPU burst length distribution for user activities/all activities. They provide a measure of how long an activity held a processor before relinquish or interrupt.

WARNING: Each report is not a true histogram of individual burst lengths but is one of burst averages over a sampling period. For example, figure 11-5 shows 37062 bursts, but only 275 samples were taken. Each sample provided a count of intersample bursts and the accumulated burst time. For each sample, the histogram was charted as though each burst had the same length, namely, the average intersample burst length equals the accumulated-burst-time divided by count-of-bursts.

11.5.2.3 Percent of Memory Time in Queue. Figure 11-6 shows a sample histogram of the percent of time spent by an activity waiting in the dispatcher's queue for an available processor. During this wait time, an activity is effectively residing in memory, but accomplishing no useful work. This figure provides a quick answer to the question, "Is my system CPU-bound?" Any system in which activities are spending more than 15% of their memory time, waiting for a given resource, is definitely constrained by that resource. This histogram is only produced if the dispatcher queuing option of the CPU Monitor has not been disabled.

11.5.2.4 CPU Time Reports (File 10). This tabular report (figure 11-7) is produced every ten minutes (default period) of elapsed time. The first line indicates how many seconds have elapsed in the analysis, the total CPU time that has elapsed (total elapsed time x the average number of processors available through this period), the system ID, time of day and date. The amount of total CPU time is adjusted for any processors that have been released.

The next block of print gives the accumulated CPU time for each system program, for TSS, TRAX, WIN, and for all slave activities (presented as a single figure under the heading USER). The CPU time used by the slave portion of the GMC collector is shown separately.

CPU AND TAPE REDUCTION, VERSION 7.2 - 11.74, 24 SEP 1982

DISTRIBUTION COLLECTED ON SYSTEM NMCC2 AT 1730:20.349 WED 82-12-27

CPU BURST DISTRIBUTION FOR USER ACTIVITIES ONLY

INDIV. NUMBER	CUMUL. NUMBER	CUMUL. PRC	INDIV. PRC	MILLI-SECOND	PERCENT OF PROBABILITY OF OCCURRENCE 00 05 10 15 20 25 30 35 40 45 50
					I....I....I....I....I....I....I....I....I....I....I....I
85	85	2.00	2.00	3.000	IXX
144	229	5.30	3.30	4.500	IXXX
1511	1740	40.40	35.10	6.000	IXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
703	2443	56.70	16.30	7.500	IXXXXXXXXXXXXXXXX
1593	4036	93.70	37.00	9.000	IXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX
75	4111	95.40	1.70	10.500	IXX
47	4158	96.50	1.10	12.000	IX
150	4308	100.00	3.50	13.500	IXXX

4308 ENTRIES TOTAL AVERAGE = 6.91398 VARIANCE = 3.543 STANDARD DEVIATION = 1.882

REPORT 3

11-10

Figure 11-4. CPU Burst Length Distribution for User Activities

11-8

CPU AND TAPE REDUCTION, VERSION 7.2 - 11.74, 24 SEP 1982

DISTRIBUTION COLLECTED ON SYSTEM NMCC2 AT 1730:20.349 WED 82-12-27

CPU BURST DISTRIBUTION FOR ALL ACTIVITIES

INDIV. NUMBER	CUMUL. NUMBER	CUMUL. PRC	INDIV. PRC	MILLI-SECOND	PERCENT OF PROBABILITY OF OCCURRENCE 00 10 20 30 40 50 60 70 80 90 100	REPORT 4
					I....I....I....I....I....I....I....I....I....I....I	
707	707	1.90	1.90	1.500	IX	
574	1281	3.50	1.50	3.000	IX	
27974	29255	78.90	75.50	4.500	IXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX	
3979	33234	89.70	10.70	6.000	IXXXXX	
1698	34932	94.30	4.60	7.500	IXX	
1834	36766	99.20	4.90	9.000	IXX	
94	36860	99.50	0.30	10.500	I	
48	36908	99.60	0.10	12.000	I	
152	37060	100.00	0.40	13.500	I	
1	37061	100.00	0.00	15.000	I	
1	37062	100.00	0.00	16.500	I	

37062 ENTRIES TOTAL AVERAGE = 4.24950 VARIANCE = 2.104 STANDARD DEVIATION = 1.451

Figure 11-5. CPU Burst Length Distribution for All Activities

CPU AND TAPE REDUCTION, VERSION 7.2 - 11.74, 24 SEP 1982

DISTRIBUTION COLLECTED ON SYSTEM TS-7.2 AT 1730:20.349 WED 82-12-08. INITIAL CMC TAPE #00673

PERCENT OF MEMORY TIME IN QUEUE

REPORT 7

INDIV. NUMBER	CUMUL. NUMBER	CUMUL. PRC	INDIV. PRC	%	PERCENT OF OCCURRENCE 00 05 10 15 20 25 30 35 40 45 50
					I....I....I....I....I....I....I....I....I....I....I....I
11	11	4.911	4.911	0- 1	IXXXXX
4	15	6.696	1.786	2- 3	IXX
0	15	6.696	0.	4- 5	I
3	18	8.036	1.339	6- 7	IX
6	24	10.714	2.679	8- 9	IXXX
4	28	12.500	1.786	10- 11	IXX
1	29	12.946	0.446	12- 13	I
5	34	15.179	2.232	14- 15	IXX
2	36	16.071	0.893	16- 17	IX
5	41	18.304	2.232	18- 19	IXX
2	43	19.196	0.893	20- 21	IX
2	45	20.089	0.893	22- 23	IX
3	48	21.429	1.339	24- 25	IX
4	52	23.214	1.786	26- 27	IXX
.	.	.	.	. .	.
.	.	.	.	. .	.
.	.	.	.	. .	.

Figure 11-6. Percent of Memory Time in Queue

CPU AND TAPE REDUCTION, VERSION 7.2 - TEST, 28 SEPTEMBER 1982
ELAPSED (PROCESSOR) TIME IS 1806.06 (3612.12) SECONDS ON SYSTEM NMCC2 AT 1241:40.061 MON 82-09-13
CPU TIME (EXCEPTING OVERHEAD AND IDLE) IN HUNDREDTHS SECOND, # OF DISPATCHES, SERVICE TIME IN CLOCK PULSES
% CPU OF TOTAL CPU, % OF TOTAL DISPATCHES SINCE RUN START (DS1) AND % OF TOTAL DISPATCHES SINCE LAST PRINT (DS2)

	CALC	PALC	SYOT	RTIN	T/DS	LOGN	FSYS	HEALS	GMC	TSS	NCP	TLNT	FTS	TLCP	TRAX	USER
CPU	493	1398	2752	221	0	367	0	0	13	23780	1499	186	1682	1	0	120361
DSP	1765	5072	4937	708	0	2499	0	0	67	100832	2983	498	3970	6	0	55355
STM	178	176	356	200	0	94	0	0	133	150	321	239	271	134	0	1391
CP%	0.14%	0.39%	0.76%	0.06%	0. %	0.10%	0. %	0. %	0.00%	6.58%	0.42%	0.05%	0.47%	0.00%	0. %	33.32%
DS1	0.80%	3.41%	2.32%	0.33%	0. %	1.18%	0. %	0. %	0.05%	49.37%	1.60%	0.90%	1.22%	0.00%	0. %	38.25%
DS2	0.78%	2.99%	2.46%	0.32%	0. %	0.68%	0. %	0. %	0.05%	39.72%	2.40%	1.64%	1.39%	0.00%	0. %	47.44%

TSS:	CPU IN HUNDREDTHS SECOND				NUMBER OF DISPATCHES				SERVICE TIME IN CLOCK PULSES			
	TS1	TS2	TS3	TS4	TS1	TS2	TS3	TS4	TS1	TS2	TS3	TS4
EXEC	11730	0	0	0	82132	0	0	0	91	0	0	0
SDSP	12049	0	0	0	18700	0	0	0	412	0	0	0

TSS:	% OF TOTAL DISPATCHES SINCE RUN START				% OF TOTAL DISPATCHES SINCE LAST PRINT			
	TS1	TS2	TS3	TS4	TS1	TS2	TS3	TS4
EXEC	36.08%	0.%	0.%	0.%	29.75%	0.%	0.%	0.%
SDSP	13.28%	0.%	0.%	0.%	9.97%	0.%	0.%	0.%

OVERHEAD, IDLE, GATE LOOP TIME IN HUNDREDTHS SECOND (GATE LOOP ALSO INCLUDED IN OVERHEAD)

	CPU 1	CPU 2	CPU 3	CPU 4	CPU 5	CPU 6	TOTAL
OHEAD	32796	10580	0	0	0	0	43377
IDLE	109875	56832	0	0	0	0	166708
GLOOP	557	657	0	0	0	0	1215

# CPUS,	% SYSTEM CPU,	% TSS CPU,	% WIN CPU,	% TRAX CPU,	% USER CPU,	% IDLE TIME &	% GATE LOOP OF BUSY...SINCE
2.00	13.46	6.58	0.93	0.	33.32	46.15	0.62 RUN START
2.00	16.19	9.69	0.05	0.	46.70	27.94	0.58 LAST PRINT

Figure 11-7. CPU Time Report (Part 1 of 2)

QUEUE TIME IN TENTHS SECOND, AVERAGE QUEUE TIME PER DISPATCH IN CLOCK PULSES AND % QUEUE OF ELAPSED-TIME SINCE RUN START

	CALC	PALC	SYOT	RTIN	T/DS	LOGN	FSYS	HEALS	GNC	TSS	NCP	TLNT	FTS	TLCF	TRAX	USER
QUE	124	188	75	16	0	8177	0	0	2	472	4227	25	2484	3746	0	16477
AVG	450	238	97	145	0	20941	0	0	205	29	9069	327	4005	3996174	0	1905
QU%	0.69%	1.05%	0.42%	0.09%	0. %	45.28%	0. %	0. %	01.01%	2.62%	23.41%	0.14%	13.76%	20.74%	0. %	

---AVERAGE QUEUE POSITION WHEN IN QUEUE (OR PROCESSOR)

AQP	2.00	1.33	1.29	2.00	0.	1.17	0.			0.99	1.07	2.00	1.36	0.		

QUEUE TIME IN TENTHS SECOND, AVERAGE QUEUE TIME PER DISPATCH IN CLOCK PULSES AND % QUEUE OF ELAPSED-TIME SINCE LAST PRINT

	CALC	PALC	SYOT	RTIN	T/DS	LOGN	FSYS	HEALS	GNC	TSS	NCP	TLNT	FTS	TLCF	TRAX	USER
QUE	58	103	12	6	0	768	0	0	1	235	5	2	3	0	0	7541
AVG	634	298	372	123	0	11525	0	0	528	31	409	770	598	409	0	1610
QU%	0.96%	1.72%	0.21%	0.11%	0. %	12.72%	0. %	0. %	0.02%	3.89%	0.10%	0.04%	0.05%	0.00%	0. %	

---AVERAGE QUEUE POSITION WHEN IN QUEUE (OR PROCESSOR)

AQP	2.33	1.44	1.00	2.00	0.	0.	0.			0.96	0.	0.	0.	0.		

TSS:

	QUEUE IN TENTHS SECOND				AVG QUEUE TIME IN CLOCK PULSES			
	TS1	TS2	TS3	TS4	TS1	TS2	TS3	TS4
EXEC	438	0	0	0	34	0	0	0
SDSP	0	0	0	0	0	0	0	0

AVERAGE QUEUE LENGTH WAS 1.39 SINCE RUN START, 1.50 SINCE LAST PRINT.

Figure 11-7. (Part 2 of 2)

The second line of print in this block gives the accumulated number of dispatches for each of the associated programs. The service time for each program is presented on the third line. This figure is calculated by dividing the accumulated CPU time by the accumulated number of dispatches. This figure is then converted to clock pulses (1/64 ms). The fourth line in this block indicates what percentage of the total processor power is being used by each of the associated programs. This figure is based on 100% power availability. Therefore, if the operating system is using 30% processor power on a 3-processor system, it would be using (30%)(3) = 90% of a processor. The final two lines of print in this block provide an indication of what percentage of total dispatches is being used by this program. A program using a large percentage of dispatches for a small percentage of processor time is an indication of a program which is getting poor utilization of the processor and possibly asking for the processor an excessive number of times. When TSS is in Priority B mode, it will display this tendency.

The third set of print provides, for each copy of Time Sharing (TSS), the CPU time attributable to the executive phase and to the subdispatch phase. In addition, for each copy of TSS, the average service time is also presented. This figure is represented in clock pulses (1/64ms). This figure represents the average amount of CPU time used by TSS whenever it received control of the processor. By tracking this figure, it is possible to see if bad periods of TSS response coincide with periods of time when TSS was receiving inadequate time slices of CPU service. The service time is calculated by taking the total CPU time accumulated by TSS and dividing it by the total number of CPU bursts accumulated by TSS.

The next block of print shows the percentage of dispatches being given to both the executive and subdispatching.

If the user has selected the option to monitor special SNUMBs, the next block of print will list each of the special SNUMBs under which will appear the total CPU time accumulated by each of the SNUMBs. This block of print is not shown on the sample figure.

The next block of print shows the amount of overhead, idle and gate-loop time accumulated by each processor and by total. Gate-loop data is not applicable in a single processor environment. Gate loop time is that amount of time a processor is locked from executing because another processor has locked a required table or blocked a given area of code. In a multiprocessor environment, there are many instances where one processor is required to alter the values in a given table. While these values are being changed, the system wants to ensure that another processor does not reference the table, while it is in the midst of being changed. To prevent this from happening, the system will lock the table while it is in the midst of being altered. If a second processor desires to reference a locked table, it is required to execute a CPU "dead" loop while it waits for the table to be opened. The GLOOP statistics display the amount of

such loop time executed by each processor. This value will not be reported for a single processor environment but should appear for any multiprocessor environment where the software release was WW7.2/4JS. If this data is not reported, it indicates a problem with the GMF collector and CCTC should be contacted.

The final block of print is a percentage breakdown of CPU usage into the categories SYSTEM, TSS, TRAX, WIN, USER, and IDLE; also shown are the percentage of gate-loop time relative to processor busy time and the time corrected number of processors in use. These figures are printed both for the current 10-minute interval and for the total current reduction interval. This allows the user to track when peak usage of CPU power is occurring and what portion of the system is using this power.

Notes: (1) SYSTEM time includes CPU time accumulated in the "functions" OVERHEAD, CALC, PALC, SYOT, RTIN, TDS, LOGN, FSYS, DMTEX and MONITR. The time attributable to WIN, TSS and TRAX is neither SYSTEM nor USER. (2) The definition of overhead time is the time spent in the interrupt handler, the dispatcher and the SWAP processor, plus all gate-loop time. (3) It should be noted that the percentage figures are based on total CPU power and therefore add up to 100% (excluding the gate-loop percentage). In order to determine the "amount of a processor" required or used by a given function, it would be necessary to multiply the percentage figure for the function by the time corrected number of processors in use.

If the user has disabled the dispatcher queuing option of the CPU Monitor, the above blocks of output will occur two per page. If the dispatcher queuing option is enabled, then several blocks of output will be presented which depict the level of queuing occurring on the system.

The first block presents three lines of output for the same set of jobs as listed in the processor usage portion of the report. The total queue time accumulated by the job is given in line 1, the average queue time per dispatch is given in line 2 (presented in clock pulses) and the percent of elapsed time for which this job was in queue is presented in the third line. This final figure is not provided for the user entry. These three lines of output are accumulated data, since the start of the run. The next line of output indicates the average position this job held in the dispatcher's queue. If the job was being serviced by the processor, when the queue was examined, its relative queue position is considered to be zero.

These same four lines of output are repeated, but the data represents the various jobs' queuing statistics only since the last print out.

The final two blocks of output provide queuing statistics for the TSS and all special jobs selected for analysis. It should be noted that all TSS queuing statistics are for the executive only. Subdispatch queuing is not currently monitored.

11.5.2.5 CPU Plot of Percent Idle (File 31). This plot report shows the percentage of CPU power that is idle during each 10-minute (default period) interval. The data is taken directly from the CPU Time Report described in subsection 11.5.2.4. One horizontal line is output for every CPU Time Report table. The horizontal line represents one increment on the X-axis and it "paints" one datum of percent idle and the change of that datum since it was last plotted. Every 10th line also displays the current time of day. By the nature of the printing mechanism, an ordinate position is a cell, a range of values. The cell width is specified by a DELTA value given in the heading. The plot contains a heading line giving the system id, starting time of reduction and the date. It also contains the report title, and identifying mark, "A" here, of the curve, and the DELTA value. A plot summary, if not user deactivated, follows each plot. Following the summary is the overall system idle percentage in three intervals; zero to 25% idle, 25 to 50% idle and 50 to 100?% idle. A final line shows the plot interval in seconds (default is 600 seconds). Figure 11-8 is a sample of this plot.

11.5.2.6 WIN Report. This report will be generated under the same time interval control as is used for the CPU Time Report (subsection 11.5.2.4) (default value every 10 minutes). Figure 11-9 is a sample of this report (no WIN software was active at the time this sample report was generated). For each WIN program, a single line of information is presented. This line indicates the total amount of CPU time accumulated by the program during the specified time interval, the number of CPU bursts accumulated during the time interval and the rate (bursts/sec) at which the bursts were generated. This report can be used to track those periods of time when WIN programs were using excessive CPU time or, on the other hand, were being denied CPU service.

11.5.2.7 CPU Access by SNUMB Report. This report (figure 11-10) summarizes CPU usage and queuing statistics for every activity processed during the monitor session. The various columns of the report are self-explanatory and require no special explanation, except for the CPU-SYST and CPU-TRACE columns. The former column indicates the amount of CPU time charged to a job from the accounting system. The latter column indicates the amount of time the processor was actually allocated to a job according to the execution of the various special traces generated by the CPU Monitor. This number will tend to be somewhat larger than the value determined by the accounting system. If a set of asterisks appear, for a given job, under the Average Queue Position column, this is an indication that this job was never caught in the processor queue. It should be stressed that the processor queue is not analyzed continually. Rather, the queue is examined under a sampling scheme and, therefore, it is possible that a given job will never be caught in the processor's queue.

11.5.3 Tape Monitor Reduction Reports. The tape reduction title page is identical to the CPU reduction title page (see subsection 11.5.2), except that the configuration described is pertinent to the tape subsystem.

Shown are the channel number, the IOM number, the number of drives configured to the channel, the type of drive and whether the drives are cross-barred. The data shown is the configuration as presented in the boot deck. If any drives have been taken off line for maintenance or repair, it will not be reflected. The title page precedes any histograms; an example is presented as figure 11-11. It is written to file 14.

11.5.3.1 Number of Tape Drives in Use (Time Corrected) (File 14). A histogram report, seen as figure 11-12, shows the number of tape drives in use at the sampling epoch. The data are corrected for the inter-sample period, so that the figures listed under the heading "INDIV. PROB." correctly represent the fraction of reduction time the corresponding "NUMBER (of) DRIVES" were in use.

11.5.3.2 Tape Activity Report (File 13). This tabular report is made for each activity of a job that used tapes. For each activity of a job, the tapes used by that activity are described by type, unit number, and channel number. The report also presents how long the activity was

THIS PAGE LEFT INTENTIONALLY BLANK

CPU AND TAPE REDUCTION, VERSION 7.2 - 11.74, 24 SEP 1982

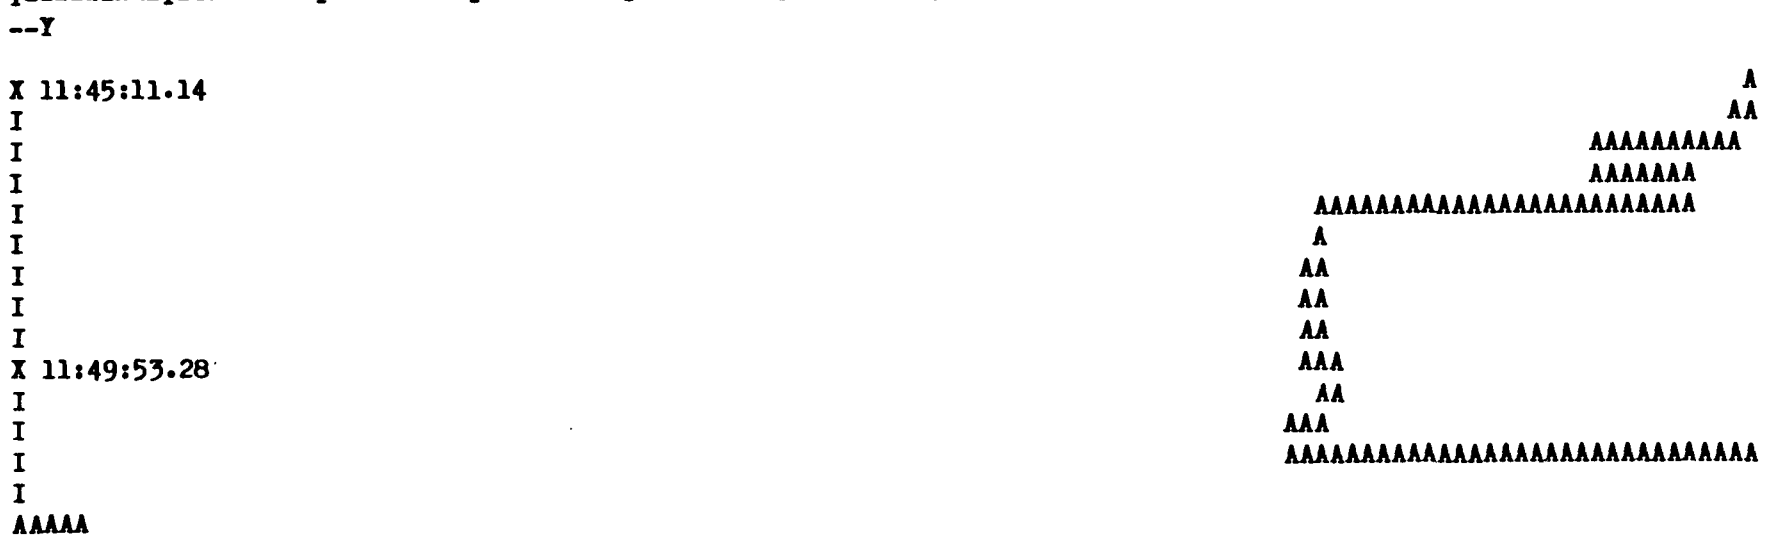
DISTRIBUTION COLLECTION ON SYSTEM NMCC2 AT 11:44:41 ON 82-12-15

CPU UTILIZATION OF SYSTEM/USER IDLE

PLOT-1

DELTA= 7.9365E-01 A=IDLE

TIME OF DAY % IDLE



CURVE	MINIMUM	MAXIMUM
IDLE	*0.6600E 02*	*0.9300E 02*

0. % 25% IDLE 0. % = 25% AND = 50% IDLE 100.00% 50% IDLE

THE PRINT PERIOD FOR THIS PLOT IS NOMINALLY 30 SECONDS

Figure 11-8. CPU Plot of Percent Idle

DISTRIBUTION COLLECTED ON SYSTEM DSCC1 AT 1005:54.955 TUE 82-01-19. INITIAL GMF TAPE # 1182

WIN REPORT -- TIME IN HUNDREDTHS SECONDS; RATE, 1/SEC

	NCP			TELNET			FTS			TLCF		
TIME	CPUTIME	BURSTS	BURST-RATE	CPUTIME	BURSTS	BURST-RATE	CPUTIME	BURSTS	BURST-RATE	CPUTIME	BURSTS	BURST-RATE
1015:54.385	0	0	0.	0	0	0.	0	0	0.	0	0	0.
1025:56.385	0	0	0.	0	0	0.	0	0	0.	0	0	0.
1035:58.074	0	0	0.	0	0	0.	0	0	0.	0	0	0.
1045:54.581	0	0	0.	0	0	0.	0	0	0.	0	0	0.
1055:59.947	0	0	0.	0	0	0.	0	0	0.	0	0	0.
1106:01.044	0	0	0.	0	0	0.	0	0	0.	0	0	0.
1116:03.416	0	0	0.	0	0	0.	0	0	0.	0	0	0.
1121:59.169	0	0	0.	0	0	0.	0	0	0.	0	0	0.

Figure 11-9. WIN Report

CPU AND TAPE REDUCTION. VERSION 7.2 - 11.74. 24 SEP 1982

DISTRIBUTION COLLECTED ON SYSTEM TS-7.2 AT 1730:20.349 WED 82-12-08. INITIAL GMC TAPE #00673

CPU ACCESS BY SNUMB

TIME OF TERMINATION	JOB SNUMB	CPU-SYST IN SECONDS	QUEUE TIME IN SECONDS	AVERAGE CPU QUEUE POSITION	MEMORY TIME (SEC)	PERCENT TIME IN QUEUE	ACTIVITY NUMBER	SWAP TIME (SEC)	CPU-TRACE IN SECONDS
1731:34	LCDTS	3.104E-01	2.704E 00	4.3	4.000E 00	67	1	0.	3.304E-01
1732:04	8247T	1.774E 01	5.883E 01	3.2	1.030E 02	57	1	0.	1.974E 01
1732:16	8278Z	1.686E 01	3.630E 01	3.8	7.700E 01	47	1	0.	1.886E 01
1733:31	8287T	3.230E 01	1.013E 02	3.5	1.890E 02	53	1	0.	3.430E 01
1734:09	8311T	6.848E-01	6.763E 00	4.6	1.000E 01	67	1	0.	7.048E-01
1734:09	8310T	8.323E 00	1.675E 01	2.8	4.100E 01	40	1	0.	8.523E 00
1735:22	8210T	3.156E 00	4.325E 01	2.1	5.700E 01	75	1	0.	3.356E 00
1736:34	8318T	1.585E 00	9.959E 00	4.1	2.100E 01	47	1	0.	1.785E 00

Figure 11-10. CPU Access by SNUMB

CPU AND TAPE REDUCTION, VERSION 7.2 - 11.74, 24 SEP 1982

***** THE CPU AND TAPE MONITOR *****

MONITORING ON 82-12-15 STARTED AT 11:44:41 AND COMPLETED AT 11:56:47 FOR A TOTAL TIME OF 0.20 HOURS
ON SYSTEM NMCC2 RUNNING W64000 ON TAPE D0020

THE SYSTEMS CONFIGURATION CONSISTED OF:

CHANNEL	IOM	# DEVICES	TYPE	SCT ADDR(SAME=XBAR)
18	0	6	9 TRACK TYPE	1936
19	0	6	9 TRACK TYPE	1960
24	1	6	9 TRACK TYPE	2032

MONITOR OVERHEAD -

MONITOR	TIME(SEC)	% OVERHEAD
EXEC	47	3.26
TAPE	1	0.01
TOTAL	48	3.27

Figure 11-11. Tape Reduction Title Page

CPU AND TAPE REDUCTION, VERSION 7.2 - 11.74, 24 SEP 1982

DISTRIBUTION COLLECTED ON SYSTEM NMCC2 AT 11:44:41 ON 82-12-15

THE TIME CORRECTED NUMBER OF TAPE DRIVES IN USE

INDIV. NUMBER	CUMUL. NUMBER	CUMUL. PROB.	INDIV. PROB.	NUMBER DRIVES	PERCENT OF PROBABILITY OF OCCURRENCE 00 10 20 30 40 50 60 70 80 90 100	REPORT 10
					I....I....I....I....I....I....I....I....I....I....I	
4769030	4769030	1.000	1.000	1- 1	IXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX	

4769030 ENTRIES TOTAL AVERAGE = 1.00000 VARIANCE = 0. STANDARD DEVIATION = 0.

Figure 11-12. Number of Tape Drives in Use (Time Corrected)

delayed in peripheral allocation and how many drives it was waiting for. The final print tells how many drives were in use when the activity started and how many were in use when it ended. At the right margin of the report are two columns of numbers. The first column is the numerical sequence in which the activities started, and the second column is the numerical sequence in which they terminated. Refer to figure 11-13 for a sample report.

11.5.3.3 Tape Status Report (File 12). This tabular report may be seen as figure 11-14. It shows the status of all configured tape drives, and is repeated each time an activity was delayed due to unavailability of tape drives. (The first report entry is listed with the initial trace captured by the monitor collector). The reason for the delay, the epoch of the delay and the program number of the activity delayed are printed. As part of the status information, the following is presented: IOM-channel number, unit number, released or not, assigned or not, dedicated or not, used for T&D's or not, number of errors on the device, and the SNUMB of the job using that drive.

11.6 Default Option Alteration

The CPU and Tape Reduction Program uses two inputs. The first is the data tape(s) produced by the GMC. The second input is a set of report option control cards. The various options and their formats are described in the following subsections. Report option control cards are not required if the default options are acceptable to the user.

11.6.1 General Format. The general format for an option request is as follows:

```
COMMAND (PARAMETER-LIST)
```

where COMMAND is a keyword specifying what action is to be taken and (PARAMETER-LIST), if required, provides data to accomplish the action. The COMMAND is interpreted through its sixth character, so the input must match the first six characters of the valid commands described below. Shorter commands must match exactly.

11.6.1.1 Valid Commands. Eleven commands are recognized by the card input interpreter. The chart below shows each command, the action to be taken, and in parentheses, the default inplications.

- o CPRINT — Print the CPU Time Report and one datum of the CPU Idle Time Plot (items A and f and g of table 11-1) at a specified repetition period. (Print every 600 seconds).
- o DEBUG — Activate the debug statements in one or more subroutines. (All debug statements are inactive).

```
CPU AND TAPE REDUCTION, VERSION 7.2 - 11.74, 24 SEP 1982
   TAPE ACTIVITY REPORT FOR SYSTEM NMCC2 ON 82-12-15 AT 11:44

 TAPES FOR PROG (22)  MUM   ACTIVITY   1          TYPE              UNIT              CHANNEL NUMBER
                                                 9 TRACK TAPE      4                 18( 0)

       THIS ACTIVITY HAD NO WAIT TIME


   WHEN ACTIVITY STARTED    1  TAPES WERE IN USE. WHEN IT ENDED     0  TAPES WERE IN USE          1      3
                  ***THIS PROGRAM WAS ACTIVE WHEN MONITOR STARTED
**********************************************************************************************************
```

Figure 11-13. Tape Activity Report

CPU AND TAPE REDUCTION, VERSION 7.2 - 11.74, 12 MAR 1982

TAPE STATUS REPORT FOR SYSTEM NMCC2 ON 82-12-15 AT 11:44

CHANNEL	UNIT	RLSED	ASGND	DED	T&D	ERRORS	SNUMB	CHANNEL	UNIT	RLSED	ASGND	DED	T&D	ERRORS	SNUMB
0-18	0	NO	NO	NO	NO	0		0-18	1	YES	NO	NO	NO	0	
0-18	2	NO	NO	NO	NO	0		0-18	3	NO	NO	NO	NO	0	
0-18	4	NO	YES	NO	NO	27	MUM	0-18	5	NO	NO	NO	NO	0	
0-19	0	YES	NO	NO	NO	0		0-19	1	YES	NO	NO	NO	0	
0-19	2	YES	NO	NO	NO	0		0-19	3	YES	NO	NO	NO	0	
0-19	4	YES	NO	NO	NO	0		0-19	5	YES	NO	NO	NO	0	
1-24	0	YES	NO	NO	NO	0		1-24	1	YES	NO	NO	NO	0	
1-24	2	YES	NO	NO	NO	0		1-24	3	YES	NO	NO	NO	0	
1-24	4	YES	NO	NO	NO	0		1-24	5	YES	NO	NO	NO	0	

****

****PROGRAM NUMBER 18 WAS DELAYED AT 11:53 FOR 2 7 TRACK TAPE

CHANNEL	UNIT	RLSED	ASGND	DED	T&D	ERRORS	SNUMB	CHANNEL	UNIT	RLSED	ASGND	DED	T&D	ERRORS	SNUMB
0-18	0	NO	NO	NO	NO	0		0-18	1	YES	NO	NO	NO	0	
0-18	2	NO	NO	NO	NO	0		0-18	3	NO	NO	NO	NO	0	
0-18	4	NO	YES	NO	NO	34	MUM	0-18	5	NO	NO	NO	NO	0	
0-19	0	YES	NO	NO	NO	0		0-19	1	YES	NO	NO	NO	0	
0-19	2	YES	NO	NO	NO	0		0-19	3	YES	NO	NO	NO	0	
0-19	4	YES	NO	NO	NO	0		0-19	5	YES	NO	NO	NO	0	
1-24	0	YES	NO	NO	NO	0		1-24	1	YES	NO	NO	NO	0	
1-24	2	YES	NO	NO	NO	0		1-24	3	YES	NO	NO	NO	0	
1-24	4	YES	NO	NO	NO	0		1-24	5	YES	NO	NO	NO	0	

****

Figure 11-14. Tape Status Report

- o HISTOGRAM — Modify a histogram report. (Histogram report parameters are as shown in table 11-2).

- o NEW — Prepare to accept another, or repeat, set of GMC data tapes and another set of option alteration commands for a following reduction run. (Reduction is complete when the current tape(s) are reduced).

- o HPRINT — Print the histograms only if the reduction interval is, at least, a specified minimum. (Print the histograms only if the reduction interval is at least 900 seconds).

- o ON / OFF
  (1) Turn on/off one or more reports. (All reports are on).
  (2) Turn on/off the CPU reduction module or the TAPE reduction module but not both. (CPU reduction is on; TAPE reduction, off).

- o PLOT — Modify a plot. (Plot report parameters are as shown in table 11-3).

- o STOP — Stop reduction after processing a specified number of physical tape records. (No limit).

- o TIMEFRAME
  (1) Accept one to five time windows for overall data reduction or for plot output. (Data reduction and report output are not time limited).
  (2) Accept one to five time windows for debug output. (Debug output is not time limited).
  (3) Treat the receipt of a specified or implied special record as the end of one timeframe and start of another - that is, print the reports and start a new reduction period immediately. (No timeframe action is taken).

- o TITLE — Accept a new title to be printed on all reports as the new banner. ("CPU AND TAPE REDUCTION, VERSION 7.2-11.74, 24 SEP 82" is printed as banner).

11.6.1.2 Parameter List. The parameter list is a sequence of data items separated one from the next by a field of spaces. Any one space in the field may optionally be replaced by a comma or colon. A data item may be an alphanumeric, a quoted string, an integer, a real number or a null. Strings are enclosed in single or double quotes. Real numbers include the standard FORTRAN forms: for example, 12.34, 1234E-2, 1.234+1 are all

Table 11-2. Histogram Default Parameters

REPORT ID	TRAILER FLAG	# INTERVALS	LOW VALUE	INTERVAL SIZE
1	ON	40	0	1
2	ON	40	0	1
3	ON	45	0.0	1.5
4	ON	45	0.0	1.5
5	ON	40	0	1
6	ON	40	0	1
7	ON	40	0	1
8	ON	40	0	1
9	ON	40	0	1
11	ON	40	0	1

Table 11-3. Plot Default Parameters

REPORT ID	# POINTS to PLOT	Y-MINIMUM	Y-MAXIMUM
10	UNLIMITED	0.0	100.0

accepted as the same number. A null is two consecutive commas or colons; it stands for a missing data item and for the separator fields on either side of the missing item.

11.6.1.3 Command Syntax. Each command with its parameter list must be complete on a single card, but more than one command may be placed on each card if spaces separate them. The parameter list may be contiguous with the command keyword or spaces may intervene. The syntax for each command is given in following subsections. The data acronym RID stands for a one or two digit integer specifying the report identification number of the report to be affected. Report identification numbers are shown in table 11-1. The specifications "integer-W", "alphanumeric-X", "real-Y", and "string-Z" require user input of an integer number in a field at most "W" wide, an alphanumeric in a field at most "X" wide, a real number in a field at most "Y" wide and a quoted string in a field at most "Z" wide, respectively. Alphanumerics and strings exceeding the field width are accepted but are truncated to the right. Integers and reals exceeding the specified field width are rejected as illegal data. Integers may not contain a decimal point. Reals must contain a decimal point or an exponent. Reals are limited, for all data input, to a maximum of 14 characters including leading sign, decimal point, and exponent. Reals without exponent are further limited to a maximum of 11 characters including leading sign and decimal point. Data items shorter than the specified field width are entirely acceptable in all commands, but a null item is acceptable only where provided for.

11.6.1.4 General Control Command Syntax. The CPU and Tape Reduction program permits reduction of either the CPU associated traces or the tape subsystem associated traces during one pass over the data tapes. The CPU reduction phase is on, and the tape reduction phase is off, by default. The following commands may be used to exercise general control over the reduction:

- DEBUG — Activates all debug statements.

- DEBUG (MODULE-NAME-1,..., MODULE-NAME-N)
  Activates the debug statements in the explicitly named modules, where MODULE-NAME-i is the FORTRAN name given to the subroutine or function. Only the first six characters are interpreted. The routines currently debuggable are:

  CLOCK - keeps the master clock

  FIGCHG - examines a reconfiguration record for CPU related change

  GETOKE - lexical analysis of card input

LOGREC - gets a logical record meeting reduction requirements

NXTRECRD - tape input handler

RECNFG - gets the initial reconfiguration trace

TIMSET - sets up the array of start-stop times for user specified reduction windows

The one-shot program GMFEXC which reads the initial tape record is executed before card input. It may be debugged by setting program switch word, bit 11 (i.e., ON6).

- HISTOGRAM (RID, TRAILER-FLAG, NUMBER-OF-INTERVALS, LOW-VALUE, INTERVAL-SIZE)
  Turns on the histogram with report identification number RID, integer-2. Print/don't print the histogram summary if TRAILER-FLAG is ON/OFF, alphanumeric-3. The new number of histogram intervals is NUMBER-OF-INTERVALS, integer-3. The new low value and interval size are LOW-VALUE and INTERVAL-SIZE both integer-10 or real-14 according to the type of histogram. (In the current version, the number of intervals may not exceed 50).

The following two options are as above, but without change to the current specification of those parameters not appearing in the list.

- HISTOGRAM (RID,,NUMBER-OF-INTERVALS, LOW-VALUE, INTERVAL-SIZE)
- HISTOGRAM (RID, TRAILER-FLAG)

- NEW (TAPE-NUMBER)
  Flags the reduction program to stop reading card input options and to proceed with data reduction. It also informs the program that another data reduction is to follow, and that the following reduction is to lead off with the tape number specified by this option (any six characters). For that following reduction, the default options will prevail unless modified by option alteration commands following the NEW command. This command must be the last command on a card (following commands on the same card are lost).

- HPRINT (X) Prints histograms if reduction interval is at least X, integer-10, seconds. Otherwise, the histogram data are lost.

- OFF Turns off all histograms and plots (if any).

- OFF (RID-1,...,RID-n)
  Turns off the reports specified by report identification numbers RID-1,...,RID-n.

For the remaining OFF options, the first two characters of the parameter are necessary and sufficient.

- OFF (PLOTS) Turns off all plots (if any).

- OFF (HISTOGRAMS)
  Turns off all histograms.

- OFF (CPU) Turns off the CPU reduction phase of the program and turns on the TAPE reduction phase.

- OFF (TAPE) Turns off the TAPE reduction phase and turns on the CPU phase. This is the default specification and is not likely to be used, but is included here for completeness.

- ON ON options use the same parameters as those shown for the OFF command, with appropriate reversal of intent. All reports are on by default.

Note: CPU and tape reduction cannot be done simultaneously. To accomplish both reductions in a single job, use the NEW option and a form of the ON/OFF option: ON(TAPE), for example.

- STOP (X) Stop data reduction after reading X, integer-10, records maximum.

- TIMEFRAME (DEBUG, START TIME, STOP TIME, START TIME, STOP TIME, ...)
  Accepts up to five start/stop pairs, where each component of the pair is a pseudo decimal number (real 7) in the form HHMM.SS where HH is the hour, MM is the minute and SS is the second. The decimal point must appear on the card, to further delimit the action of debug statements. The limiting action becomes effective as soon as the program determines "what time it is." This command must be accompanied by some form of the DEBUG command or no debug output will be produced. The timeframes are independent of, though not unaffected by, any other user selected timeframes.

- TITLE (X) Accepts the string X, string-72, as the new banner for all reports. X must begin and end with a quotation mark, single or double, and may contain any FORTRAN acceptable character except the matching quote with which it is delimited.

11.6.1.5 CPU Reduction Command Syntax. In addition to those commands given in section 11.6.1.4, the following are pertinent to the CPU reduction phase.

- o CPRINT (X) Prints the CPU Time Report and one datum of the CPU percent-idle plot every X, integer-10, seconds. The specified period is nominal and the actual period will vary with trace times. This command will not affect the on/off status of these reports.

- o PLOT (RID, NUMBER-OF-POINTS, Y-MINIMUM, Y-MAXIMUM)
  Turns on plot with report identification number RID, integer-2. The new number of points to plot is NUMBER-OF-POINTS, integer-10, maximum. The y-axis of plot is bounded below by Y-MINIMUM, real-14, and above by Y-MAXIMUM, real-14.

The following two options are as above, but without change to the value of those parameters not specified.

- o PLOT (RID, NUMBER-OF-POINTS)
- o PLOT (RID,, Y-MINIMUM, Y-MAXIMUM)

- o TIMEFRAME (RID, START TIME, STOP TIME, START TIME, STOP TIME, ...)
  Turns on report with report identification number RID, integer-2. Sets start and stop times according to the pair
  START TIME, STOP TIME, START TIME, STOP TIME,
  where each component is a pseudo decimal number (real 7) in the form HHMM.SS where HH is the hour, MM is the minute and SS is the second. The decimal point must appear on the card. Up to five start/stop pairs may be specified for a given report. The count of five may be achieved in one or several commands; however, the times must be in increasing order in order to get the results wanted. The final quadruple in a command may be abbreviated - if only a start time is given, the stop time is permanently unbounded. If RID is zero, the timeframe(s) bound the entire reduction but the status of individual reports (on or off) is unaffected. Histograms may not be individually delimited with this command, thus the command affects the entire reduction (RID equals zero) or the CPU plot (RID equals 10).

- o TIMEFRAME (LOSTDATA)
  Treat the receipt of a GMC lost data trace as the end of one timeframe and the start of another.

- TIMEFRAME — Treat the receipt of a reconfiguration trace in which any CPU related data has changed (processor has been released or assigned) as the end of one timeframe and the start of another.

11.6.1.6 Tape Reduction Command Syntax. In addition to those commands given in subsection 11.6.1.3, the following is pertinent to the tape reduction phase.

- TIMEFRAME (0,0,0, STOP TIME)
  By the nature of its construction, the Tape Reduction Module can only produce a nonwindowed set of reports (items M and N of table 11-1). A timeframe set for overall reduction (RID equals zero) will affect only the final stop time for these reports. The start time will always coincide with the start of the GMC collector. For consistency, the histogram of tape drive use (item L of table 11-1) will cover the same period. The stop time must be in the same format as described in previous timeframe commands.

11.6.1.7 Card Input Errors. If the card interpreter cannot understand a command keyword or data item, it will echo the entire card and mark the last correctly interpreted character position with the message "ILLEGAL COMMAND (or DATA) FOLLOWING...." Depending on the state of the interpreter, it may attempt to find the end of the incorrect command in order to pass on to a following command on the same card, or it may reject the remainder of the card; it will report which of the two actions was taken, and will then continue with card interpretation. In no case will an action confirmed by the interpreter be reversed or "undone" by a following error. Due to the manner in which cards are interpreted, the JCL file used to execute the CPUTRP should have the line numbers removed prior to execution. If this is not done, the interpreter will try to process the line numbers on the data cards as valid input requests and generate error messages when unable to do so. Despite the error messages, the CPUTRP will process correctly.

11.6.1.8 Examples of Command Use. A few examples will show the simplicity of the scheme.

- ON (99 1 5 2, 3, 6)
  Action: Turns on reports 1, 2, 3, 5 and 6. Report 99 doesn't exist, so that parameter is ignored.

- ON ("1", 002, THREE, 4.0)
  Action: Card is ignored since report id's are not one or two digit integers.

o TIMEFR (0 2102.00, 2304.00, 0500.00)
Action: Sets overall data reduction window to start at 2102 hours, stop at 2304 hours, and start again at 0500 hours with a permanently unbounded stop time.

o HISTOG (12, OFF, 42, 1.2, 1-4)
Action: If report ID #12 is a histogram, turns the trailer print flag off. If histogram #12 is of the real continuous type, turns the report on, sets the number of intervals to 42, beginning at 1.2 (Units) with each interval 10,000.00 (Units) wide.

o NEW (*A2345) DEBUG OFF (PL)
Action: Tape number *A2345 is remembered to be the lead off tape for a following run. Both DEBUG and OFF(PL) commands are lost since they appear on the same card. If these commands are desired for the next processing, they should appear on a set of following cards.

o DEBUG DEBUG (TIMSET)
Action: Turns on all debug statements. DEBUG(TIMSET) is a valid command, is correctly interpreted, and is repetitious.

11.7 JCL

The CPU and Tape Reduction Program is controlled by the following JCL:

```
$    SNUMB        ...
$    IDENT        ...
$    USERID       ...$...
```

These are required by the installation.

```
$    SELECT       B29IDPX0/OBJECT/CPU-TAPE
$    LIMITS       99,30K,,20K
```

These two cards load the reduction program object code and begin execution.

The following card,

```
$    TAPEn        01,X1D,,data-tape-number
```

describes the tape (7- or 9- track) generated by the data capture procedure, and

```
$    DATA         05
```

is used to identify the data cards that follow and that specify the input options as described in subsection 11.6.1. When options are not required, no $ DATA card is required.

Figure 11-15 shows a sample deck setup.

11.8 Multireel Processing

If more than a single reel of data has been collected, a series of messages will be outputted to the console informing the operator that a new data reel is required. The following are the messages produced.

a. DISMOUNT REEL #XXXXX THEN MOUNT REEL NUMBER YYYYY ON ZZZZZ

In this case, XXXXX is the old reel number, YYYYY is the new reel number, and ZZZZZ is the tape drive ID.

If the operator fails to mount the new tape, the above message will be repeated three times, after which the program will terminate, and all reports produced.

b. IS TAPE XXXXX MOUNTED ON DRIVE ID YYYYY (Y/N)

In this case, XXXXX is the tape number being requested for mounting and YYYYY is the tape drive ID.

This message occurs when the data reduction program finds the wrong tape has been mounted (by comparing internally generated tape labels). If the operator answers N, the message in (c) below is produced. If the operator answers Y, the data reduction program will terminate and all reports will be produced. In this case, the data reduction program is unable to process the tape. Even though the operator is mounting the correct tape, the internal label on the new tape does not match that being requested by the old tape. The user should check the data collection session to insure that the operator did not respond with an incorrect tape number during multireel change.

After entering the Y or N, the operator will need to hit the EOM key twice in order for the response to be transmitted.

c. WRONG REEL JUST MOUNTED, DISMOUNT AND MOUNT REEL XXXXX ON ZZZZZ

In this case, XXXXX is the new reel number, and ZZZZZ is the tape drive ID.

d. CAN TAPE XXXXX BE MOUNTED ON DRIVE YYYYY (Y/N)

In this case, XXXXX is the new desired reel and YYYYY is the tape drive ID.

```
$       IDENT       18020251/30/3044
$       MSG2        THIS JOB WILL USE THE FOLLOWING TAPES
$       MSG2        IN THE GIVEN ORDER, ALL RING-OUT
$       MSG2        AJ283, AD281, DC020
$       SELECT      B29IDPKO/OBJECT/CPU-TAPE
$       LIMITS      99,30K,,2CK
$       TAPE        01,X1D,,AJ283,,DATA
$       DATA        05
        DATA cards
$       ENDJOB
```

Figure 11-15. Sample Deck Setup

If the operator fails to answer this message it will be repeated until he responds with a "Y" for YES or "N" for NO. If he types in "Y", then message (a) will be repeated. If he types in "N", then the program will be terminated and all reports will be produced.

11.9 Tape Error Aborts

During the course of processing it is possible that the operator will be required to abort the data reduction program due to an irrecoverable tape error. If such a condition occurs, the operator should abort the job with a "U" abort. This will allow the data reduction program to enter its wrap-up code processing and produce all reports generated prior to the tape error.

SECTION 12. TRANSACTION PROCESSING SYSTEM DATA REDUCTION PROGRAM (TPSDRP)

12.1 Introduction

The Transaction Processing System Data Reduction Program is a FORTRAN program that sequentially processes the data tape written by the General Monitor Collector (GMC). It produces several reports pertaining to response times for user/TPS interactions, in order to help the site analyst fine-tune the TPS parameters for maximum overall efficiency. A list of the information collected is found in table 12-1.

12.2 Transaction Processing System Trace Number

When work on the TPS Monitor began, all available system trace numbers (1 to 77 octal) were used either by the dispatcher itself of by other monitors of the GMF. It was decided to use trace 74 for the TPS Monitor, as this was least likely to cause a conflict.

12.3 Transaction Processing System Trace Collection

Operator messages to the TPS are routed through the Core Allocator (POPR). To avoid altering the system software outside of that being monitored, the TPS Monitor is turned on or off by the console operator via the TP MESS command; i.e., once the TPS is running, the operator requests "TP MESS." When the console responds "TP MESS?" the message "TRACE ON" is entered to start collecting traces, or "TRACE OFF" to discontinue trace collection.

12.3.1 Sample Decks. Figures 12-1 and 12-2 are sample decks for collecting TPS traces. Figure 12-1 is a compile of the GMF with only the transaction processor data collector. Figure 12-2 illustrates a run of that program. Figure 12-21, at the end of this section, lists the alters that must be incorporated into the WWMCCS TPS in order to collect the GMF traces.

12.4 TPE Tuning Guide Reports

This section and the following section will contain the running instructions for producing the TPE reports. There are two separate programs producing two separate sets of outputs. This section deals with the Tuning Guide formatting dump. Both sections will also contain the input and output formats.

12.4.1 Input. To produce the Transaction Processing Tuning Guide Reports from the TPE System tape, the program B29IDPXO/OBJECT/TPETG must be run. Input parameter cards may be used. They must be inserted immediately after the $DATA I* card in the JCL. If the input parameter cards are omitted, all the data will be reported. The input option consists of START/STOP times

Table 12-1. Transaction Processing System Traces (Part 1 of 2)

SUBTRACE NUMBER TRXD	INFORMATION
1	TPE scans terminals for a connect request
2	TPE connects a new terminal
3	TPE notices a break request from user
4	TPE disconnects a terminal
5	Number of UNDOT table entries increases or decreases by 1. UNDOT is an acronym for UNDelivered OuTput. If a user runs a TPAP that returns information to the terminal, but the terminal is busy or disconnected when the output is ready for transmission, the output is saved until that terminal becomes available (Transaction Processing System Site Manual [DD36B], p. 3-15). If memory is unavailable, the undelivered output is written to the journal file (if present)
6	Reason that TSP is unable to continue and is about to abort itself
7	Name of TPAP (Transaction Processing Applications Program) SPAWNed at user request
8	Name of TPAP removed from "quick-access" file in order to SPAWN another TPAP requested by user
9	Number of blocks of buffer space acquired by TPE, via a MME GEMORE, in order to run a TPAP requested by user
10	Number of blocks of buffer space refused to TPAP; i.e., that TPAP cannot be run at this time due to insufficient core
11	Invalid keyword (request for a specific TPAP) entered
13	Number of blocks of buffer space released after the specified TPAP has finished running
14	Output transmitted
16	One entry of undelivered output (see #5 above) has been retrieved from temporary storage on the journal file
17	INTERCOM message (message sent from one TPAP to another) is reissued, because confirmation that message was received did not arrive at sending TPAP. The need to resend a message could indicate a line problem or a hardware problem
18	TPAP receives INTERCOM line, enabling TPAP-to-TPAP communication to commence
19	Write instruction performed by a TPAP
20	Read instruction performed by a TPAP
21	SNUMB of TPAP which terminates abnormally due to operator intervention or because TPS terminates
22	SNUMB of TPAP which terminates abnormally due to some error within the TPAP itself

Table 12-1. (Part 2 of 2)

SUBTRACE NUMBER TRXD	INFORMATION
31	Request for TPAP status
32	Table reaching threshold size and the number of times this table has reached threshold size. The TPE threshold and impasse traces are taken when these messages are sent to the operator's console (DD36B, p. 6-19). Thresholds are error conditions correctable by TPE. For instance, if an output buffer is full, TPE will direct output to the online printer to prevent an overflow situation. Threshold conditions are counted. When the count reaches certain predetermined levels, the operator is informed and a trace taken. The number of threshold conditions can be reduced by adjusting the buffer sizes specified in the .TPC. macro (DD36B, p. 6-10). If TPE encounters an error condition from which it cannot recover, it notifies the operator that it has reached an impasse and terminates itself
33	Master terminal sign-on accepted
34	IDs of master terminal and second terminal attempting to sign on as a master terminal
35	Master terminal disconnect
40	Compression of TPAP information within TPS's core space attempted. This is done only when the master terminal operator is entering dynamic information for a new TPAP and there is insufficient contiguous storage available
41	Name of new TPAP being dynamically removed from TPS from master terminal
42	Name of TPAP being dynamically removed from TPS from master terminal
43	Name of TPAP being modified by master terminal user
44	Command issued at operator console
45	Command issued at master terminal
52	Initial buffer sizes of major TPS files

```
$           FILEDIT SOURCE,OBJECT,INITIALIZE,NONE
$           DATA    *C,COPY
$           LOWLOAD
$           OPTION  ERCNT/500/
$           GMAP    NDECK,NLSTOU
$           SELECTA B29IDPXO/GMFCOL/GMF/GMF.TOP
$           SELECTA B29IDPXO/GMFCOL/TPE/TP.INIT
$           SELECTA B29IDPXO/GMFCOL/GMF/GMF.MID
$           SELECTA B29IDPXO/GMFCOL/GMF/GMF.MON
$           SELECTA B29IDPXO/GMFCOL/GMF/GMF.BTM
$           GMAP    NDECK,NLSTOU
$           SELECTA B29IDPXO/GMFCOL/TPE/TPE200
$           EXECUTE
$           ENDEDIT
$           ENDCOPY
$           PRMFL   R*,R/W,S,B29IDPXO/GMFCOL/GMF.OBJ
```

Figure 12-1. TPS GMF Compile JCL

```
$           OPTION  ERCNT/500/
$           LOWLOAD
$           EXECUTE DUMP
$           PRMFL   R*,R,S,B29IDPXO/GMFCOL/GMF.OBJ
$           LIMITS  15,16K
$           PRIVITY
$           TAPE    OT,X2D,,,,RING-IN
$           FILE    DK,X1R,300R
$           DATA    I*
MO M1 M2 M3 M4 M5 M6 M8
```

Figure 12-2. Control Card for Running Only the TPS GMF

for the reports. The first card gives the number of times to follow on the second card. The second card contains a maximum of five pairs of START/STOP times. The input is in free format. A description of each follows.

- START - This is a time of day on a 2400-hour clock. The format is HH:MM. It can be used to skip the printing of traces created before a certain time. The data for the traces skipped will not be reported at all. To start at the beginning of the tape when using a STOP time, use 00:00 for the START time.

- STOP - This is also a time of day on a 2400-hour clock. The format is HH:MM. It can be used to skip the printing of traces created after a certain time. The data for the traces skipped also will be eliminated from all reports. At midnight, the clock will roll over to 00, so this time may be less than the START time.

  An example of the input cards are as follows:

  2 (Card 1)
  22:00,23:30 (Card 2)

After inserting the options desired, the $TAPE card must be altered to insert the TPE Data Collector tape number before executing the program. A list of the JCL is shown in figure 12-3.

12.4.2 Output. Up to 15 reports can be produced by this program. If no data pertaining to a particular report is found, the report is not produced. All 16 reports listed here are written to report code 08. Example reports are shown in figures 12-4 through 12-18. A list of each report with its description follows.

12.4.2.1 TPE Summary Report. This report includes the total number of reads, writes, invalid keywords, keyword table suppression attempts and Transaction Processing Applications Programs (TPAPs) termed. Also included are the maximum number of UNDOT (undelivered output) entries encountered and the overall time-weighted average number. The maximum number of the largest number of entries in the trace 5 (UNDOT Table Change). The trace 5 obtains the number of entries from the UNDOT table every time a change is made to the table. The average is the sum of the number of entries of each trace 5 times the corresponding time interval divided by the sum of the time intervals.

The minimum, maximum and time-weighted average number of teminals on TPE are reported, where the minimum number is the least number of terminals connected at any one time and the maximum number is the most terminals connected at any one time. These numbers are obtained by adding the terminal connect traces (2) and subtracting the terminal disconnect traces (4). The average is the sum of the number of terminals (trace 2 or 4) multiplied by the corresponding time interval divided by the sum of the time intervals.

```
$           IDENT
$           LOWLOAD
$           SELECT  B29IDPX0/OBJECT/TPETG
$           LOWLOAD
$           EXECUTE DUMP
$           MSG2    1,PLEASE ENTER "U" IF TAPE
$           MSG2    1,ERRORS OCCUR
$           LIMITS  10,45K,-2K,30K
$           FFILE   P*,LGU/(06,08,09,29,42)
$           TAPE    10,X1D,,,,NO-RING
$           DATA    I*
2
22:00,23:30
$           ENDJOB
```

Figure 12-3. JCL for TPE Data Reduction

The minimum, maximum and average interval in seconds between connect request scans are reported, where the minimum interval is the least amount of time between terminal connect request scans (trace 1) and the maximum interval is the largest amount of time between terminal connect request scans. The average is the sum of the time intervals divided by the total number of scans. Seven different buffer tables (see 12.4.2.5) and the initial buffer sizes in blocks, the total number of thresholds reached by each buffer are produced.

Lastly, the minimum, maximum and average interval in seconds between break requests encountered are also reported. The minimum interval is the least amount of time between noticed break requests (trace 3) and the maximum interval is the largest amount of time between noticed break requests. The average is the sum of the time intervals divided by the total number of noticed break requests. See figure 12-4.

12.4.2.2 NONEW Option Report. This report lists the times in hours, minutes and seconds that an operator console request or master terminal request of NONEW (do not accept new terminals/transactions) or NEW (equivalent of NONEW off) was encountered. See figure 12-5.

12.4.2.3 Profile ID Report. This report lists profile IDs and the time in hours, minutes and seconds that each was either stored or erased. A TPAP profile is information assembled into the TPS, consisting of the TPAP name, keywords (entry points), I/O buffer sizes, and option flags (TPS Users Manual [DD41], p. 4-8). See figure 12-6.

12.4.2.4 Buffer Space Report. This report lists the times in hours, minutes and seconds that blocks were either acquired or released, and the number of blocks acquired or released each time. Each block is 64 words. See figure 12-7.

12.4.2.5 Table Reaching Threshold Size. This report lists the buffer table names and the time in hours, minutes and seconds that table threshold size was reported to the operator ("TPE THRSHLD ON ID xx LK xx RY xx OT xx TR xx UD xx DB xx"). This message is issued when requested by the operator or the master terminal user, when TPE terminates (unless all counts (xx) are zero), and when any count reaches a predetermined value. The standard values are 8, 64, 512, 4096, 32768, 262144. The message refers to the following buffers: INPUT (buffer for input buffer assignment, DD36B, p. A-3), LINK (message queue entries, DD36B, p. A-3), RECOVY (recovery table, DD36B, p. A-5), OUTBUF (output buffer queue area, DD36B, pp. A-6, 7), TRABUF (transmission buffer queue area, DD36B, p. A-8), UNDOT (undelivered output catalog, DD36B, p. A-9), and dynamic INTERCOM (interslave communications facility, DD41, chapter 5). See figure 12-8.

12.4.2.6 Master Terminal Report. This report lists the Master Terminal IDs, the sign-on times, the disconnect times, other terminal IDs attempting sign-on and those times. All times are in hours, minutes and seconds. See figure 12-9.

TPE SUMMARY REPORT

NO. OF READS = 52	# KYWD TABLE	# TPAPS	UNDOT ENTRIES		NO. OF TERMINALS		
NO. OF WRITES = 42	SUPP. ATTEMPTS	TERMED	MAXIMUM	AVERAGE	MINIMUM	MAXIMUM	AVERAGE
NO. INVALID KYWDS = 1							
	0	0	40	33.7	2	5	3.0

CONNECT REQUESTS SCANS (SECS)						BREAK REQ. ENCOUNTERED (SECS)		
MINIMUM INTERVAL	MAXIMUM INTERVAL	AVERAGE INTERVAL				MINIMUM INTERVAL	MAXIMUM INTERVAL	AVERAGE INTERVAL
			INBUF = 83 BLKS	INBUF THRES. TOTAL = 000				
			RECOVY = 83 BLKS	RECOVY THRES. TOTAL = 004				
			OUTBUF = 55 BLKS	OUTBUF THRES. TOTAL = 000				
			TRABUF = 36 BLKS	TRABUF THRES. TOTAL = 000				
0.011	228.86	0.38	UNDOT = 55 BLKS	UNDOT THRES. TOTAL = 000		0.001	297.80	3.30
			BPOOL = 55 BLKS	BPOOL THRES. TOTAL = 000				
				LINK THRES. TOTAL = 000				

Figure 12-4. TPE Summary Report

11/26/80 14.374 TRANSACTION PROCESSING TUNING GUIDE REPORTS START TIME 111980 158.2392
TAPE(S) 4976 STOP TIME 111980 308.7935

NONEW OPTION REPORT

COMMAND	TIME
NEW	1:59:42
NONEW	2:00:50
NONEW	2:02:37
NONEW	2:27:33
NEW	2:29:45
NEW	2:31:47
NONEW	2:54:45
NONEW	2:55:32
NEW	2:57:51
NEW	2:59:05
NEW	2:59:54
NEW	2:08:42

Figure 12-5. NONEW Option Report

02/28/80 14.766 TRANSACTION PROCESSING TUNING GUIDE REPORTS START TIME 011080 2208.6098
TAPE(S) 5565 STOP TIME 011080 2302.5013

PROFILE ID REPORT

PROFILE	STORED AT	ERASED AT
GM2		20:43:23
GM5	20:44:24	
GM5	20:45:13	

Figure 12-6. Profile ID Report

11/26/80 14.374 TRANSACTION PROCESSING TUNING GUIDE REPORTS START TIME 111980 158.2392
TAPE(S) 4976 STOP TIME 111980 300.7955

BUFFER SPACE REPORT

TIME	# OF BLOCKS ACQUIRED	# OF BLOCKS RELEASED
22:18:00	0	8
22:18:01	0	1
22:18:14	0	1
22:18:14	1	0
22:18:22	9	0
22:18:26	0	8
22:18:37	2	0
22:18:40	0	1
22:18:44	1	0
22:18:44	9	0
22:18:44	1	0
22:18:44	0	1
22:18:44	0	1
22:18:48	0	8
22:18:48	1	0
22:18:48	1	0
22:18:48	1	0
22:18:48	0	1
22:18:48	0	1
22:18:48	0	1
22:18:51	9	0
22:18:54	0	8

Figure 12-7. Buffer Space Report (Part 1 of 2)

BUFFER SPACE REPORT (contd)

TIME	# OF BLOCKS ACQUIRED	# OF BLOCKS RELEASED
22:18:55	0	1
22:19:06	0	1
22:19:06	1	0
22:19:06	9	0
22:19:11	0	8
22:19:24	9	0
22:19:27	0	8
22:19:36	1	0
22:19:36	9	0
22:19:36	1	0
22:19:37	0	1
22:19:37	0	1
22:19:47	0	8
22:19:47	1	0
22:19:47	1	0
22:19:47	1	0
22:19:47	0	1
22:19:47	0	1
22:19:47	0	1
22:19:56	0	1
22:20:05	9	0
22:20:09	0	8
22:20:14	0	1
22:20:14	1	0
22:20:20	9	0
22:20:23	0	8
22:20:43	2	0

Figure 12-7. (Part 2 of 2)

02/28/80 14.766 TRANSACTION PROCESSING TUNING GUIDE REPORTS START TIME 011080 2208.6098
TAPE(S) 5565 STOP TIME 011080 2302.5013

TABLE REACHING THRESHOLD SIZE

TABLE REACHING THRESHOLD	TIME
RECOVY	22:42:12

Figure 12-8. Table Reaching Threshold Size Report

02/28/80 14.766 TRANSACTION PROCESSING TUNING GUIDE REPORTS START TIME 011080 2208.6098
TAPE(S) 5565 STOP TIME 011080 2302.5013

MASTER TERMINAL REPORT

MASTER TERM ID	SIGN-ON TIME	DISCONNECT TIME	TERM ID ATTEMPTING	TIME ATTEMPTED
PB	20:34:01	20:50:02	PF	20:40:58
			PF	20:46:09
PB	20:59:57	21:00:27		

Figure 12-9. Master Terminal Report

12.4.2.7 Impasse Report. This report lists the time in hours, minutes and seconds that an impasse occurred in TPS, and an error message was sent to the operator's console. Possible problem areas are LINK, RECOVY, OUTBUF, TRABUF, UNDOT, INPUT and JC-STS. The first six names specify tables in which the fault may have occurred and are described in paragraph 12.4.2.5 above. JC-STS refers to an unrecoverable error on a journal file and should actually appear as JI-STS for the JOUR-IN file or JO-STS forthe JOUR-OUT file (DD36B, p. 3-5). Suggestions for avoiding IMPASSE faults are described in DD36B, p. 6-19. See figure 12-10.

12.4.2.8 Output Transmission Report. This report lists the terminal IDs that received output, the time transmitted in hours, minutes and seconds and the record pointers indicating the starting address within TPS of the output buffer. See figure 12-11.

12.4.2.9 TPAP ABORT Report. This report lists the TPAP transaction numbers that aborted, the time that they aborted in hours, minutes and seconds, and the reason for aborting. See figure 12-12.

12.4.2.10 UNDOT Entry Report. This report lists the times in hours, minutes and seconds that an UNDOT (undelivered output) entry was retrieved from the journal file. See figure 12-13.

12.4.2.11 Reissuing Intercom Message Report. This report lists the time in hours, minutes and seconds and the SNUMBs of programs reissuing an intercom message between TPE and the TPAPs or between TPAPs. See figure 12-14.

12.4.2.12 Receiving Intercom Message Report. This report lists the time in hours, minutes and seconds and program names receiving an intercom message. See figure 12-15.

12.4.2.13 TPAP Modify Report. This report lists each TPAP modifed, the modify code, and the time modified in hours, minutes and seconds. The items that can be modified are keywords, priority, flags and buffers. See figure 12-16.

12.4.2.14 TPAP Status Report. This report lists TPAPs of which the status was requested by the master terminal. The status, the status code and the time in hours, minutes and seconds the status was requested are also printed. See figure 12-17.

12.4.2.15 TPAP Run Times Report. This report lists all TPAPs that have started and ended within the time limit of this data reduction run. Listed are the start and stop times in hours, minutes and seconds, the elapsed time in seconds, and the processing time in seconds. A blank EOJ status indicates normal termination. "ABORTED" indicates the TPAP aborted. "TERMED" indicates the TPAP was terminated by an operator. See figure 12-18.

02/28/80 14.766 TRANSACTION PROCESSING TUNING GUIDE REPORTS START TIME 011080 2208.6098
TAPE(S) 5565 STOP TIME 011080 2302.5013

IMPASSE REPORT

REASON FOR "IMPASSE"	TIME
TRABUF	20:37:20
LINK	21:21:23

Figure 12-10. Impasse Report

OUTPUT TRANSMISSION REPORT

RECEIVING TERM ID	TRANSMITTED OUTPUT TIME	RECORD POINTER
PD	22:18:44	053526
PD	22:18:48	052626
PD	22:19:36	054026
PD	22:19:47	053126
PD	22:20:44	054426
PD	22:21:17	053526
PF	22:22:40	050526
PD	22:36:24	053126
PD	22:37:02	053126
PF	22:39:29	050426
PA	22:48:31	050726
PA	22:50:02	047726
PA	22:50:47	051026
PA	22:50:53	050026
PA	22:53:48	051126
PA	22:54:00	050126

Figure 12-11. Output Transmission Report

11/26/80 14.374 TRANSACTION PROCESSING TUNING GUIDE REPORTS START TIME 111980 158.2392
TAPE(S) 4976 STOP TIME 111980 308.7935

TPAP ABORT REPORT

ABORTED TPAP TRANS. NO	TIME	REASON CODE	REASON
22075	22:32:56	12	INTERACT MSG
22075	22:32:56	2	TPAP REPORT

Figure 12-12. TPAP ABORT Report

02/28/80 14.374 TRANSACTION PROCESSING TUNING GUIDE REPORTS START TIME 011080 2208.6098
TAPE(S) 5565 STOP TIME 011080 23∪2.5013

UNDOT ENTRY REPORT

UNDOT ENTRY RETRIEVED FROM JOURNAL FILE AT 20:50:36

Figure 12-13. UNDOT Entry Report

02/28/80 14.374 TRANSACTION PROCESSING TUNING GUIDE REPORTS START TIME 011080 2208.6098
TAPE(S) 5565 STOP TIME 011080 2302.5013

REISSUING INTERCOM MESSAGE REPORT

PROGRAM SNUMB REISSUING INTERCOM	TIME
GM500	20:15:12
GM200	20:18:06

Figure 12-14. Reissuing Intercom Message Report

RECEIVING INTERCOM MESSAGE REPORT

PROGRAM NAME RECEIVING INTERCOM	TERM ID	TIME OF TRANSFER
GM5	PC	20:17:42

Figure 12-15. Receiving Intercom Message Report

11/26/80 14.374 TRANSACTION PROCESSING TUNING GUIDE REPORTS START TIME 111980 158.2392
TAPE(S) 4976 STOP TIME 111980 308.7935

TPAP MODIFY REPORT

TPAP	MODIFIED TIME	MODIFIED ITEM
GMS	20:48:12	FLAGS
		BUFFERS

Figure 12-16. TPAP Modify Report

02/28/80 14.766 TRANSACTION PROCESSING TUNING GUIDE REPORTS START TIME 011080 2208.6098
TAPE(S) 5565 STOP TIME 011080 2302.5013

TPAP STATUS REPORT

TPAP SNUMB	STATUS CODE	STATUS	TIME STATUS REQUESTED
GM500	1	DATABASE HOLD	22:23:24
GM400	2	PURGE	22:30:23

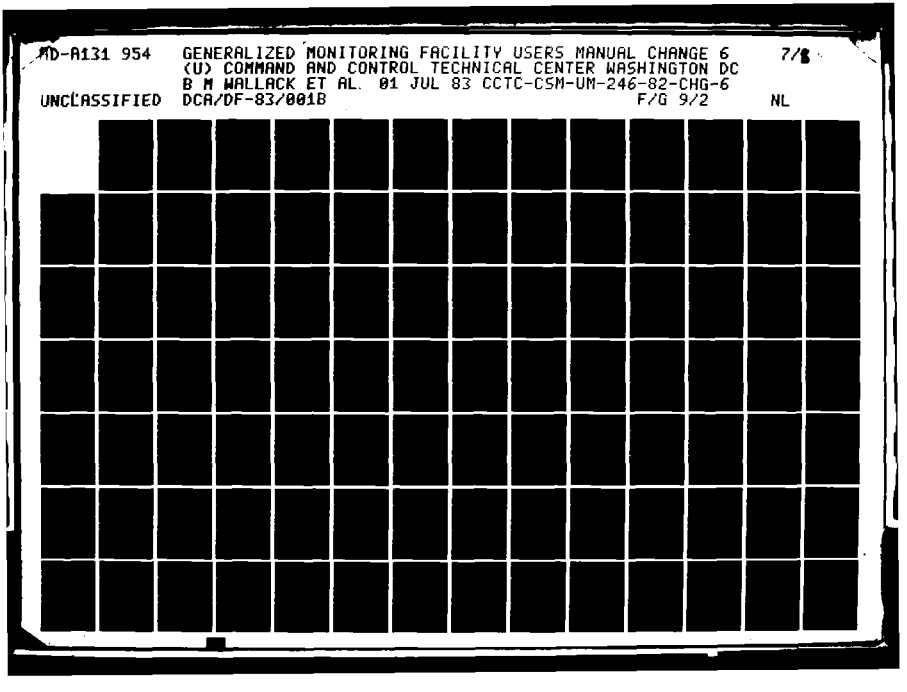
Figure 12-17. TPAP Status Report

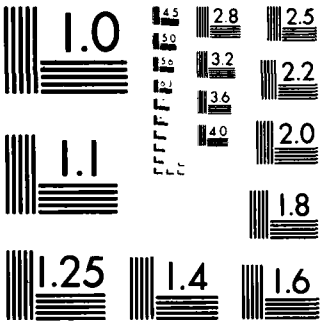
11/26/80 14.374 TRANSACTION PROCESSING TUNING GUIDE REPORTS START TIME 111980 308.7935
TAPE(S) 4976 STOP TIME 111980 322.3402

TPAP RUN TIMES

TPAP NAME	START TIME	STOP TIME	LAPSE	PROC TIME	EOJ STATUS
GM300	2:00:44	2:00:47	2.785	4.1547	
TP200	2:01:29	2:01:31	1.637	2.2912	
TP200	2:02:22	2:02:24	2.498	1.4500	
INT00	2:20:18	2:20:20	2.050	0.	
INT00	2:53:32	2:53:34	1.988	0.	
GM300	2:54:41	2:54:44	2.539	0.	
GM300	2:55:41	2:58:44	182.912	0.	
TP200	3:01:10	3:01:12	1.834	0.	
TP200	3:01:50	3:01:52	1.656	0.	

Figure 12-18. TPAP Run Times Report


AD-A131 954
GENERALIZED MONITORING FACILITY USERS MANUAL CHANGE 6
(U) COMMAND AND CONTROL TECHNICAL CENTER WASHINGTON DC
B M WALLACK ET AL. 01 JUL 83 CCTC-CSM-UM-246-82-CHG-6
UNCLASSIFIED
DCA/DF-83/001B
F/G 9/2
NL



MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS 1963-A

12.5 TPE Tuning Guide Formatted Dump

This section will contain the running instructions for producing the TPE Tuning Guide formatted dump. It will also include the input and output formats.

12.5.1 Input. To produce the TPSDRP formatted dump, the file B29IDPX0/OBJECT/TPEDUMP must be run. Input parameter cards may be used. They must be inserted immediately after the $DATA I* card in the JCL. If the input parameter cards are omitted, all the data will be dumped. There are five input options available: MAXOUT, TIME ,SKIP, OPTION and LIMIT. The first card has the input option type to follow (e.g., MAXOUT). The second card gives the number of entries to follow on the third card (applies only to the TIME and LIMIT options). The third card contains the input parameter. The input is in free format. A description of each follows:

- MAXOUT - This option is a limit to the maximum number of traces to be printed. Even though the data will not be printed, it still will be collected for the summary report. The default is 30,000. The second card will contain the new limit.

- TIME - This is a time window option. The second card is used for this option. It gives the number of times to follow. The third card contains a maximum of 5 pairs of START/STOP times.

  a. START - This is a time of day on a 2400-hour clock. The format is HH:MM. It can be used to skip printing traces created before a certain time. The data for the traces skipped will also be collected for the summary report. To start at the beginning of the tape when using a STOP time, use 00:00 for the START time.

  b. STOP - This is a time of day on a 2400-hour clock. The format is HH:MM. It can be used to skip printing traces created after a certain time. The data for the traces skipped will also be skipped for the summary report. At midnight, the clock will roll over to 00, so this time may be less than the START time.

- SKIP - This option is for skipping traces before printing is to begin. The data for these traces will also be collected for the summary report. The default is 0. The second card will contain the number of traces to be skipped.

- OPTION - This is a print option. If this option is not present, the traces will be printed. If present, only the summary report will be printed. No additional cards are required.

- LIMIT - This option is a limit to the maximum number of a particular trace type to be printed. The second card is used for this option.

It gives the number of trace types to follow. The third card contains FIRSTX and the trace types (NOPRNT).

a. FIRSTX - This works in conjunction with the NOPRNT. It determines how many to print of each trace type specified in NOPRNT. The default is 0.

b. NOPRNT - This specifies up to 10 trace types for which the printing can be limited. The default is 0.

Examples of the input cards are as follows:

```
example A
MAXOUT (Card 1)
40000 (Card 2)

example B
OPTION (Card 1)

example C
LIMIT (Card 1)
2 (Card 2)
100 1,3 (Card 3)
```

In example A, the user is setting the maximum number of traces to be printed at 40,000. In example B, the user is eliminating printing the traces and only the summary report will be printed. In example C, the user is printing only 100 each of trace type 1 and trace type 3.

After choosing the options desired, the $TAPE card must be altered to insert the TPE Data Collector tape number before executing the program. A list of the JCL is shown in figure 12-19.

12.5.2 Formatted Output. The formatted output contains the TPE trace number in decimal, and the TPE trace number (in octal) in the first half of the first word and the TSS trace number (in octal) in the second half of the first word of the trace itself. Also, the remaining two to six words of the trace itself are printed. Printed next is the title of the TPE trace (a 20-character field). Also printed is the trace time converted to seconds (trace time is microseconds originally) minus the initialization time, and the trace time of day in a 24-hour day. Any BCD representations in the various traces will be displayed in readable format also. This report is written to report code 08. An example formatted dump is shown in figure 12-20.

12.5.3 Summary Report. There is a small summary report written to report code 09 which lists a count of all the traces found and the times (minus the initialization time) in seconds of the first four traces of each trace type (also shown in figure 12-20).

```
$           IDENT
$           LOWLOAD
$           SELECT  B29IDPX0/OBJECT/TPEDUMP
$           EXECUTE DUMP
$           MSG2    1,PLEASE ENTER "U" IF TAPE
$           MSG2    1,ERRORS OCCUR
$           LIMITS  10,26K,-2K,30K
$           FFILE   P*,LGU/(06,08,09,23,24,25,29,42)
$           FFILE   10,NSTDLB,NOSRLS,BUFSIZ/4094,FIXLNG/4094,ERRXIT/KILLFM
$           TAPE    10,X1D,,,,NO-RING
$           DATA    I*
MAXOUT
40000
$           ENDJOB
```

Figure 12-19. JCL for TPE Formatted Dump

9	000011000074	000000000011	004544536115	BUFFER SPACE ACQUIRD		630.373451	22.3186
13	000015000074	000010000033	004566316117	BUFFER SPACE RELEASED		635.018318	22.3199
3	000003000074	202020205252	004650242171	BREAK REQUEST NOTICED		648.103035	22.3235
9	000011000074	000000000011	004650245025	BUFFER SPACE ACQUIRD		648.104469	22.3235
13	000015000074	000010000034	004664664256	BUFFER SPACE RELEASD		651.389099	22.3244
9	000011000074	000000000001	004726525600	BUFFER SPACE ACQUIRD		660.253571	22.3269
20	000024000074	001515777777	004726535664	READ INSTR PERFORMED		660.257713	22.3269
9	000011000074	000000000011	004726552631	BUFFER SPACE ACQUIRD		660.264343	22.3269
9	000011000074	000000000001	004726560461	BUFFER SPACE ACQUIRD		660.264992	22.3269
14	000016000074	054026004724	004726560461	OUTPUT TRANSMITTED		660.267311	22.3269
3	000003000074	202020205252	004730167613	BREAK REQUEST NOTICED		660.664200	22.3270
13	000015000074	000001000034	004730177137	BUFFER SPACE RELEASD		660.667999	22.3270
13	000015000074	000001000046	004730200363	BUFFER SPACE RELEASD		660.668655	22.3270
13	000015000074	000010000036	004776117310	BUFFER SPACE RELEASD		670.605003	22.3298
9	000011000074	000000000001	004776126464	BUFFER SPACE ACQUIRD		670.608688	22.3298
9	000011000074	000000000001	004776157554	BUFFER SPACE ACQUIRD		670.621552	22.3298
5	000005000074	000004000041	004776166163	UNDOT ENTRY CHANGES		670.624886	22.3298
9	000011000074	000000000001	004776204326	BUFFER SPACE ACQUIRD		670.632149	22.3298
14	000016000074	053126004724	004776210564	OUTPUT TRANSMITTED		670.634354	22.3298
13	000015000074	000001000035	004776263120	BUFFER SPACE RELEASD		670.656082	22.3298
13	000015000074	000001000036	004776345545	BUFFER SPACE RELEASD		670.681953	22.3298
13	000015000074	000001000037	004777706240	BUFFER SPACE RELEASD		671.059105	22.3299
13	000015000074	000001000001	005041522113	BUFFER SPACE RELEASD		679.912521	22.3324
5	000005000074	000004000041	005041523333	UNDOT ENTRY CHANGES		679.913177	22.3324
44	000054000074	000000000001	005054566235	OPERATOR COMMAND		682.814621	22.3332
3	000003000074	202020205252	005103735644	BREAK REQUEST NOTICED		688.896935	22.3349
9	000011000074	000000000011	005103740510	BUFFER SPACE ACQUIRD		688.898376	22.3349
19	000023000074	001515777777	005104261026	WRITE INSTR PERFORMD		689.005081	22.3349
13	000015000074	000010000036	005123235552	BUFFER SPACE RELEASD		692.927338	22.3360
7	000007000074	634702000000	005123407021	NAME OF TPAP SPAWNED	TPAP	692.981262	22.3360
44	000054000074	000000000024	005134361521	OPERATOR COMMAND		695.329620	22.3366
19	000023000074	001515777777	005134362724	WRITE INSTR PERFORMD		695.330261	22.3366
20	000024000074	001515777777	005144475430	READ INSTR PERFORMED		697.465622	22.3372
13	000015000074	000001000035	005144476563	BUFFER SPACE RELEASD		697.466225	22.3372
9	000011000074	000000000001	005144606033	BUFFER SPACE ACQUIRD		697.502747	22.3372
5	000005000074	000004000041	005144614422	UNDOT ENTRY CHANGES		697.506065	22.3372
20	000024000074	001515777777	005144620042	READ INSTR PERFORMED		697.507874	22.3372
3	000003000074	202020205252	005175507103	BREAK RQUEST NOTICED		704.024132	22.3391

Figure 12-20. TPE Formatted Dump (Part 1 of 2)

9	000011000074	000000000011	005175511740	BUFFER SPACE ACQUIRD	704.025566	22.3391
13	000015000074	000010000036	005211371213	BUFFER SPACE RELEASD	707.129997	22.3399
3	000003000074	202020205252	005323473716	BREAK RQUEST NOTICED	726.562767	22.3453
9	000011000074	000000000002	005323476600	BUFFER SPACE ACQUIRD	726.564224	22.3453
9	000011000074	000000000001	005327262142	BUFFER SPACE ACQUIRD	727.540833	22.3456
20	000024000074	001515777777	005327273155	READ INSTR PERFORMED	727.545456	22.3456
9	000011000074	000000000011	005327310150	BUFFER SPACE ACQUIRD	727.552101	22.3456
9	000011000074	000000000001	005327311361	BUFFER SPACE ACQUIRD	727.552750	22.3456
14	000016000074	054426004724	005327316000	OUTPUT TRANSMITTED	727.555069	22.3456
3	000003000074	202020205252	005330721553	BREAK RQUEST NOTICED	727.950188	22.3457
13	000015000074	000001000040	005330731040	BUFFER SPACE RELEASD	727.953949	22.3457
13	000015000074	000001000052	005330732253	BUFFER SPACE RELEASD	727.954605	22.3457

TRACE	TOTAL	FIRST FOUR TIMES			
1	8454	2.855432	3.109425	3.363381	3.618357
2	6	49.237865	359.761200	1851.296417	1942.519775
3	807	53.812205	75.882420	97.370882	108.188192
4	3	1460.888000	1916.062317	2735.896271	0.
5	38	66.907244	103.061275	113.680939	143.884554
7	20	6.100223	57.483391	58.140926	90.054187
9	131	53.816041	66.904400	75.883858	90.052096
11	1	228.329462	0.	0.	0.
13	138	6.072093	58.028220	66.780342	89.983652
14	26	96.972860	103.070342	173.941208	189.204172
19	42	8.542897	60.723680	62.633680	93.981119
20	52	66.779792	66.909051	96.961090	143.844147
21	1	1459.802460	0.	0.	0.
22	2	1459.799347	1459.840607	0.	0.
32	7	2015.739761	2015.740311	2015.740854	2015.741394
33	2	869.425728	2414.070435	0.	0.
44	41	8.542246	12.512819	60.723033	62.633066
45	23	928.069572	1028.087509	1046.402167	1084.747772
52	1	2225.681030	0.	0.	0.

9795 TOTAL T74 TRACES

THE PRECEDING FORMATTED DATA PERTAINS TO TAPE(S) 5565
PRODUCED FROM 011080/ 2215.0000
TO 011080/ 2300.0000

Figure 12-20. (Part 2 of 2)

```
$       GMAP                                                                    TRXA
$       ALTER   530,530
***SITE***SITE***SITE***SITE***SITE***SITE***SITE***SITE***SITE***SITE*****GMF**
*                                                                               ****GMF**
*         THIS SYMBOL CONTROLS ASSEMBLY OF THE TPE INTERFACE WITH               ****GMF**
*         THE CCTC/WAD GENERALIZED MONITORING FACILITY (GMF)                    ****GMF**
*                                                                               ****GMF**
*         SET THIS SYMBOL TO THE SAME VALUE IN TRXB,C, AND TRXD                 ****GMF**
GMF     SET     1                                                               ***GMF**
*                                                                               ****GMF**
***SITE***SITE***SITE***SITE***SITE***SITE***SITE***SITE***SITE***SITE*****GMF**
    (NTPE ALTER cards)
$       ALTER   5652                                                            ***GMF**
        IFE     GMF,1,%.GMF                                                     ***GMF**
        TTLS    PERFORMANCE ANALYSIS TRACES OF TPS                              ***GMF**
        SYMDEF  STTPTR,TREP,MAKTR,EXIT1,JSTOFF,GMFTPO                           ***GMF**
        SYMREF  UDOTVL,TP3015,TP3016,TP3023,TP3115,.LLINE,TP2055                ***GMF**
        SYMREF  TC0950,OUTAVL,MAPTRA,MAPPIO,MAPINP,RECAVL                       ***GMF**
STTPTR  NULL                                                                    ***GMF**
        INHIB   ON                                                              ***GMF**
JSTOFFESTC1     IC              SAVE IC+I                                       ***GMF**
        TRA     *+1             GET OUT OF XED                                  ***GMF**
        SREG    SREGS           SAVE REGISTERS                                  ***GMF**
        EAA     190             PLACE T190 TYPE IN AR                           ***GMF**
        STCA    TRACE1+1,70     STORE SUBTRACE NO                               ***GMF**
        LDA     TRACE1+1                                                        ***GMF**
        EAQ     0                                                               ***GMF**
        TRA     MAKTR           GO MAKE TRACE T190                              ***GMF**
TREP    ESTC1   IC              SAVE IC+I                                       ***GMF**
        TRA     *+1             GET OUT OF XED                                  ***GMF**
        SREG    SREGS           SAVE REGISTERS                                  ***GMF**
        LDA     IC,I            GET PATTERN #, TRACE TYPE                       ***GMF**
        LDX0    1,DU                                                            ***GMF**
        ASX0    IC                                                              ***GMF**
        SZN     GMFTPO          CHECK TP TRACE FLAG ON                          ***GMF**
        TZE     EXIT1           NO                                              ***GMF**
        STCA    TRACE1+1,70     STORE SUBTRACE TYPE                             ***GMF**
        EAX4    0,AU                                                            ***GMF**
        TSX1    TABLE-1,4*      TRANSFER TO INDIVIDUAL TRACE                    ***GMF**
MAKTR   MME     .EMM                                                            ***GMF**
        XEC     .CRTRC+2        CHECK TRACE MECHANISM ON/OFF                    ***GMF**
        TSS     EXIT1           TRACE IS OFF/SEE NOTE BELOW                     ***GMF**
* NOTE-DURING OPERATION ABOVE SHOULD NEVER BE TAKEN SINCE IF                     ***GMF**
* TRACE MECHANISM IS OFF, SO WILL BE T74 AND WE WON'T GET                       ***GMF**
* THIS FAR.                                                                     ***GMF**
        LDX0    .CRTRC+3        ADDRESS —| X0                                   ***GMF**
        TSX1    1,0*            MAKE TRACE ENTRY                                ***GMF**
EXIT1   LREG    SREGS,$         RELOAD REG;CNTRL RETURNS HERE AFTER TRACE***GMF**
```

Figure 12-21. Transaction Processing System ALTER Cards (Part 1 of 11)

```
        RET     IC,$            RETURN TO TPS                                   ***GMF**
        EJECT                                                                   ***GMF**
*                                                                               ***GMF**
******************      TRXD TRACES      ******************                     ***GMF**
*                                                                               ***GMF**
T74.1   NULL                                                                    ***GMF**
        LCQ     1,DL            WORD 3 OF TRACE IS EMPTY                        ***GMF**
        LDA     TRACE1+1                                                        ***GMF**
        TRA     0,1                                                             ***GMF**
T74.2   NULL                                                                    ***GMF**
        LDQ     SREGS+5         MAX # TERMS ALLOWED, # TERMS CONNECTED          ***GMF**
        LDA     TRACE1+1                                                        ***GMF**
        TRA     0,1                                                             ***GMF**
T74.5   NULL                                                                    ***GMF**
        LDQ     UDOTVL          NO. UNDOT ENTRIES AVAILABLE                     ***GMF**
        LDA     TRACE1+1                                                        ***GMF**
        TRA     0,1                                                             ***GMF**
T74.7   NULL                                                                    ***GMF**
        LXL4    SREGS                                                           ***GMF**
        STZ     TRACE1+2                                                        ***GMF**
        STX4    TRACE1+2        TPAP ID IN 0-17 WORD 3                          ***GMF**
        LDAQ    TRACE1+1                                                        ***GMF**
        TRA     0,1                                                             ***GMF**
T74.11  NULL                                                                    ***GMF**
        LDAQ    SREGS+4         STORE INVALID KEYWORD                           ***GMF**
        STA     TRACE1+4          IN TRACE WORDS 5 AND 6                        ***GMF**
        STQ     TRACE1+5          IN ASCII                                      ***GMF**
        LCQ     1,DL                                                            ***GMF**
        STQ     TRACE1+2        IN 18-35 OF TRACE WD 3                          ***GMF**
        LXL3    SREGS+1                                                         ***GMF**
        LXL4    .ID,3           ID OF USER ENTERING                             ***GMF**
        STX4    TRACE1+2          INVALID KEYWORD IN 0-17 WD 3                  ***GMF**
        LDAQ    TRACE1+1                                                        ***GMF**
        EAX3    TRACE1+4                                                        ***GMF**
        TRA     0,1                                                             ***GMF**
T74.14  NULL                                                                    ***GMF**
        LDQ     TP2055          RECORD POINTER                                  ***GMF**
        ANQ     =0777777007777    AND                                           ***GMF**
        STQ     TRACE1+2          TERMINAL ID                                   ***GMF**
        LDA     TRACE1+1                                                        ***GMF**
        TRA     0,1                                                             ***GMF**
T74.17  NULL                                                                    ***GMF**
        LDX4    SREGS+2                                                         ***GMF**
        LDQ     .SPWN+1,4       SNUMB                                           ***GMF**
        LDA     TRACE1+1                                                        ***GMF**
        TRA     0,1                                                             ***GMF**
T74.18  NULL                                                                    ***GMF**
        LDA     TP3015                                                          ***GMF**
```

Figure 12-21. (Part 2 of 11)

```
         ALS     24                                          ***GMF**
         ARL     6                                           ***GMF**
         STCA    TRACE1+2,70     CURRENT LINE ID             ***GMF**
         LDX2    TP3016                                      ***GMF**
         SXL2    TRACE1+2        NEW PROGRAM NAME            ***GMF**
         LDAQ    TRACE1+1                                    ***GMF**
         TRA     0,1                                         ***GMF**
T74.19   NULL                    ***GMF**
         LDX4    TP3023          ADDRESS OF FILE             CODE:***
*        LCQ     1,DL                                        ***GMF**
*        STQ     TRACE1+2                                    ***GMF**
         LDA     0,4                                         ***GMF**
*        ALS     18                                          ***GMF**
         STA     TRACE1+2        FC WRITTEN FROM             ***GMF**
         LDAQ    TRACE1+1                                    ***GMF**
         TRA     0,1                                         ***GMF**
T74.20   NULL                                                ***GMF**
*        LCQ     1,DL                                        ***GMF**
*        STQ     TRACE1+2                                    ***GMF**
         LDX4    TP3115          GET ADDRESS OF FC           ***GMF**
         LDA     0,4                                         ***GMF**
*        ALS     18                                          ***GMF**
         STA     TRACE1+2        FC ISSUING READ             ***GMF**
         LDAQ    TRACE1+1                                    ***GMF**
         TRA     0,1                                         ***GMF**
T74.21   NULL                                                ***GMF**
         LDQ     SREGS+4         TPAP NAME                   ***GMF**
         LDA     TRACE1+1                                    ***GMF**
         TRA     0,1                                         ***GMF**
T74.22   NULL                                                ***GMF**
         LDAQ    SREGS+4                                     ***GMF**
         STA     TRACE1+2        TRANSACTION #               ***GMF**
         QLS     6                                           ***GMF**
         LXL3    TRACE1+1        SAVE FORMAT SANS CODE       ***GMF**
         STCQ    TRACE1+1,06     CODE FOR REASON BOMBED      ***GMF**
         LDAQ    TRACE1+1                                    ***GMF**
         SXL3    TRACE1+1        RESTORE PRISTINE FORMAT     ***GMF**
         TRA     0,1                                         ***GMF**
*                                                            ***GMF**
**********          TRXC TRACES          **********         ***GMF**
*                                                            ***GMF**
T74.31   NULL                                                ***GMF**
         LDA     SREGS+5         FLAG IN 3L, SNUMB IN 4      ***GMF**
         LDQ     SREGS+4                                     ***GMF**
*        ANQ     =0777777777700  GET SNUMB                   ***GMF**
         STQ     TRACE1+2                                    ***GMF**
*        ANA     7,DL                                        ***GMF**
         STCA    TRACE1+2,01     FLAG                        ***GMF**
```

Figure 12-21. (Part 3 of 11)

```
        LDAQ    TRACE1+1                                                ***GMF**
        TRA     0,1                                                     ***GMF**
T74.32  NULL                                                            ***GMF**
        LXL3    SREGS                                                   ***GMF**
        LDA     TC0950+4,3      BUFFER ENCOUNTERING THRESHOLD           ***GMF**
        LDQ     TC0950+5,3      # THRESHOLDS REACHED                    ***GMF**
        STCA    TRACE1+2,07                                             ***GMF**
        STCQ    TRACE1+2,70                                             ***GMF**
        LDAQ    TRACE1+1                                                ***GMF**
        TRA     0,1                                                     ***GMF**
T74.33  NULL                                                            ***GMF**
        LXL3    SREGS+1                                                 ***GMF**
        LDQ     .MIT,3          MASTER TERMINAL TYPE + ID               ***GMF**
        LDA     TRACE1+1                                                ***GMF**
        TRA     0,1                                                     ***GMF**
T74.34  NULL                                                            ***GMF**
        LDX0    SREGS                                                   ***GMF**
        SXL0    TRACE1+2        MASTER TERMINAL ID                      ***GMF**
        LXL0    SREGS+4                                                 ***GMF**
        STX0    TRACE1+2        NEW SIGNON ID                           ***GMF**
        LDAQ    TRACE1+1                                                ***GMF**
        TRA     0,1                                                     ***GMF**
*                                                                       ***GMF**
***************    INITIALIZE TRACE     ****************              ***GMF**
*                                                                       ***GMF**
T74.52  NULL                                                            ***GMF**
        LDX0    OUTAVL          OUTBUF                                  ***GMF**
        STX0    TRACE1+2                                                ***GMF**
        LDX0    MAPTRA+1        TRABUF                                  ***GMF**
        SXL0    TRACE1+2                                                ***GMF**
        LDX0    MAPPIO+1        BPOOL                                   ***GMF**
        STX0    TRACE1+4                                                ***GMF**
        LDX0    MAPINP+1        INPUT                                   ***GMF**
        SXL0    TRACE1+4                                                ***GMF**
        LDX0    UDOTVL          UNDOT                                   ***GMF**
        STX0    TRACE1+5                                                ***GMF**
        LDX0    RECAVL          RECOVY                                  ***GMF**
        SXL0    TRACE1+5                                                ***GMF**
        LDAQ    TRACE1+1                                                ***GMF**
        EAX3    TRACE1+4        2 EXTRA WORDS IN THIS TRACE             ***GMF**
        TRA     0,1                                                     ***GMF**
*                                                                       ***GMF**
**********************     VARIABLES     **********************         ***GMF**
*                                                                       ***GMF**
TRACE100CT      000004000174  RECORD SIZE, MEDIA CODE, TRACE TYPE        ***GMF**
        OCT     74              SUBTRACE, TRACE                         ***GMF**
        OCT     202020202020    VARYING INFORMATION                     ***GMF**
```

Figure 12-21. (Part 4 of 11)

```
        OCT     0               TIME FROM RSCR INSTRUCTION              ***GMF**
        OCT     0               USED IN TRACES 11, 52 ONLY              ***GMF**
        OCT     0               USED IN TRACES 11, 52 ONLY              ***GMF**
SREGS 8BSS      8                                                       ***GMF**
IC      BSS     1                                                       ***GMF**
TABLE   TRA     T74.1           TABLE OF TRANSFERS TO SUBTRACE ROUTINES ***GMF**
        TRA     T74.2                                                   ***GMF**
        TRA     T74.1                                                   ***GMF**
        TRA     T74.1                                                   ***GMF**
        TRA     T74.5                                                   ***GMF**
        TRA     T74.21                                                  ***GMF**
        TRA     T74.7                                                   ***GMF**
        TRA     T74.1                                                   ***GMF**
        TRA     T74.2                                                   ***GMF**
        TRA     T74.2                                                   ***GMF**
        TRA     T74.11                                                  ***GMF**
        TRA     T74.1                                                   ***GMF**
        TRA     T74.2                                                   ***GMF**
        TRA     T74.14                                                  ***GMF**
        ZERO                    T74.15                                  ***GMF**
        TRA     T74.1                                                   ***GMF**
        TRA     T74.17                                                  ***GMF**
        TRA     T74.18                                                  ***GMF**
        TRA     T74.19                                                  ***GMF**
        TRA     T74.20                                                  ***GMF**
        TRA     T74.21                                                  ***GMF**
        TRA     T74.22                                                  ***GMF**
        ZERO                    T74.23                                  ***GMF**
        ZERO                    T74.24                                  ***GMF**
        ZERO                    T74.25                                  ***GMF**
        ZERO                    T74.26                                  ***GMF**
        ZERO                    T74.27                                  ***GMF**
        ZERO                    T74.28                                  ***GMF**
        ZERO                    T74.29                                  ***GMF**
        ZERO                    T74.30                                  ***GMF**
        TRA     T74.31                                                  ***GMF**
        TRA     T74.32                                                  ***GMF**
        TRA     T74.33                                                  ***GMF**
        TRA     T74.34                                                  ***GMF**
        TRA     T74.2                                                   ***GMF**
        ZERO                    T74.36                                  ***GMF**
        ZERO                    T74.37                                  ***GMF**
        ZERO                    T74.38                                  ***GMF**
        ZERO                    T74.39                                  ***GMF**
        TRA     T74.1                                                   ***GMF**
        TRA     T74.2                                                   ***GMF**
        TRA     T74.2                                                   ***GMF**
        TRA     T74.2                                                   ***GMF**
```

Figure 12-21. (Part 5 of 11)

```
        TRA     T74.21                                                  ***GMF**
        TRA     T74.2                                                   ***GMF**
        ZERO                    T74.46                                  ***GMF**
        ZERO                    T74.47                                  ***GMF**
        ZERO                    T74.48                                  ***GMF**
        ZERO                    T74.49                                  ***GMF**
        ZERO                    T74.50                                  ***GMF**
        TRA     T74.1                                                   ***GMF**
        TRA     T74.52                                                  ***GMF**
.ID     EQU     4               TERM ID IN INPUT BUFFER                 ***GMF**
.MIT    EQU     6               WORD 6, MASTER TERMTYPE AND TERM ID     ***GMF**
.SPWN   EQU     5               POSITION OF SPAWN INFO IN .QTT          ***GMF**
GMFTPO  DEC     0                                                       ***GMF**
.GMF    MARK                                                            ***GMF**
$       GMAP                                                            TRXB
$       ALTER   2417
        TTLS    ****BYE TERMINAL DISCONNECT COMMAND***
 GMF    SET     1
TRMAK   MACRO
        IFE     GMF,1,2
        XED     TREP
        VFD     18/#1,8160      0-17 DB.TYPE 18-35 TYPE74
        ENDM    TRMAC
    (NTPE ALTER CARDS)
        IFE     GMF,1,1
        TRMAC   4
    (NTPE ALTER CARDS)
        LDA     0,2
        ARL     18
        IFE     GMF,1,1
        TRMAC
$       GMAP                                                            TRXC
$       ALTER   125
    (NTPE ALTER cards)
***SITE***SITE***SITE***SITE***SITE***SITE***SITE***SITE***SITE***SITE*****GMF**
*                                                                       ****GMF**
*         THIS SYMBOL CONTROLS ASSEMBLY OF THE TPE INTERFACE WITH       ****GMF**
*         THE CCTC/WAD GENERALIZED MONITORING FACILITY (GMF)            ****GMF**
*                                                                       ****GMF**
*         SET THIS SYMBOL TO THE SAME VALUE IN TRXA,C, AND TRXD         ****GMF**
GMF     SET     1                                                        ***GMF**
*                                                                       ****GMF**
***SITE***SITE***SITE***SITE***SITE***SITE***SITE***SITE***SITE***SITE*****GMF**
        IFE     GMF,1,2                                                 ***GMF**
        SYMDEF  TC0950,TC0030,MCOM                                      ***GMF**
        SYMREF  TREP                                                    ***GMF**
        IFE     OPMESS,0,1                                              ***GMF**
GMF     SET     1                                                       ***GMF**
TRMAC   MACRO                                                           ***GMF**
        IFE     GMF,1,2                                                 ***GMF**
        XED     TREP            MAKE THE TRACE                          ***GMF**
        VFD     18/#1,18/60     18/SUBTRACE,18/MAJOR TRACE  =074        ***GMF**
        ENDM    TRMAC                                                   ***GMF**
```

Figure 12-21. (Part 6 of 11)

```
$       ALTER   187                                                     ***GMF**
        TRMAC   44              OPERATOR COMMAND                        ***GMF**
$       ALTER   1095                                                    ***GMF**
        TRMAC   31                                                      ***GMF**
$       ALTER   1448                                                    ***GMF**
*                                                                       ***GMF**
*   TRACE ON/OFF                                                        ***GMF**
*                                                                       ***GMF**
        IFE     GMF,1,%.GMF                                             ***GMF**
        SYMREF  GMFTPO,JSTOFF                                           ***GMF**
        STQ     STORQ           SAVE QREG                               ***GMF**
        LDA     TC0693                                                  ***GMF**
        CMPA    =6HTRACE        IS THIS A TRACE COMMAND                 ***GMF**
        TNZ     TCGMF1          NO                                      ***GMF**
        LDA     TC0693+1        YES                                     ***GMF**
        CMPA    =2HON           TRACE-ON REQUEST                        ***GMF**
        TZE     TCGMF2          YES                                     ***GMF**
        CMPA    =3HOFF          TRACE-OFF REQUEST                       ***GMF**
        TNZ     TCGMF1          NOT A TRACE COMMAND AFTER ALL           ***GMF**
        SZNC    GMFTPO          IS TRACE ALREADY OFF                    ***GMF**
        TZE     *+2             YES, FORGET THIS                        ***GMF**
        XED     JSTOFF          MAKE TRACE TYPE 190                     ***GMF**
        EAA     GMFTOF          DCW ADDRESS                             ***GMF**
        TRA     TCGMF3          GO SEND ACKNOWLEDGMENT                  ***GMF**
TCGMF2  STC2    GMFTPO          TRACE ON, SET FLAG NONZERO              ***GMF**
        TRMAC   52              TRACE BUFFER SIZES CALCULATED           ***GMF**
        EAA     GMFTON          DCW ADDRESS                             ***GMF**
TCGMF3  ARL     18              POSITION DCW ADDRESS                    ***GMF**
        STCA    GMFRIT,07       STORE IN GEINOS SEQUENCE                ***GMF**
        MME     GEINOS          SEND ACKNOWLEDGMENT TO CONSOLE          ***GMF**
        WTYP                                                            ***GMF**
GMFRIT  ZERO    TC0120,**                                               ***GMF**
        ZERO                                                            ***GMF**
        MME     GEENDC                                                  ***GMF**
GMFTON  IOTD    *+1,5                                                   ***GMF**
        OCT     770117171717                                            ***GMF**
        BCI     4,*TPE GMF TRACE STARTED                                ***GMF**
GMFTOF  IOTD    *+1,5                                                   ***GMF**
        OCT     770117171717                                            ***GMF**
        BCI     4,*TPE GMF TRACE STOPPED                                ***GMF**
STORQ   BSS     1                                                       ***GMF**
TCGMF1  LDQ     STORQ                                                   ***GMF**
.GMF    MARK                                                            ***GMF**
$       ALTER   1748                                                    ***GMF**
        TRMAC   32              TRACE OF TABLE REACHING THRESHOLD       ***GMF**
$       ALTER   2235                                                    ***GMF**
        TRMAC   33              TRACE LID OF MASTER TERMINAL SIGNED ON  ***GMF**
$       ALTER   2253                                                    ***GMF**
        TRMAC   34              TRACE 2 TERMS WANTING TO BE MASTER      ***GMF**
$       ALTER   2332,2332
```

Figure 12-21. (Part 7 of 11)

```
      (NTPE ALTER cards)
         IFE     GMF,1,2                                                ***GMF**
         LDQ     MBUF+.MIT       TERM TYPE AND TERM ID                  ***GMF**
         TRMAC   35              TRACE LID OF MASTER TERM DISCONNECT    ***GMF**
      (NTPE ALTER cards)
$        ALTER   2402,2403
         IFE     GMF,1,2                                                ***GMF**
         LDQ     MCOM,$                                                 ***GMF**
         TRMAC   45              MASTER TERM COMMAND                    ***GMF**
      (NTPE ALTER cards)
$        ALTER   3695,3696
      (NTPE ALTER cards)
         IFE     GMF,1,2
         LDQ     .MID
         TRMAC   35
$        ALTER   4530                                                   ***GMF**
         TRMAC   40              TRACE # KEYWD TABLE COMPRESSIONS ATTEMPTED0000000
$        ALTER   4821                                                   ***GMF**
         TRMAC   41              TRACE ID OF PROFILE STORED             ***GMF**
$        ALTER   5289                                                   ***GMF**
         TRMAC   42              TRACE ID OF PROFILE ERASED             ***GMF**
$        ALTER   5444                                                   ***GMF**
         IFE     GMF,1,1                                                ***GMF**
         SXL2    TRCFLG          TPAP NAME IN BCD                       ***GMF**
$        ALTER   5729                                                   ***GMF**
         IFE     GMF,1,2                                                ***GMF**
         EAA     8               KEYWORDS                               ***GMF**
         ASA     TRCFLG                                                 ***GMF**
$        ALTER   5751                                                   ***GMF**
         IFE     GMF,1,2                                                ***GMF**
         EAA     4               PRIORITY                               ***GMF**
         ASA     TRCFLG                                                 ***GMF**
$        ALTER   5808                                                   ***GMF**
         IFE     GMF,1,2                                                ***GMF**
         EAA     2               FLAGS                                  ***GMF**
         ASA     TRCFLG                                                 ***GMF**
$        ALTER   5817                                                   ***GMF**
         IFE     GMF,1,1                                                ***GMF**
         TZE     TM169Z          NO BUFFER SIZE CHANGES - DONE - WRITE REP***GMF**
         INE     GMF,1,1                                                ***GMF**
$        ALTER   5827                                                   ***GMF**
         IFE     GMF,1,3                                                ***GMF**
         EAA     1               FIXED BUFFERS                          ***GMF**
         ASA     TRCFLG                                                 ***GMF**
         TRA     TM169Z          DONE - WRITE REPORT                    ***GMF**
         INE     GMF,1,1                                                ***GMF**
$        ALTER   5832                                                   ***GMF**
         IFE     GMF,1,2                                                ***GMF**
         EAA     1               DYNAMIC BUFFERS                        ***GMF**
         ORSA    TRCFLG                                                 ***GMF**
```

Figure 12-21. (Part 8 of 11)

```
$       ALTER   5838                                                    ***GMF**
        IFE     GMF,1,1                                                 ***GMF**
        TZE     TM169Z          NO CHANGE IN I                          ***GMF**
        INE     GMF,1,1                                                 ***GMF**
$       ALTER   5841                                                    ***GMF**
        IFE     GMF,1,5                                                 ***GMF**
TM169Z  LDQ     TRCFLG          TPAP ID IN QL                           ***GMF**
        STZ     TRCFLG          RESET FLAG INDICATOR                    ***GMF**
        TRMAC   43              TRACE TPAP MODIFIED                     ***GMF**
        TRA     TM1475                                                  ***GMF**
TRCFLG  DEC     0               FLAGS SHOW WHAT IN TPAP IS MODIFIED     ***GMF**
        INE     GMF,1,1                                                 ***GMF**
$       GMAP                                                            TRXD
$       ALTER   76
    (NTPE ALTER cards)
***SITE***SITE***SITE***SITE***SITE***SITE***SITE***SITE***SITE***SITE*****GMF**
*                                                                       ****GMF**
*         THIS SYMBOL CONTROLS ASSEMBLY OF THE TPE INTERFACE WITH       ****GMF**
*         THE CCTC/WAD GENERALIZED MONITORING FACILITY (GMF)            ****GMF**
*                                                                       ****GMF**
*         SET THIS SYMBOL TO THE SAME VALUE IN TRXB,C, AND TRXD         ****GMF**
GMF     SET     1                                                       ***GMF**
*                                                                       ****GMF**
***SITE***SITE***SITE***SITE***SITE***SITE***SITE***SITE***SITE***SITE*****GMF**
        IFE     GMF,1,1                                                 ***GMF**
        SYMREF  TREP                                                    ***GMF**
TRMAC   MACRO                                                           ***GMF**
        IFE     GMF,1,2                                                 ***GMF**
        XED     TREP            MAKE THE TRACE                          ***GMF**
        VFD     18/#1,18/60     18/SUBTRACE,18/MAJOR TRACE  =074        ***GMF**
        ENDM    TRMAC                                                   ***GMF**
$       ALTER   108                                                     ***GMF**
        TRMAC   1               TRACE TERMINAL CONNECT SCAN TIME        ***GMF**
$       ALTER   432                                                     ***GMF**
        IFE     GMF,1,2                                                 ***GMF**
        LDQ     TERM,$                                                  ***GMF**
        TRMAC   2               TRACE # TERMINALS CONNECTED, # ALLOWED  ***GMF**
$       ALTER   941                                                     ***GMF**
        TRMAC   4               TIME OF DISCONNECT                      ***GMF**
$       ALTER   1965                                                    ***GMF**
        TRMAC   3               TRACE TIME BREAK REQUEST NOTICED        ***GMF**
$       ALTER   2111                                                    ***GMF**
        IFE     GMF,1,2                                                 ***GMF**
        LLR     36                                                      ***GMF**
        TRMAC   6               TRACE REASON FOR IMPASSE                ***GMF**
$       ALTER   2350                                                    ***GMF**
        IFE     GMF,1,3                                                 ***GMF**
        TSS     *+1             LEAVE MM                                ***GMF**
        TRMAC   7               TRACE TPAP SPAWNED, TERMINAL ID         ***GMF**
        MME     .EMM            RETURN TO MM                            ***GMF**
$       ALTER   2400,2407
```

Figure 12-21. (Part 9 of 11)

```
(NTPE ALTER cards)
      LDA     .SPWN+1,1        TRACE # TPAPS TERMED TO SPAWN ANOTHER    ***GMF**
      TRMAC   8                                                         ***GMF**
      TRMAC   21                                                        ***GMF**
$     ALTER   2781                                                      ***GMF**
      IFE     GMF,1,1                                                   ***GMF**
      SYMDEF  TP0868                                                    ***GMF**
$     ALTER   2843                                                      ***GMF**
      IFE     GMF,1,2                                                   ***GMF**
      LDQ     TP0868                                                    ***GMF**
      TRMAC   9                TRACE # WORDS BUFFER SPACE ACQUIRED      ***GMF**
$     ALTER   2847                                                      ***GMF**
      IFE     GMF,1,2                                                   ***GMF**
      LDQ     TP0868                                                    ***GMF**
      TRMAC   10               TRACE # WORDS BUFFER SPACE DENIED        ***GMF**
$     ALTER   3083                                                      ***GMF**
      TRMAC   11               TRACE WHO ISSUES WHAT INVALID KEYWORD    ***GMF**
$     ALTER   3088                                                      ***GMF**
      TRMAC   12               TRACE TPAP UNAVAILABLE TO USER           ***GMF**
$     ALTER   4025                                                      ***GMF**
*     IFE     GMF,1,2                                                   ***GMF**
*     LDQ     TP1538                                                    ***GMF**
*     TRMAC   9                                                         ***GMF**
$     ALTER   4093                                                      ***GMF**
      IFE     GMF,1,1                                                   ***GMF**
      SYMDEF  TP2055                                                    ***GMF**
$     ALTER   4110                                                      ***GMF**
      IFE     GMF,1,2                                                   ***GMF**
      LLR     36                                                        ***GMF**
      TRMAC   13               TRACE # BLOCKS BUFFER SPACE RELEASED     ***GMF**
$     ALTER   4985                                                      ***GMF**
      TRMAC   5                TRACE NO OF UNDOT ENTRIES                ***GMF**
$     ALTER   5565                                                      ***GMF**
      TRMAC   14               TRACE OUTPUT TRANSMITTED FOR WHOM        ***GMF**
$     ALTER   5688                                                      ***GMF**
      TRMAC   5                TRACE # AVAILABLE UNDOT ENTRIES          ***GMF**
$     ALTER   6202                                                      ***GMF**
      TRMAC   5                TRACE # AVAILABLE UNDOT ENTRIES          ***GMF**
$     ALTER   6584                                                      ***GMF**
      TRMAC   16               TRACE # UNDOT ENTRIES RETRIEVED FROM JOURNAL FILE
$     ALTER   6850                                                      ***GMF**
      TRMAC   17               TRACE SNUMB FOR WHICH INTERCOM Z---Z REISSUED0000
$     ALTER   7550                                                      ***GMF**
      TRMAC   21               TRACE TPAP TERMED                        ***GMF**
$     ALTER   8131                                                      ***GMF**
      IFE     GMF,1,1                                                   ***GMF**
      SYMDEF  TP3015,TP3016,TP3023                                      ***GMF**
$     ALTER   8176                                                      ***GMF**
      TRMAC   18               TRACE SNUMB RECEIVING CONTROL OF LINE    ***GMF**
```

(Figure 12-21. Part 10 of 11)

```
$       ALTER   8258                                                    ***GMF**
        TRMAC   19              TRACE SNUMB ISSUING WRITE               ***GMF**
$       ALTER   8300                                                    ***GMF**
        IFE     GMF,1,1                                                 ***GMF**
        SYMDEF  TP3115                                                  ***GMF**
$       ALTER   8338                                                    ***GMF**
        TRMAC   20              TRACE SNUMB ISSUING READ                ***GMF**
$       ALTER   8624                                                    ***GMF**
        TRMAC   5               TRACE # AVAILABLE UNDOT ENTRIES         ***GMF**
$       ALTER   9278                                                    ***GMF**
        TRMAC   21              TRACE TPAP ABORTED                      ***GMF**
$       ALTER   10707
    (NTPE ALTER cards)
        TRMAC   22              TRACE TPAP BOMBED, WHY                  ***GMF**
    (NTPE ALTER cards)
```

Figure 12-21. (Part 11 of 11)

SECTION 13. EXTENDED USE OF GMC

13.1 Introduction

Throughout this document the Generalized Monitor Collector (GMC) has been described as a trace based performance evaluation tool used to evaluate the performance of certain preselected components of the H6000 computer system. The collectors are all prewritten, as are the data reduction programs and report formats. The user of the GMC has had no clearly defined way of choosing what data is being collected or how to expand the monitoring to other portions of system activity.

In this section, an attempt will be made to describe the procedure that must be followed if a user wants to develop his own monitor. The design of GMC provides the user with a relatively simple method of creating additional monitors. As used in this section, the word monitors does not necessarily mean a performance monitor but rather it is being used in the broader sense of a tool used to check on the operation of a given program or machine component.

13.2 Monitor Structure

Every monitor within the GMC is comprised of an initialization section and a data collection subroutine. The initialization section is used to insure that the correct traces are active, all needed subroutines are present and, finally, to absolutize necessary code in the data collection subroutine. The data collection subroutine is used to process the occurrence of a given H6000 trace, gather whatever data is desired and request the GMC Executive to buffer the collected data to tape, for later processing.

13.2.1 Initialization Section. Prior to coding this section, the user must decide what traces he desires to capture and must select a monitor number. The monitor numbers can range from the letter A through letter F (Monitor numbers 10-15). The H6000 traces are described in the H6000 System Tables Manual (#DD14A). For purposes of this discussion, it is assumed that monitor #C has been selected and it is desired to capture trace OCTAL 27. Figure 13-1 shows the skeleton format required for the initialization section. By following this format, and looking at the initialization sections of the other GMC Monitors, the user should have no problem in writing this portion of the code.

13.2.2 Monitor Subroutine. The monitor subroutine is a master mode program that collects all required data upon execution of the desired trace, and requests the GMC executive to buffer that data to tape. Figure 13-2 provides a skeleton for a monitor subroutine.

In the section where the data is captured, if the user performs the following two statements, every record collected will be time stamped. The data reduction program can then easily convert this time stamp into a wall clock time. The two instructions are:

```
        SYMREF      (Any variables or locations in Monitor routine that
                    need to be absolutized because of register
                    modification)

        LDA     ONOFF+#,$       # = Monitor number (12)
        TZE     USEROF,$        Monitor is off
        LDA     =#,DL           # = decimal equivalent of trace number
                                (23) (OCTAL 27)
        SZN     CVECTR+#,$      # = decimal equivalent of trace number
                                (23) (OCTAL 27)
        TZE     SUBERR,$        Monitor routine missing from R*
        LDAQ    .CRTCT
        ANAQ    USERTC,$        test for required traces being on
        TNZ     TRCERR,$        required trace was off
        AOS     ONOFF+#,$       # = monitor number (12)
        ASX5    XXX,$           absolutize code in the monitor
                                subroutine, if required
            .
            .
            .
            .
            .
        SZN     .CRMBA          Is extended memory on?
        TNZ     USEREN,$        Transfer means yes
        LDA     EXTOFF,$
        ANSA    XXX,$           modify any necessary instructions
                .               because extended memory is off
                .
                .
                .
        TRA     USEREN,$
USEROF  STZ     CVECTR+#,$  # = decimal equivalent of trace number (23)
        (repeat for each trace to be captured)
        TRA     USEREN,$
USERTC  OCT     (12 Octal digits)   set appropriate bit
        OCT     (12 Octal digits)   to match the trace #'s
                                    being collected
USEREN  NULL
```

Figure 13-1. New Initialization Section

```
        LBL     T#C,T#C             # = OCTAL equivalent of trace
                                    number (27)
        LODM    .G3MAC
        EIGHT
        .ENTRY  XSA2,,,T#C,PRG      # = OCTAL equivalent of trace
                                    number (27)
        SYMDEF  T#,T#C,_ _ _ _ _    # = OCTAL equivalent of trace
                                    number (27)
                                    In addition, any other locations
                                    referenced by the initialization code
        SYMREF  BUFCTL,SVRQ,SVRA    SVRQ contains the Q register from
                                    the trace.  SVRA contains the A
                                    register from the trace.
        INHIB   ON
        TRA     WRAP,$              wrapup code for when monitor is
                                    terminated
T#      NULL                        # = OCTAL equivalent of trace
T#C     NULL                        number (27)
        (Code to capture any data desired.  Data is written to a user
        defined buffer area.  Code is in Master Mode).
        LDX3    =#,DU               number of words of data generated
                                    (including time stamp if used)
        STX3    TOPBUF,$
        EAX3    1,3                 count the header word
        EAX2    TOPBUF,5
        TSX1    BUFCTL,5            Tell executive to buffer data
        TRA     GOOD,$              data written to buffer
        TRA     LOST,$              internal buffers were full, data is
                                    lost

GOOD    data has been buffered.
        Perform any needed housekeeping before returning to system
        TRA     0,0

WRAP    monitor is terminating.  Perform any required housekeeping and
        write a final record, if desired
        TRA     0,0

TOPBUFEZERO  0,#                    # = decimal equivalent of trace
        TIME    BSS     1           number (23)
        BUF     BSS     #           # = number of words of data

LOST    data was not able to be placed in buffer.
        Perform any needed housekeeping and
        return to system
        TRA 0,0
```

Figure 13-2. Monitor Subroutine

```
RSCR    32
STQ     TIME, $    store the time.  This word should be counted in
                   the total number of words of data collected.
```

13.3 User Generated Trace

If a user has the capability of modifying program code, either directly in source, or through patching, and can execute in master mode, he can create his own system traces anywhere within a program. Thus, in effect, he can use the GMC system to monitor the execution of any program of his choice. If the new monitor is being executed in conjunction with the CAM, CPU, GRTS, or MSM monitors, the user must be sure not to use the trace numbers 14, 70, 62, 73 or 76 octal, created by these GMC monitors. Figure 13-3 provides the skeleton code for creating system traces. It should be noted that the time returned from executing the trace code in figure 13-3 is not the same time as obtained from the RSCR instruction described in 13.2.2.

It is also possible to directly transfer into the GMC program. However, this procedure is too complicated to be explained in this document. If the user is interested in using this procedure, he should examine the manner in which the CPU monitor interfaces with the system. Special attention should be paid to files CPU.PAT, CPU.REMO, and CPUDOIT. When using this method the user will need to alter his monitor routine with the addition of three instructions. These instructions check whether the user generated trace, or a system trace is being processed. System traces will be ignored. In figure 13-2, following label T#C the following three instructions should be added:

```
LXL2    SVRA,5      get the trace number
CMPX2   =05757XX    where XX = octal trace number
TNZ     0,0         system trace ignore
```

13.4 Data Reduction

In order to analyze the tape generated by the new monitor, a data reduction program must be coded. This step will prove to be the most difficult and time consuming step to execute. In order to proceed with the coding of a data reduction program the user should examine the source code of the already existing GMC monitors. Much of the code needed to process the GMC tape is already written and can be used with minimum modification by the user. For purposes of this document, the CM data reduction program will be used as an example. The remaining description assumes that the user has a source listing of CM available to him.

13.4.1 Subroutine INITZ. This subroutine performs many functions within the program but the areas of interest at present begin at the first lines of executable code. At this point in the code the initial data record on the tape is read, the initial start time of data collection is determined and the system configuration is determined. This initial record was written to tape by the GMC executive and must be processed before the user can begin to reduce the records generated by his monitor. Much of

```
        XED     .CRTRC
        TRA     DONE,$              code to be executed when trace
                                    is completed or if trace is off

        Code to set up A&Q
             registers

        ORA     =05757XX,DL         where XX is the octal trace number
                                    desired to be used.  The 5757 insures
                                    that the GMC will be able to distinguish
                                    this user generated trace from a system
                                    trace of the same number.
        TRA     -1,0 (Time of Day in Q register)
            or
        TRA     1,0 (A & Q registers not changed)
DONE    NULL        return code or trace off code
```

Figure 13-3. Creation of User Trace

the data in this record can be ignored by the user, if he is not interested in the system configuration. However, execution of this code exactly as written will insure proper execution. The meaning of each word within this initalization record can be found in section 5 of this document. All code down to the beginning of the 1000 do loop should be executed. The key points that should be noted are as follows:

- o DATPTR - This word will now point to the header word of the first monitor trace on the tape.
- o INDATA - Array which contains a complete GMC record, i.e., 4090 words.
- o HOUR - Starting hour of the data collection
- o MINUTE - Starting minute of data collection
- o PMHOUR - Current hour
- o PMINUT - Current minute
- o PMSECD - Current second
- o PTIME - Current RSCR time stamp
- o ORGTIM and STRTIM - Initial RSCR time stamp
- o OLDT,BEGINT - Variables needed for later time calculations

upon leaving subroutine INITZ the user can begin to process trace data. The variable DATPTR points to the header record of a monitor trace with bits 30-35 of this word indicating the trace number of the data following.

13.4.2 Subroutine NXTRECRD. After having processed a monitor trace or determining that this monitor trace is of no interest, a call to subroutine NXTRECRD will advance DATPTR to the next monitor record.

13.4.3 Subroutine RLOVER. If the user is interested in maintaining a wall clock time for each monitor trace executed, the following code should be executed:

```
PTIME = AND (INDATA(DATPTR+X),MAKPLS)  Where X is the offset for
        the RSCR time word in the trace.

STARTM = ORGTIM

SBEG = BEGINT

CALL RLOVER

TIMES = TIMES/100000 + SBEG
```

At this point, TIMES will contain the total run time accumulated thus far in 1/10 seconds. MAKPLS is a variable that is defined within the CM program.

13.4.4 Subroutine TIMEDAY. In order to convert the variable TIMES into a wall clock time, the following code should be executed:

```
CALL TIMEDY2
PMHOUR = PHOUR
PMINUT = FMINUT
PMSECD = FSECOND
```

PMHOUR, PMINUT and PMSECD will now contain the current wall clock time.

13.4.5 Subroutine MAIN. The GMC executive will create several additional trace records that must be checked for by the user and must be processed correctly. The first is a lost data record which indicates that one or more buffer records were lost and not written to tape. Another record is a monitor termination record and indicates that data collection has been completed. All these record types are processed in the MAIN subroutine and the user should reference this code for proper handling of these record types.

13.4.6 Other Routines. The CM data reduction program also contains routines for creating histograms and plots. However, for purposes of this document, these routines are too complicated to explain. This chapter should allow a user to be able to write his own monitor and begin the process of writing a data reduction program. The user should note that two subroutines required for all data reduction programs are ALIGN and EVENUP.

13.5 Utility Tape Dump

If the user is interested in quickly being able to see the data that was collected he can run a standard Utility Tape Dump. The following JCL should be used:

```
$    UTILITY
$    LIMITS        10,40K,,30K
$    FFILE         01,NSTDLB,NOSRLS,BUFSIZ/4094,FIXLNG/4094
$    FUTIL         01,,RWD/01/,DUMP/XR/
                        where X = number of records to be dumped
$    TAPE          01,X1D,,Tape#
$    ENDJOB
```

SECTION 14. CONDUCTING A SITE COMPUTER PERFORMANCE EVALUATION USING THE GMC

14.1 Introduction

A general plan for using the GMC System and document when conducting a performance evaluation of a Honeywell 6000 computer system is presented in this section. Detailed analysis procedures that guide the analyst in applying specific techniques to analyze system performance are introduced. The primary purpose is to aid the analyst in his use of the GMC reports. This section will indicate which reports the analyst should reference, what data he should extract from those reports, and some guidance as to what the reports are indicating. In some cases, possible corrective ("tuning") recommendations will be made, but tuning guidance is not the primary purpose behind this section. If a tuning study should be made, the user must rely on his own personnel skills and the tools of GMF to perform such a study.

14.2 General Definitions

14.2.1 Computer Performance Evaluation (CPE). Computer performance evaluation is a generic term applied to many techniques for determining the performance characteristics of both a computer system and its associated site processing activities. The performance characterisitics may be compared to many criteria, including: (1) standards of economic operation, (2) technical norms, or (3) measures of service provided.

14.2.2 Computer System Performance Variables. The performance of a computer system is influenced by nearly every facet of the data processing function. The following examples define the scope of the computer performance function.

14.2.2.1 System Design. Computer application system design and development can be the starting point of performance degradation. Errors in original design with respect to I/O media selection, file structures, frequency of run, etc., may result in less than optimal performance for as long as an application is in existance.

14.2.2.2 Programming. A computer programmer's proficiency and the availability of program optimization tools, for example, will influence program design and coding, and affect system performance.

14.2.2.3 Hardware Configuration. Specific components of a computer system may be mismatched to the system as a whole, causing major subsystems (or the entire configuration) to operate at a reduced performance level. Even if the performance capabilities of the individual subsystems are reasonably well-matched, the system may be poorly configured for the site's workload, resulting in poor performance.

14.2.2.4 System Software. System software is considered to be those programs supplied by the mainframe vendor. This software may be inappropriately parameterized to fit the site workload, or may be a source of high overhead or bottlenecks to efficient workload processing.

14.2.2.5 Operations. Operations consists of functions such as the scheduling of the workload, providing job assembly (and library) services, and operating the system through the console. All of these functions are vital to the proper operation of the system. Mistakes, insufficient training, poor documentation, and a variety of other factors may contribute to operational problems which substantially decrease system performance.

14.2.2.6 Communications Hardware and Software. A communications network, its interface to a central system, and the software used to control the on-line applications, may have a significant impact on the system's overall performance.

14.2.2.7 Computer System Performance Tuning. The process of analyzing and appropriately adjusting computer system performance variables is known as computer system performance tuning. After becoming proficient in the use of GMF, the user should be able to conduct his own tuning study.

14.2.2.8 Turnaround Time. This is the total elapsed time taken by a job (or set of jobs) submitted to a WWMCCS site for batch processing. Total batch turnaround time is comprised of computer system processing and physical input and output handling in the machine room, both before and after system processing. A job's total turnaround time therefore includes all processing and waiting points through which the job must pass from submission until return to a user. The GMC measures only a batch job's computer processing turnaround time. The other measure of total turnaround time (input/output handling) are important and can be a bottleneck to total turnaround time.

14.2.2.9 CPE User Objectives. The definitions associated with any or all of the above data processing functions can be considered variables that affect computer system performance. The GMC reports should be used to determine which of the above functions need further investigation.

14.3 Solutions to Performance Problems

Particular resource bottlenecks may be confirmed as elongating turnaround or response time. Several solutions can usually be applied to remove a particular bottleneck. In general, four kinds of solutions exist to remove identified bottlenecks.

14.3.1 Scheduling Solutions. These solutions change the way that either batch or TSS workloads are scheduled for processing. They shift particular workloads to more evenly distribute system resources across the workload.

14.3.2 Parameter Solutions. These changes involve adjustments to system or subsystem functions. Examples include: (1) changes to the

parameters of the GCOS Dispatcher or (2) a change in the placement of GCOS libraries. A solution may also include specific changes to GCOS code, made through authorized software patch procedures.

14.3.3 Programming Solutions. These changes can involve modification of one or more processing jobs running in the system. Recommendations may be made to speed application of jobs by changing particular file locations discovered as delaying the job.

14.3.4 Sizing Solutions. These types of system changes involve an increase, decrease, or realignment in the system's hardware configuration.

14.4 Structure of the Analysis Process

The analysis process in figure 14-1 is a flow chart that should be referenced during a performance study. The analysis process is comprised of two phases: (1) a Problem Definition Phase and (2) a Problem Analysis Phase. The activities of the Problem Definition Phase are directed toward determining whether a batch turnaround time or TSS response time problem actually exists. The activities of the Problem Analysis Phase are directed toward revealing causes of the identified turnaround time or response time problem.

14.4.1 Starting The Process. The evaluation process may begin in one of three ways: (1) by direct request, (2) by an internal review of site-selected service elongation metrics, (3) by user requests.

14.4.1.1 Direct Request. The request may by initiated by WWMCCS management, directing a facility to perform an analysis and tuning effort. Many reasons could cause management to request a CPE study. Examples include: (1) desire to extend the life of a system, (2) pressure applied from users through higher authority, (3) awareness of the potential for performance gain after an evaluation, and (4) a desire for management to evaluate cost reduction programs.

14.4.1.2 Internal Review. The decision to initiate a study may result from the output of a performance exception reporting system or from the desire to conduct a periodic review of site operations.

14.4.1.3 User Requests. Users of site services may request an evaluation. Complaints of unacceptable batch turnaround time or TSS response time can initiate a search for their causes. Unfavorable comparisons with service rates provided at other installations can point out the need for a study to determine if service can be improved.

14.4.2 Problem Definition Phase. The Problem Definition Phase (see figure 14-1, part 1) is comprised of four activities: (1) define and verify the problem, (2) gain understanding of facility environment, (3) understand installation service objectives and priorities, and (4) specify current evaluation objectives. Some of these activities are initiated as evaluation begins. Others are maintained as on-going activities. The activities of the Problem Definition Phase are introduced below.

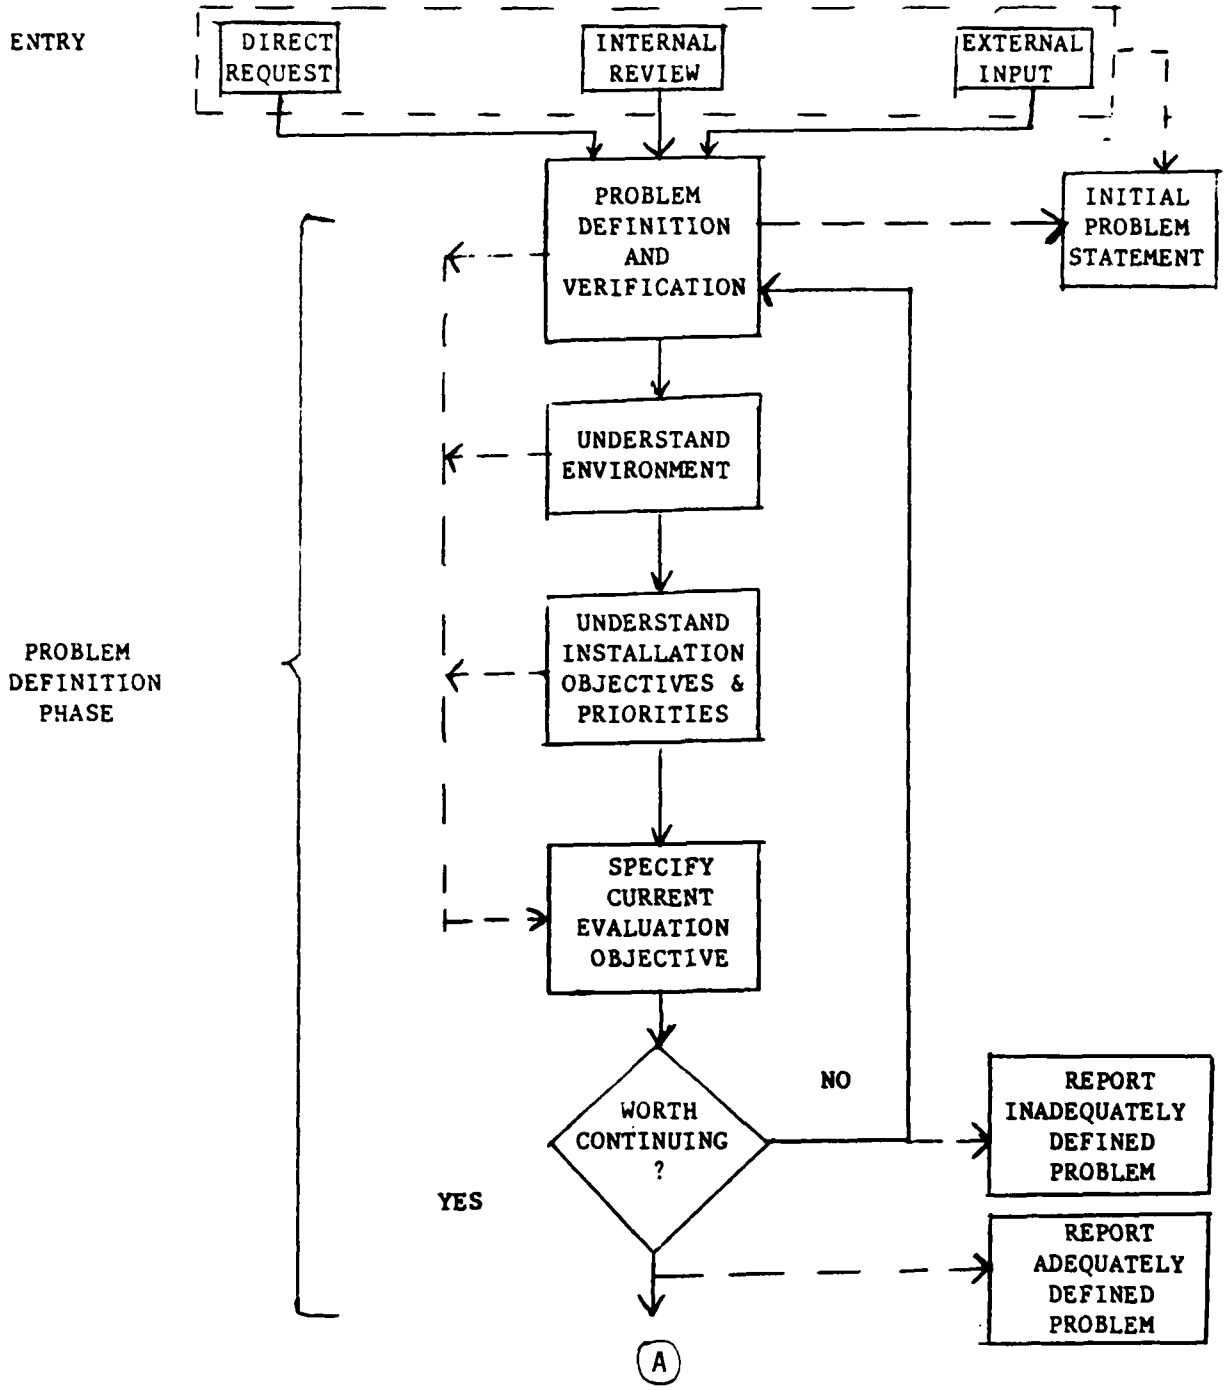


Figure 14-1. Flowchart Of Evaluation/Tuning Process (Part 1 of 2)

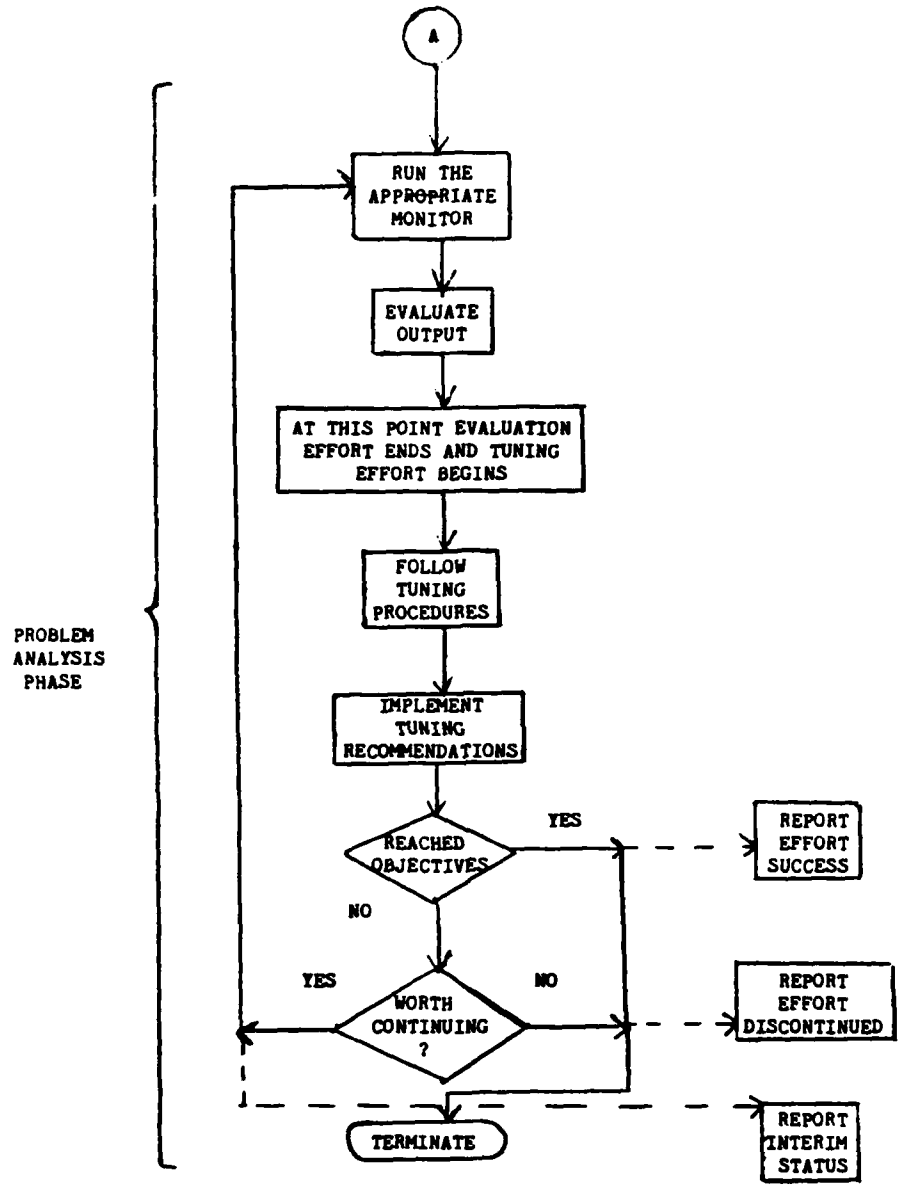


Figure 14-1. (Part 2 of 2)

14.4.2.1 Define and Verify the Problem. For whatever reason the evaluation was initiated, a starting point or premise must be established. The problem must be described as well as it can at this point, even if at a very general level. "Batch turnaround time is poor" or "TSS response time must be improved" are valid statements of problems at this time. Site management must then verify that this is the problem they wish to pursue. The facility staff and users should verify that, in fact, this problem does exist and that they view it as a problem of importance. Note that a specific evaluation objective is not yet defined; only the basic problem statement is verified.

14.4.2.2 Gain Understanding of Facility Environment. This second step helps the CPE analyst relate the system environment to the problem statement. This activity may be time-consuming if the analyst has little personal experience with the facility and the system. Information collected during this activity will help determine the reason for the problem's importance; it may explain a reason for the problem's existence and help rank one problem against another. The activity forces an analyst to view the entire facility; this is important since many performance problems are caused by combinations of factors. The activity assists the analyst in understanding how to narrow, refine, modify, and improve the problem's definition in order to attempt a valid solution. Each area described in the following paragraphs should be examined. It may help the analyst to examine them in the order they are given below.

14.4.2.2.1 Hardware Configuration. The exact system configuration should be determined and a diagram of the overall structure of the system should be created. Collection of reliability statistics on major configuration components and a history of hardware changes to the configuration should be done.

14.4.2.2.2 Software Configuration and Development Practices. The exact operating system configuration should be determined. Any local changes made to the system and any specialized major subsystems that are running should be documented. Document software monitors and other CPE measurement techniques used. Determine program optimization techniques used and note standards imposed on operations and programming staffs.

14.4.2.2.3 Existing CPE Practices. Responsibility for computer performance evaluation at the site must be determined. Identify sources of CPE data that are employed and determine how the data are used to evaluate system performance. Determine if personnel in all site functional areas (e.g., programming and operations) can relate to the performance data. Research how previous CPE studies conducted at the site in the past have been documented.

14.4.2.2.4 Site Workload Characteristics. The existing workload uses of system resources should be scribed. Identify patterns of resource use by selected jobs. Obvious bottlenecks to the handling of jobs within the installation should be determined. The resource monitor, the RMC portion of GMF, should prove very useful during this phase.

14.4.2.2.5 System Users. The workload analysis should be expanded by investigating the practices of the major users in the installation. Determine how special priorities are assigned to these users and whether the users can directly control system resources. Note chargeback schemes used (if any) to level the workload across the operating day and night. Evaluate the levels of user satisfaction (or dissatisfaction) exhibited.

14.4.2.2.6 Operations Practices. The operating shift schedule for the site should be examined. Analyze pre-scheduled and nonproduction time and unscheduled nonproduction time. Evaluate training provided for operators and systems programmers. Study site library maintenance procedures. Examine the production control points established by operations. Observe how formalized logs are maintained to track work as it passes through the installation.

14.4.2.2.7 Batch Job Scheduling. Observe how batch work is accepted for processing at the site. Study the techniques used to schedule the batch workload that are either automated or are manually implemented.

14.4.2.2.8 Site and Computer Facility Organization. Study the organization structure and reporting authority at the site. This includes the organization of the site itself. Determine the extent to which system sizing activities are organizationally separated from operations and applications systems development.

14.4.2.3 Understand Installation Service Objectives and Priorities. This third step determines site processing objectives and priorities to be applied to decisions made during both the Problem Definition and Problem Analysis Phases.

14.4.2.3.1 Installation Service Objectives. To be effective, a CPE effort must consider the hardware, software, personnel, and service objectives of the computer installation being analyzed. Although the first three areas are readily identifiable, the site's service objectives are often misunderstood. An installation's objectives can emphasize production or availability, or a mixture of both. A mixture is generally dominated by either a production or availability objective. A production objective attempts to take full advantage of the capabilities of the system in terms of throughput. With a customer-oriented availability objective, however, the interests of the users take priority over the most efficient use of system resources. In either case, the computer installation's management objective is return on investment. For the production objective, the most important investment

is the computer; for the availability objective, the most important investment is the people or system using the computer. The difference between these two objectives is reflected in how the performance parameters are evaluated.

14.4.2.3.1.1 Production Objective. Traditional data processing is production-oriented. The computer installation is viewed as a large investment for capacity to do routine work. Much of that capacity depends on a high degree of simultaneous use of many system components. Management's production objective is to use as many of the components as much of the time as possible by scheduling compatible jobs. Under these conditions, low batch turnaround and TSS response times are secondary to high machine utilization. Management's success is generally defined in terms of a relatively high number of jobs contending for computer resources (a high multiprogramming level), high central processor unit (CPU) activity, nearly full memory utilization, and highly active channels. Schedule-driven processing uses internal and external priority allocations to sustain efficiency. Programs must be written to share available resources. This can be done with program segmentation techniques, use of a minimum number of devices, the indirect use of input/output through spooling, and adherence to rigid standards -- all at the expense of the individual job. Data must be so distributed that device activity is economically justified with the attainment of a low system wait, low unit cost, and high simultaneity.

14.4.2.3.1.2 Availability Objective. Service-oriented systems are more concerned with return on investment for users (managers, scientists, or programmers) than with the computer itself. Low TSS response and batch turnaround times are critical. Their increase delays user operations on such applications as on-line command and control systems, real-time update systems, scientific support services, and program development systems. Demand-driven processing varies in activity levels, both by users and type of work. Minimum emphasis is placed on scheduing. The system must therefore have available capacity ready to respond to demand. Such a system must have utilization rates well below the 100% limit in order to accommodate the variations in demand. Success is measured in terms of user satisfaction and little emphasis is placed on reporting high utilization figures.

14.4.2.3.1.3 Mixed Objectives. Usually computer installations encounter (1) demands for highly responsive services and (2) pressure from management for high production rates. The two objectives are not mutually exclusive. Predominance of either objective can be identified within operational time periods, and management must evaluate whether the satisfaction of one objective might deleteriously affect the satisfaction of another.

Computer performance evaluation can serve the management of either type of computer installation and can also serve mixtures that might have different objectives, depending on the time of day or day of the week. However, it is important that managers, auditors, and executives recognize the implications of the different objectives when they compare the performance of one installation to another.

14.4.2.3.2 Installation Priorities. The installation's priorities derive from the predominance of either the production objective or the availability objective at a particular site. They also, by implication, determine the sequence in which system tuning solutions are applied.

14.4.2.3.2.1 Service Priorities. A site may feel that low TSS response time is more desirable than low batch turnaround time for a certain period of the operating day. The analysis procedures presented in this section are not generally directed toward determining which of the two is "more important". However, the sequence of examining either of the two service areas may be affected by this priority.

14.4.2.3.2.2 Evaluation/Tuning Solution Priorities. Tuning solutions are presented to correct system bottleneck conditions. In nearly all cases, more than one solution is specified to correct a problem. The solutions are generally proposed in a sequence that recommends the more quickly (or easily) applied solution be implemented before others. Installation requirements may change this sequence.

14.4.2.4 Specify Current Evaluation. This fourth step is used to determine whether the originally specified problem can actually be investigated with currently-available procedures.

14.4.2.4.1 Objective. An objective is a stated (i.e., documented) goal of an analysis. An objective must be stated in specific and quantified terms. The objective statement is used to determine when a particular analysis has been completed. Examples of well-stated objectives are: (1) "reduce mean response time for (stated) non-trivial TSS commands to five seconds" and (2) "reduce the mean batch turnaround time for (stated) jobs to 1.5 hours." A well-stated objective includes: (1) definition of the workload catagory, (2) a description of the process to be investigated, and (3) a service metric value. Examples of badly-stated objectives are "reduce TSS response time" and "improve turnaround time." Note the missing components of these objectives.

14.4.2.4.2 Objective Decision. Determine whether the objective as documented is realistic, attainable, and a cost effective target.

14.4.2.4.2.1 Attainable Objectives. Objectives must be attainable. It might be possible to reduce a program's elapsed time to certain limits, but not to a desired limit, simply because of the quantity of I/O and computation that occurs within the program. If this is the case, re-evaluate the objective.

14.4.2.4.2.2 Realistic Objectives. It might be possible to reduce a program's elapsed time to the desired amount. However, particular site constraints may prevent its being a realistic goal. These constraints

might have their source outside of the tuning effort, but directly affect the internal performance of the program. These might include, for example, the requirement to give certain other jobs higher priority. If this is the case, re-evaluate the objective.

14.4.2.4.2.3 Cost-Effective Objectives. The additional effort, cost, and time required to achieve an attainable and realistic objective might not be worth it. A particular increment of performance improvement might simply not be cost effective. Determine the amount of improvement likely with additional effort. Decide whether the effort might be better off abandoned.

14.4.2.4.3 Determine if Worth Continuing. The final part of step four is to determine whether the process itself is worth continuing. There may be potential performance problems of greater magnitude or immediate importance that may have been uncovered during the Problem Definition Phase. Document the decision and the objective and obtain management approval of them.

14.4.2.4.4 Begin Problem Analysis. Having obtained concurrence from management, begin the second half of the performance evaluation process (see figure 14-1, part 2): the Problem Analysis Phase. Activities of this phase involve the execution of the various GMC monitors and the analysis of their output.

14.4.3 Problem Analysis Phase. The Problem Analysis Phase (see figure 14-1, part 2) is comprised of five activities briefly described below:

14.4.3.1 Run Appropriate Analysis Tool. Analysis requires collection of data to start an investigation. The various GMC monitors have been developed for this purpose.

14.4.3.2 Evaluate System Output. The output from the various GMC monitors must be analyzed.

14.4.3.3 Follow Tuning Procedures. The user must develop his own skill in developing provide specific procedures to test hypotheses by examining the reports produced by the GMC monitors. Specific system tuning steps can result from the tests.

14.4.3.4 Evaluate Need to Continue the Tuning. This decision is made after the relevant recommendations have been implemented. If the current objective (specified in the last activity of the Problem Definition Phase) has not been met, the user analysis can define areas for further investigation.

14.5 Composition of a Performance Evaluation Team

While this section has been developed to aid WWMCCS ADP Managers in the application of various computer performance evaluation tools and techniques, the actual team conducting this analysis must have a reasonable education and experience level on the H6000. The team should be comprised of at least one individual familiar with system software and its operation; one individual generally familiar with system hardware (at least in terms of functionality); one individual familiar with the procedures for executing jobs on the H6000 and at least one individual (probably a manager) who is familiar with the objectives and priorities of the installation. It certainly is possible that one individual may be able to handle several of the above functions.

14.6 System Evaluation

14.6.1 Introduction. A total system evaluation should be performed at least once a year. It may have to be performed more often, depending upon the results of daily resource statistics. A steady increase in resource utilization would possible imply the need for an evaluation. The Resource Monitor (RMON) of GMF should be used to track daily resource utilization.

The data collected for the detailed system evaluation must be representative (i.e., typical) of the data that would be collected on a "normal" day. This requires the collection of data over several time periods. The suggested schedule for collecting data for a total system analysis is to run the GMC monitors intermittently for two weeks. The monitors should be run as two separate groups, on alternate days. Group one would consist of the Memory Utilization Monitor (MUM) (possibly the Idle Monitor (IDLEM)), CPU Monitor (CPUM), Communications Analysis Monitor (CAM), and the GRTS Monitor (GRTM). Group two would consist of the Mass Storage Monitor (MSM), Channel Monitor (CM), and possibly the Idle Monitor (IDLEM). This sequence provides a good representative sample of the system workoad. To limit the amount of data collected, it is advisable to run Group Two only during prime time, heavy usage (4-6 hours per day). It may also be desirable to run the combined set of monitors at least once per week during the two-week monitoring session. This allows the analyst to get a complete, unified picture of the entire performance of the system. It must be realized that a large number of tapes can be generated under such a monitor configuration. Therefore, data collection should be limited to no more than four hours (during heavy prime time usage).

14.6.2 Selecting a Representative Value From GMC Histogram Reports. Some test procedures require an analyst to pick a "Representative Value" to describe a frequency distribution. The following paragraph gives guidelines for choosing Representative Values based on the type of distribution observed.

Figure 14-2 shows a hypothetical "Total Elapsed Time an Activity Was in Memory" report. Variations of this report will be used to illustrate types of distributions. This chapter will reference only the pictorial part of the report. All other parts of the histogram report are described in Chapter 6 of this document.

14.6.2.1 Symmetric Distribution Closely Clustered Around a Single Point. The "Representative Value" for the distribution shown in Figure 14-2 is easy to select. The values are clustered around the "Average": 17.6 tenths of a

DISTRIBUTION COLLECTED ON SYSTEM NMCC at 14:41:24 on 81-05-15

The Total Elapsed Time An Activity Was In Memory

Percent of Occurance Report 21

Tenths Second	00 /	05 /	10 /	15 /	20 /	25 /	30 /	35 /	40 /	45 /	50 /
10-10	I										
11-11	I										
12-12	IX										
13-13	IXX										
14-14	IXXXX										
15-15	IXXXXXXXXXX										
16-16	IXXXXXXXXXXXXXX										
17-17	IXXXXXXXXXXXXXXXXX										
18-18	IXXXXXXXXXXXXXXXXX										
19-19	IXXXXXXXXXXXXXX										
20-20	IXXXXXXXXX										
21-21	IXXXXX										
22-22	IXX										
23-23	IX										
24-24	I										

9276 Entries Total Average 17.62 Variance 4.85 Standard Deviation 2.20

Figure 14-2. Sample Symmetric Distribution

second, or 1.76 seconds. The shape of the distribution is symmetric -- about the same number of activities had values over 1.76 seconds as under 1.76 seconds, and the standard deviation is small when compared to the average. The absence of a second line under the "Entries Total" line indicates that no activities stayed in memory longer than 2.4 seconds. When a distribution resembles this example, use the "Average" printed at the bottom as the Representative Value.

14.6.2.2 Skewed Distribution. Figure 14-3 shows another distribution. This distribution is "skewed" (i.e., not symmetric), because most of the activities spent around 0.1 to 0.7 seconds in memory, while some spent as much as four or five seconds. Care should be taken when selecting a "Representative Value" from this distribution. If the analyst wants to emphasize the "typical" activity, which stayed in memory 0.3 seconds or less, he could select the "median" (the value which evenly divides the activities in the distribution -- half spent less time in memory, and half spent greater time in memory). In figure 14-3, the median is about 0.29 seconds. The median can be estimated from these reports by descending down the "Cumulative" column (not displayed in the figure) until the value first exceeds 50. The median falls within the time range of this row.

14.6.2.3 Distribution With Outliers. There are some instances when the distribution will also have some values that were too big to fit in the histogram. This condition will be indicated by an additional output line at the bottom of the report. This line will indicate the number of occurrences that were outliers, the average for just the outliers, and the average for the values that fit into the report, minus the outliers. The three important factors about these types of distributions are: (1) the amount of times they occur; (2) the percent of the total values that are outliers; (3) the "in-range average."

If the percent of outliers is greater than 10%, the analyst should use the overall average, given in the first line of the report, for any comparisons. If the percentage of out-of-range values is less than 10%, the analyst can use the "in-range average" value for his comparisons since the effect of the outliers will most likely be minimal on the total system performance.

14.6.3 Memory Evaluation. The first step in a memory evaluation is to summarize all the pertinent information. The easiest way to do this is for the user to create a table similar to figure 14-4. The information for each column is collected from various MUM reports. Once the chart is filled out, the analyst can then ascertain a reasonable idea of the overall memory status of the system. The reports required to collect the statistics will be discussed, as well as an analysis procedure. The current version of the MUM will automatically produce this table (see subsection 6.3.13 - Memory Statistics Report).

DISTRIBUTION COLLECTED ON SYSTEM NMCC AT 14:14:24 ON 81-05-15

The Total Elapsed Time An Activity Was In Memory

Percent of Occurence

```
Tenth      00   05   10   15   20   25   30   35   40   45   50
Second      /    /    /    /    /    /    /    /    /    /    /

  0-1      Ixxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx
  2-3      Ixxxxxxxxxxxxxxxxxxxxxxxxxxx
  4-5      Ixxxxxxxxxxxxxxxx
  6-7      Ixxxxxxxxxxx
  8-9      Ixxxxxx
 10-11     Ixxxx
 12-13     Ixxx
 14-15     Ixx
 16-17     Ix
 18-19     Ix
 20-21     I
   .
   .
   .
 56-57     I
```

118669 Entries Total Average 4.247 Variance 18.52 Standard Deviation 4.3

Figure 14-3. Sample Skewed Distribution

14.6.3.1 Obtaining the Data. The first step is for the user to examine the monitor data collection reports for the total time of the run. Any monitoring session of less than 2 hours in duration should be discarded. The Title Page will indicate the overhead generated by each of the GMF monitors active during the data collection phase. While this information is not used in this analysis, it is an item of information that is usually requested. Columns 1 and 2 of figure 14-4 can be obtained directly from the MUM Title Page. Under the System Configuration section of the Title Page, the amount of memory configured on the system will appear. This figure should be tentatively noted under column 5. Following this figure, the Title Page will list the amount of memory used by Hard Core, Core Allocator, SSA Cache and if any memory releases occurred during the monitoring session. All these functions have the effect of reducing the available user memory. These values should be summed and tentatively recorded under column 6.

Following the above information are several lines of statistics concerning the processing characteristics of the system. Columns 17 and 18 may be filled from this information. When determining the number of activities processed per hour, the analyst has two figures available. The analyst may choose to record the total number of activities processed per hour or he may record the number of (non-system schedular) activities processed per hour. During the course of the day many system schedular activities (activity 0 of any batch job) are processed. These activities are not really user generated, but rather are generated from the system. Therefore, by removing these activities from the total activities processed, a more realistic figure will be generated. The final three lines of the title page can be used to provide the data for columns 3 and 7 on the chart. Column 4 is filled from Report 1, columns 8 and 9 from Reports 11 and 12, columns 10 and 11 from Reports 16 and 17, column 12 from Report 19, column 13 from Report 51, column 14 from Report 37, column 15 from Report 50, and column 16 from Report 36. The two remaining columns that need to be completed are columns 5 and 6. The System Program Usage of Memory Report should be used to complete column 6. When processing the MUM data reduction program, the user should seriously consider using the MASTER input option. This will provide the user with a much better indication of the true system program load. In order to complete column 6, the user should record the total value appearing under the "WEIGHTED TTM" column of the System Program Usage of Memory Report. This value should then be added to the value already recorded under column 6. Finally column 5 can be determined by subtracting the value reported in Column 6 from the value previously reported in column 5.

14.6.3.2 Evaluating the Data. Figure 14-4 should be used in the following manner to determine if memory is a system constraint.

Step 1 - If column 7 shows a surplus of memory greater than 15% of the total available memory or greater than two times the value reported in column 4, then the implication is that there is no memory constraint on the system.

#1	#2	#3	#4	#5	#6	#7	
Date	Hours	CPU/IO Ratio	Avg Activity Size (K)	Amt of Memory available for User activities	Amt of Memory used by System Functions	Memory Surplus or Memory Shortfall	with and without the PALC Queues

#8	#9	#10	#11	#12
Avg Number of User Activities Waiting Memory	Avg Number of System Activities Waiting Memory	Avg Number of User Activities in Memory	Avg Number of System Activities in Memory	Ratio of Duration vs Memory Time

#13	#14	#15	#16	#17	#18
% Slave Memory Used	Amount of Time User Activity Swapped	Amount of Time Activity Waited for Original Memory Allocation	Number of User Swaps	Activities Per Hour (Throughput)	Total Swaps Per Hour

Figure 14-4. Memory Statistics

If Column 7 shows a surplus of memory, but does not exceed the aforementioned limits, the implication is that the current system has sufficient memory, but that the system is approaching memory saturation. If Colulmn 7 shows a shortfall of memory, and the value is greater than 15% of the total available memory or greater than two times the value reported in column 4, then the implication is that memory is a constraint on this system. Finally, if Column 7 shows a shortfall of memory, but does not exceed the aforementioned limits, the implication is that the system has reached memory saturation, but is still able to process the current workload.

Step 2 - It should be stressed that the value reported in column 7 is calculated over the entire measurement period and therefore could be biased by periods of heavy or light activity. It is for this reason that the user is urged to run the monitor during those periods of time where the workload is considered to be heavy and constant. In order to determine if the above type of biasing is occurring, the user may want to check Plots 1-3. If it appears that there is a mixing of light processing and heavy processing the user may want to re-run the data reduction program, using the time-frame option, to separate the heavy processing time.

Step 3 - Calculate the ratio of column 5 divided by column 4. This is an indication of the maximum number of user jobs that your system can support at any one time, without the occurrence of significant swapping. If the value in column 10 is equal to or exceeds this ratio, then the implication is that the system has reached memory saturation. If the value in Column 10 is within 2 units of the ratio, then the system probably has sufficient memory but is approaching saturation. Finally, if the value in column 10 is less than the ratio by more than 2 units, the current system has sufficient memory.

This step can be further verified by checking columns 14, 16 and 18 for indication of significant swapping.

Step 4 - If column 13 is less than 85, the current system should have sufficient memory and the other steps should not be showing indications of memory problems. If column 13 is between 85 and 95 then the current system is approaching saturation and at times may be showing some indications of a backlog. If the figure exceeds 95, then other steps should be indicating signs of moderate to severe memory problems.

Step 5 - If column 8 is greater than or equal to 3, the indication is that memory has become a constraining factor.

Step 6 - If Column 12 is greater than 2, the indication is that memory wait time is high and that memory is probably a constraining factor.

At this point, the user should have a fairly good indication as to whether or not memory is a constraining factor. The following steps will indicate some additional reports that the user should reference to determine those jobs that might be causing the memory problem.

Step 7 - One of the largest users of resources are jobs that abort and then must be rerun. Aborts usually occur due to user errors, but hardware aborts are not uncommon. If management is aware of aborting jobs and the reasons for them, they can possibly save substantial system resources. The Abort Report is described in subsection 6.3.8 and gives an indication of the system resources being wasted by aborting jobs.

Step 8 - It is important for management to be aware of jobs that are either misusing system resources or are requesting large amounts of system resources. Upon identifying such jobs, these jobs could be redesigned, scheduled for non-peak processing, or, in the case of wasted resources, the waste could be eliminated. The Excessive Resource Report allows the user to uncover jobs of this type and is described in subsection 6.3.10. When using this report, the following are suggested parameter values:

wasted core - 10K
memory - either 35K (WWMCCS standard) or 2 times the value in column 4
CPU time - 15 minutes
IO time - 30 minutes
URG - 40
RATIO - 2

Step 9 - By examining the System Program Usage of Memory Report, the user can determine those system type jobs that are requiring memory. It is possible that some of the system jobs can be eliminated or at least reduced in size. This is especially true for the TSS. However, it must be realized that a limitation on Time Sharing size may adversely effect TSS response. In many cases, if large file transfers are being processed during prime time, the size of the FTS WIN subsystem can rise to 70 or 80K. By not allowing WIN file transfers to run during prime time, significant memory savings can result.

Step 10 - As is explained in great detail in subsection 6.3.15, it is vitally important that the overall urgency level of jobs being processed remain low. The Distribution of Urgency Report can be used to determine the overall urgency level of jobs being processed. This report should show that 60 percent of the jobs being processed at any one time have an urgency level below 40 and that a substantial proportion of these should have an urgency level between 5-10. The summary at the bottom should indicate that 75-80 percent of all activities processed had an urgency level below 20.

If this report indicates a large percentage of high-urgency jobs, then the SNUMB/IDENT report, or the Excessive Resource Report, can be used to identify those particular activities processing with a high urgency.

Step 11 - If the analyst wants to track the memory performance of a given set of jobs, the use of the SPECL input option and the generation of the Special Job Memory Reports will provide sufficient data for detailed memory tracking. This procedure is especially useful in analyzing the memory requirements of TS1, FTS and the special JDA-developed software (JDSIP, JDSUP). Refer to subsections 6.1.24 and 6.3.14 for complete descriptions of these Special Job Memory Reports.

Step 12 - Another indication of poor system performance possibly caused by memory shortfall, tape drive shortfall, poor operator performance or a poor system scheduler design is the long delay of jobs as they pass through the various allocation phases prior to core allocation. The Allocation Status Report, the System Scheduler Delay Time Histogram and the Delay Time Until Core Allocation Histogram can all be used to determine which jobs, and how many jobs, are being significantly delayed during the various allocation phases. These reports are all fully discussed in section 6.

Step 13 - Memory problems may also be occurring as a result of jobs being delayed due to CPU constraints or I/O constraints. In these cases, jobs tend to sit in memory due to a lack of other system resources. Because these jobs are being delayed, other jobs cannot enter memory, and memory demands begin to backlog. Therefore, if memory is a constraint, the user should consider conducting a CPU analysis as well as an I/O analysis.

14.6.4 CPU Evaluation. The CPU evaluation will determine the general utilization level of the processor and then determine if the CPU is dominated by GCOS or user program execution. In addition, the CPU evaluation can be used to determine if jobs are being significantly delayed by a lack of processor power. A CPU data reduction is required for this evaluation. It is also beneficial to have an associated MUM data reduction available for the same time period.

14.6.4.1 Data Recording. The heading page of the CPU data reduction report provides the dispatcher options currently in effect on the system. Recent tests have shown that the Urgency Thruput Option should be enabled, as well as the In-Core Push Area and Dynamic Buffering of SSA Modules. The TSS and the various WIN subsystems should not be placed in Priority 3 processing. In addition, sites should try to avoid enabling the I/O Thruput Option, unless strong justification exists to decide otherwise. The CPU Time Report is produced every 10 minutes of elapsed time and the data of interest should be found in the last 10-minute report.

THIS PAGE LEFT INTENTIONALLY BLANK

14.6.4.2 Evaluating the Data.

Step 1 - The CPU may by considered a bottleneck if the % Idle CPU column of the CPU Time Report is less than 20 or the summary column at the end of the CPU Idle plot indicates that 50% of the time or more, the processor is less than 25% idle. It is a fairly agreed upon standard that average processor utilization should not exceed 80%. By maintaining average processor utilization at this level, the processor will have sufficient remaining capacity to be able to handle those peaks of processor demand that normally occur during the day. The three category breakdowns at the end of the CPU plot represent the three conditions of insufficient power, sufficient power, and excess power respectively.

Step 2 - The % Gate Loop column of the CPU Time Report provides an indication of the percentage of CPU power being lost because multiple processors are interfering with each other. Within GCOS, there are many tables that must be updated and/or referenced by a processor during its execution. When these tables are being used, the processor must insure that they are not altered in any manner. In order to accomplish this, the processor will lock a gate. This locked gate will prevent any other processor from accessing this table. When the first processor has completed its use of the table, the gate will be opened. If a processor must access a gated table, it will simply perform a "CPU loop" at that table, waiting for it to be opened. The amount of time being spent in this gate locked-CPU loop code is depicted by this column. If this value is greater than 5%, then this is an indication that the multiple processor configuration is beginning to lose its cost effectiveness.

Step 3 - The MUM Excessive Resource Report can be used to determine those jobs requiring excessive CPU resources. In addition, the CPU Plot report can be used to determine those times of day when processor availability is in the critical range. Using these two items of information, system schedular classes can be created based on CPU requirements.

Step 4 - Another indication that the CPU is a bottleneck can be determined from the CPU Queue Length histogram. If the average queue length is greater than two times the number of processors configured, then the processors are being requested to handle an excessive amount of work. This queuing of jobs at the CPU is an indication that during some period of the day, there is insufficient CPU power to handle the workload. Once again, the CPU Plot can be used to determine those periods of time.

In addition, at the bottom of each 10-minute section of the CPU Time Report, a line is printed indicating the CPU queue length during the last 10-minute period, as well as since the start of the run. These figures provide an excellent indication of those times during the day that the processor is being overloaded. Control of the scheduler queues is one method of limiting the amount of work being entered into the system.

FIGURE ON THIS PAGE HAS BEEN DELETED

FIGURE ON THIS PAGE HAS BEEN DELETED

Step 5 - The CPU Time Report and WIN Report can be used to determine those periods of time when TSS and/or WIN programs are using excessive amounts of processor time or, on the other hand, do not appear to be requesting sufficient CPU service.

Step 6 - If the Percent of Memory Time in Queue histogram shows that the average activity is spending more than 30% of its memory time waiting for the processor, this is a strong indication that processing power is a constraining factor. Once again, in order to relieve this constraint, it will be necessary to acquire an additional processor, a faster processor, or else the workload will need to be controlled via the scheduler classes. The CPU Access By SNUMB Report can be used to determine, on an activity-by-activity basis, exactly which activities are being delayed the most due to the lack of processor power.

Step 7 - If the CPU Time Report shows that the percent of system CPU exceeds 30%, this is another indication that the system software is being requested to handle an excessive workload. In all probability, the system queues are increasing to such a size that the system software is expending excessive resources managing its queues.

Step 8 - The CPU Time Report (dispatcher queuing portion) indicates the percentage of time the various system programs spend waiting for the processor and the average queue position of these programs. System programs should not spend more than 2% of the time waiting for service and their average queue position should not exceed 2. If this is not the case, the analyst should ensure that the Urgency Thruput Option is enabled. In addition, the Urgency Report and Excessive Resource Report of the Memory Monitor should be examined to ensure that there is not an excessive number of user activities processing at very high urgencies (see section 14.6.3 for memory evaluation details).

14.6.5 I/O Evaluation. The I/O evaluation will determine whether the mass storage subsystem, or tape channel subsystem is the cause of system degradation. This evaluation requires the user to have processed the Mass Storage Monitor and Channel Monitor data reduction programs.

14.6.5.1 Data Recording. All output from the Mass Store Monitor and Channel Monitor are required. No individual work tables are required, but the user may generate some if he feels that it will help in his analysis.

14.6.5.2 Evaluating the Data. Chapters 7 and 8 provide a fairly detailed description of the procedure to be followed in analyzing the associated reports. In this section, reference will be made to those chapters indicating actual data values that should be used as a reference for comparison.

Step 1 - Read subsection 7.3 and subsections 8.2 and 8.3.

Step 2 - Check the crossbar configuration using the procedure described in subsection 8.2.

Step 3 - Examine the Proportionate Device Utilization Report produced by either the MSM or CM. Check for devices which have significantly higher utilization than other devices in the system. These devices are potential bottlenecks and should be more closely analyzed. It is desirable, even though perhaps not possible, to have equal utilization across all disk packs. Read subsections 7.5.18 and 7.5.20 for further details on this step. Once a pack(s) is identified, further analysis should be performed to determine the actual files being referenced on the pack (see subsection 7.6.1).

Step 4 - The histogram displaying Data Transfer Sizes for TSS Swap Files can give a strong indication of the sizes of TSS subsystems being used. TSS subsystems of over 25K can cause significant increase in overall TSS response, especially if several of these subsystems are being executed simultaneously. If more than 20 percent of the entries in this report fall in the bucket ranges above 25000, this is a strong indication that TSS response might be a problem. This problem can be further confirmed with the CAM.

Step 5 - Seek Elongation - Subsectons 7.5.6 and 7.5.7 describe in detail the reports used to investigate seek elongation problems. An average seek of over 50 cylinders for DSS191s and 100 cylinders for DSU450s should be considered significant.

Step 6 - Analyze the Channel Monitor Idle Report. This report can be generated only if the Idle Monitor was run in conjunction with the Channel Monitor. If the "% of Idle Time During Which I/O Was Active" value exceeds 25%, then substantial benefit may be obtained by eliminating I/O contention. The above value is an indication that even though the CPU is going idle (i.e., has no useful work to perform) there really is potential CPU work available. However, under current conditions, this potential CPU work is being delayed because of I/O contention.

Even though the above figure exceeds 25%, the system may not have sufficient CPU power available to handle the increased work generated by removing the I/O contention. Therefore, the analyst should also check that the "Average System % Idle" figure exceeds 15%. If this proves to be the case, then removal of any I/O contention should prove beneficial. On the other hand, if the figure is lower than 15%, then removal of any I/O contention will probably result in additional CPU contention. The Idle Report will also indicate those devices causing most of the contention. Make a record of the device numbers.

Step 7 - Examine I/O queue length and I/O queue time histograms for individual devices and channels. Queues greater than one and queue times greater than 15 MS should be considered significant. Record those devices with high contention.

Step 8 - If certain devices have been determined as bottlenecks under the procedures described in Steps 1 and 2, the Job Conflict Report should be obtained for those devices following the procedures described in Chapter 8.

Step 9 - Execute the Mass Store Monitor Data Reduction Program. Following the procedure described in subsection 7.6.1 for monitoring a specific device, the analyst should be able to determine the exact files that are causing the contention found under the earlier steps.

Step 10 - Using the CM, it is possible to perform a detailed analysis on channel queuing for a particular job. Details for this procedure can be found in subsections 8.5.13, 8.5.14 and 8.6.10.

Step 11 - This step outlines procedures for relocating files identified as candidates for file relocation. Because of automatic load-leveling activity by the GCOS operating system, an analyst has only limited flexibility for the placement of system, permanent, and temporary files:

a. System Files. The device name on which a system file is to be placed can be specified at system startup. Care should be taken to insure that multiple high-used system files are not placed on the same disk device. If possible, separate high-use system files across disk subsystems. In addition, ensure that SSA Cache memory and FMS cache are enabled to reduce disk I/O activity to certain system files. Details for this analysis can be found in subsections 7.5.9, 7.5.10 and 7.5.19.

b. Permanent Files. The device name for a permanent file can be specified at creation, whether through FMS or the ACCESS subsystem of Time Sharing. Files can be moved by changing their names, creating a new file with the old name, and moving the data. The new file can be created with a DEVICE specification.

c. Temporary Files. The device name for a temporary file can be specified in the second field of the $FILE card in the job control deck. Jobs which run frequently can have their $FILE cards changed. Other jobs can be controlled by policies governing the use of $FILE cards.

Additionally, sites that have different device types may specify preferred device types to be used for temporary files. This procedure will allow activities requiring disk storage to take advantage of higher speed devices.

Step 12 - This step identifies possible seek contention problems attributable to inadequate temporary file space. This procedure uses the SPUTIL feature of the GCOS FMS facility and the Disk Fragmentation Report available at most sites. If such a report is not available, contact CCTC/C751. It is necessary to analyze temporary disk capacity on all disk units rather than just the units identified in previous tuning steps. This analysis is necessary because the disk units exhibiting high activity due to temporary file use often have more available temporary space. The increased utilization of these disk units may be caused by inadequate temporary storage on other disk devices. For this analysis a form as shown in figure 14-5 may prove useful.

a. Report Values. The SPUTIL Report contains the following information on each disk device: (1) device identification, (2) overall capacity, (3) available disk space, (4) disk space dedicated to permanent storage. The Disk Fragmentation Report contains additional information on each disk device: (1) number of disk fragments, (2) average fragment size, (3) maximum fragment size, (4) percentage distribution of fragments by size, and (5) total fragmented space.

b. Form Entry. Enter the device identification for each disk device on the Temporary Storage Test Form. For each disk device enter: (1) the LLINK capacity as indicated by the SPUTIL Report in the Total Capacity column; (2) the temporary capacity from the SPUTIL Report in the Temp Capacity column; (3) the number of fragments from the Disk Fragmentation

DEVICE	TOTAL CAPACITY	TEMP CAPACITY	RATIO 1	RATIO 2	NUMBER OF FRAGMENTS	MAX SIZE	PERCENT 1-12	FRAGMENTS 13-121	TOTAL PERCENT	FRAGMENT

Figure 14-5. Temporary Storage Test Form

Report in the Number of Fragments column; (4) the maximum fragment size in the Max Size column; (5) the percentage of fragments (1-12 and 13-121, respectively) in the Percent Fragments column.

c. Calculations. Divide the temporary capacity of each disk device by the total capacity, and place the result in the "Ratio 1" column. Add the Percent Fragments 1-12 and the Percent Fragments 13-121 columns and place the sum in the Total Percent column.

d. Decision. Place a check mark in the "Ratio 2" column if the ratio of temporary storage to total storage is (.15). Place a check mark in the "Fragment " column if (1) the number of fragments is 100, (2) the maximum size of a fragment is 2000, and (3) the total percentage of fragments less than 121 LLINKS (10 LINKS or less) is 40.

Step 8 - If devices have been checked in Step 6, then the temporary file space on these devices is constrained by either (1) insufficient file space ("Ratio "), or (2) disk fragmentation ("Fragment ").

If a large number of devices need to be checked, then overall temporary disk space availability may be a constraint to system performance. At this point a site should either institute procedures to recover permanent disk space (purging of unused files) or re-evaluate disk capacity relative to system workload.

If certain disk devices have been checked because of their temporary to total ratio, or if many devices have been checked because of disk fragmentation, then the following procedures to balance the file system should be instituted.

a. Fragmentation. Temporary disk space may be compacted by a full restore of the file system (cold boot). The Disk Fragmentation Report, run daily, could help determine the frequency at which a full file system restore would be necessary.

b. Amount of Available Space. The full file system restore redistributes permanent files to equalize the amount of temporary storage on each pack. The FMS SPUTILST Report should be examined after the restore to verify that disk units have been balanced. Temporary disk space may still not be evenly distributed if (1) users have specified specific devices for permanent files or (2) unusually large files are present on certain disk units. Consider such factors as intentional file placement for maximum seek activity and operational requirements of data bases before moving permanent files.

14.6.5.3 Mass Storage Operation. The following article appeared in the September 1981 WWMCCS H6000 World Magazine. It presents an excellent description of disk operation and some problem areas that a system analyst should be aware of.

"We won't discuss the gyrations that GCOS performs in setting up an I/O request; e.g., management of PAT pointers, PAT bodies, etc. (We recommend the HIS GCOS Analysis course). Our interest begins at the point when an I/O request generated by some workload element is presented to the I/O Supervisor (IOS) for service. We don't care whether the I/O came via GFRC or directly from a MME GEINOS. We do care about the major logical/physical events from this point to I/O termination which are important to understanding system performance.

"As an example, let's assume that we have a dual-channel micro-programmable controller (MPC) cross-barred to its disk devices. Dual-channel means that the MPC has two physical data paths for transferring data to the IOM. Cross-barred to the device means that there are two physical paths from the MPC to each of its devices.

"Let's also assume that the IOM is cross-barred at two-to-one. IOM cross-barring refers to the logical-physical hardware channel relationship of the IOM. This relationship is defined to GCOS in the STARTUP deck and maintained in the Secondary Configuration Table and elsewhere. For the example, assume that PUB addresses 0-8 and 0-9 are associated with one physical MPC channel, and that 0-10 and 0-11 are associated with the second MPC channel as is shown below:

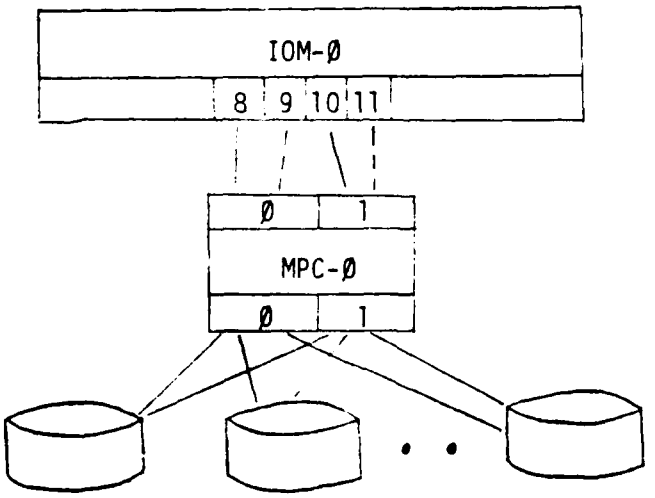


Figure 1

"Defined in the STARTUP deck is the primary PUB address for the disk subsystem and the order in which logical paths will be assigned by GCOS during processing. The card

```
$       IOM-0           PUB-8,MSO450,UNITS-6,
$       ETC             UNIT-1,ST1
          :
```

defines PUB 0-8 as the primary PUB address for a MSS451 subsystem. The card

```
$       XBAR            IOM-0,PUB-8,PUB-10,PUB-9,PUB-11
```

defines the order in which PUB assignments will be made. This information is maintained by GCOS in the I/O stream table. All requests for a device are made with its primary PUB address, e.g., an I/O request for device 6 on this subsystem would request IOM-PUB-DEVICE = 0-8-6.

"Suppose that IOS receives a request for I/O to 0-8-6. If the device is currently performing an I/O (i.e., busy), the I/O request is queued for later service. If the device is free, IOS will assign a logical channel to the request if one is available. It will attempt to assign PUB 0-8. If 0-8 is busy, it will attempt to assign 0-10, etc., in the order prescribed on the $ XBAR card. If no logical channel is available, the I/O request will be queued for later service as before.

"To recap, for an I/O request to become active (in-service), both the requested device and a logical channel associated with the MPC which controls the device must be free. Otherwise the request is placed in the I/O queue. Normally the request will be placed at the end of the queue; i.e., behind earlier I/O requests. However, certain I/O requests are linked in front of the queue, e.g., TSS requests for access to its swap files (#S, #T, #U, #V). This can impact system performance and will be discussed later.

"Whether or not our request for I/O to 0-8-6 becomes active, IOS will scan the I/O queue to determine if any earlier I/O requests can now be activated (device and logical channel free). When an I/O request can be activated, IOS loads the mailbox in Hard Core Memory for the logical channel and signals the IOM. This mailbox contains information required by the IOM to process the I/O independently of the mainframe processor. The IOM loads the mailbox information into its scratch pad memory and begins processing the physical I/O.

"Thus far, the events which have occurred were independent of the type of disk device involved. The three phases of physical delay for a disk device are seek time, rotational latency, and data transfer. The IOM issues a seek request to the MPC for 0-8-6, but now performance implications arise which depend both on the physical characteristics of the device, and on device-dependent interactions with IOS which might be required.

"First let's look at seek time, defined as the time required to move the read/write heads from their current position to the cylinder requested by the current I/O. For a given device, seek time is proportional to the number of

cylinders traversed. Table 1 gives seek times for H6000 disk devices. The table provides sufficient information to derive the average seek time for a device from the average number of cylinders moved (available from MSM).

Device	DSS181	DSS191	MSU450
Minimum	10 ms	10 ms	8 ms
Average	34 ms	30 ms	25 ms
Maximum	60 ms	55 ms	55 ms
Cylinders/pack	200	404	808

Seek Specifications
Table 1

"Now that "seek" has been defined, we can discuss the motivation of implementing the logical-physical channel architecture. The disk subsystem can have as many seeks and data transfers ongoing as there are logical channels assigned to the subsystem. The number of simultaneous data transfers that can be performed is limited by the number of physical channels. In our example, we could see two devices connected to the physical channels transferring data and two additional devices moving the read/write heads to the requested cylinder (seeking). Or, we could have one data transfer ongoing with three seeks, or we could have four seeks ongoing, etc. So the purpose of IOM cross-barring is to increase the simultaneity or overlapping of device seeks.

"When the seek to the requested cylinder has been completed, the device is ready to transmit or receive data. Here again a device-dependent performance consideration comes into play. For a DSS191 or MSU450 disk subsystem, once IOS has loaded the channel mailboxes and signalled the IOM, it performs no further service to the I/O request until an I/O termination interrupt is received for that device. The determination of seek completion and assignment of the physical channel for the data transfer takes place completely in the MPC. Not so for the DSS181 subsystem.

"The DSS181 does not provide a seek complete interrupt, and so IOS must check for a seek complete status each time it gets control for an I/O request as described above or for servicing an I/O termination interrupt. This introduces another delay for the 181s from the time seek completes until the next time IOS is activated. When IOS detects seek complete for a 181, it must issue another command to connect the physical channel and begin data transfer.

"If both of our physical channels are busy transferring data, there is another delay until a physical channel is free. But let's suppose the seek has completed, a physical channel is available for connecting to the device (whether by the MPC or by IOS) and we are ready to transfer data. Another delay is imposed now due to the rotational latency of the device. This is the time from physical channel assignment until the requested sector rotates to the read/write heads. This delay is a function of the rotational velocity of

the device. The minimum latency for a device is always 0. (The data is at the heads at channel connect time.) Maximum latency for a device is the time required for 1 revolution of the disk, and average latency is generally defined as the time required for 1/2 revolution (see Table 2).

Device	DSU181	DSU191	MSU451
Rotational Speed	2400 RPM 40 RPS	3600 RPM 60 RPS	3600 RPM 60 RPS
Avg. Latency	12.5 ms	8.3 ms	8.3 ms
Max. Latency	25 ms	16.7 ms	16.7 ms

Rotational Latency
Table 2

"A word about Rotational Positional Sensing (RPS). RPS is a feature available on some disk subsystems -- a required option (?) on MSU451s -- which keeps the MPC informed about the rotational position of each disk. In a situation where the physical channel is free and more than one device has completed it's seek, the MPC can assign the physical channel to the unit whose requested sector is closest to the heads. The idea is to reduce rotational latency thereby increasing the physical channel time which can be utilized for data transfer. What does RPS buy you if only one disk unit has completed its seek when the physical channel is assigned? Not a thing.

"Now that device 0-8-6 has been assigned the pnysical channel and the requested sector is under the heads, the physical transfer of data begins. The time to complete the transfer of data depends on the number of words to be transferred, the rotational velocity of the device, and the recording density of the media. Data transfer rates per second are given in Table 3.

Device	DSU181	DSU191	MSU451
Transfer	416KC	1074KC	1074KC
Rate	69.33KW	179KW	179KW

KC = 1000 - 6 BIT CHARACTERS
KW = 1000 - 35 BIT WORDS

Data Transfer Rates
Table 3

"When data transfer has completed, an I/O terminate interrupt is issued to the H6000 CPU (CPU 0). The GCOS dispatcher suspends execution of the program in operation, performs a few housekeeping duties and then dispatches CPU 0 to the disk channel module in IOS. A bit more housekeeping and, for our purposes, the I/O to 0-8-6 is complete. IOS then checks the I/O queue for seeks to start -- note that at this point it has a free logical channel, and so can start at least one seek if an I/O request is pending against a free device on this subsystem. If a DSS181 subsystem is present, it must also check for a seek complete status as described above. When IOS can perform no more actions, it relinquishes control to the dispatcher.

"In terms of capacity, a DSU1191 can hold as much data as approximately 4.3 DSU181 packs (see Table 4). A MSU451 can store roughly 2 DSU191 packs or 8.6 DSU181 packs worth of data. In terms of data transfer rates, the 191s and 451s are about 2.6 times faster than 181s.

Device	DSU181	DSU191	MSU451
Char. per pack	27.6mc	117.9mc	235.8mc

Device Capacities
Table 4

Here are a few WATCHITs.

"WATCHIT 1: When replacing 181s with 451s. The vast capacity of the 451s make it feasible to replace up to 8 181s with a single 451. In a situation where 4 I/Os were pending to files residing on 4 different 181s (our 2 to 1 cross-
barred dual channel example), as opposed to the 451, the 181 configuration might be preferable. Although the 451 is 2.6 times as fast, we would sacrifice overlapping of the seeks, and the 181 subsystem can transfer data from 2 of the devices in parallel. The CM Device I/O Queue Length Histrogram and the Channel Free-Device Busy Report can indicate if a device contention problem exists.

"WATCHIT 2: When pondering vendor specs for 181s. We modified the MSM data reduction program to obtain the average data transfer size to each disk device. MSM also provides the average cylinders seeked per I/O for each device. Assuming the average rotational latency, one should be able to use the vendor specs provided in this article to compute the I/O service time for a given device. The I/O service time for each device is provided as a CM report. For MSU451s the computed value matches the CM report to within a few milliseconds. However the 181 computed and monitored values differ by as much as 30ms and more -- the computed value always being less. Since channel utilization is fairly low, the most likely reason for this discrepancy is that IOS latency time we discussed earlier.

"WATCHIT 3: When placing TSS swap files (#S, #T, #U, #V). We mentioned earlier that TSS requests for access to it's swap files were linked to the front of the I/O queue. TSS is not one to wait on user I/Os, nor does it break it's swap action into nice little 320-word standard system format I/Os.

It's gonna swap the whole thing. Say, there's a nice little TSS COBOL subsystem -- only 40K! Time to swap? So soon? Oh yeah, that memory-time quantum. Lot of memory -- not much time. Oh well, a 40K transfer is OK. My swap files are on fast 451s. Friend, you just tied up that device for almost a quarter of a second, and more importantly, a physical channel as well. That's only half of it. TSS swapped it out, you can bet it'll swap it back in. So we're up to .5 sec of physical channel time to free up some TSS user memory. This scenario impacts total system throughput, not just TSS. That's not all the havoc that larger user subsystems create, but that is enough heartburn for now."

14.6.6 Communication Evaluation. The communication evaluation will determine the overall terminal usage of a system. It can also be used to examine the DN355 usage. This evaluation can be done using either the CAM or the GRTS monitor, or both. Figures 14-6 through 14-9 are sample table formats that may be used to display the gathered data.

14.6.6.1 Data Recording. For figure 14-6, the Terminal Session Report is used. All users with TSS subsystems of >35K are recorded. Column 7 is obtained from scanning the Excess Think/Response Time Report. Figure 14-7 is obtained from the Response Time Report. All periods of time when the response for TSS is greater than 15 seconds are recorded. Figure 14-8 is obtained directly from the High Terminal Usage Report. All DAC terminals with over 90 percent usage, except WIN lines, are recorded.

Figure 14-9 is obtained from the H6000-DN355 Reject Report and the Abort Report. The H6000-DN355 Reject Report is used if the average number of reject commands per hour of running time is greater than 50 (total number of reject commands/number of hours in run). Terminals with more than 30 percent of the rejects should also be listed. For the Network Control Program (NCP) disconnects, all NCP 01 line disconnects listed in the Abort Report should be tallied.

14.6.6.2 Evaluating the Data. The following procedure should be followed in order to analyze the data.

Step 1 - TSS response time is dependent upon certain TSS parameters. The TS1 Initial Parameter Report contains current settings of these parameters. The critical parameters are:

a. Initial TS1 Max Size - If operators must increase max size during the day, TSS can slow its processing of user responses until it can grow. This should be set to the normal maximum size TSS reaches.

b. Size Growth/Reduction Factor - These sizes should be identical or growth can be twice reduction.

c. Max Number of Terminals - Should be large enough to satisfy all possible users.

#1	#2	#3	#4	#5	#6	#7
DATE	TERMINAL ID	USERID	LOGON TIME	LOGOFF TIME	TSS SUBSYSTEM SIZE	NUMBER OF EXCESSIVE RESPONSES

Figure 14-6. Large TSS Subsystem Users

DATE	TIME	RESPONSE TIME	IN-RANGE RESPONSE	NUMBER OF OUT-OF-RANGE RESPONSES	NUMBER OF USERS

Figure 14-7. Poor TSS Response Log

DATE	TERMINAL ID	PERCENT USAGE	NUMBER SESSIONS	NUMBER INPUTS	NUMBER OUTPUTS

Figure 14-8. High Usage Terminals

DATE	TERMINAL ID	OPCODE REJECTS	PERCENT REJECTS	NCP DISCONNECTS

Figure 14-9. Terminal/Line Error Report

d. Large Subsystem Size/Wait Time - The average subsystem size should be less than 35K and all large users (greater than 35K) should be penalized.

e. Number Swap Files - On an active TSS system, should be four.

f. Allocated Devices - Swap Files (#S, #T, #U, #V) should be on specific devices, as determined by MSM.

Step 2 - Users of large TSS subsystems cause TSS to swap other users out and therefore generate extra TSS and system I/O overhead. If an excessive number of users are using subsystems of greater than 35K (figure 14-6), these users should be queried as to what subsystem they are using. Any large subsystem that has high usage should be investigated for possible rewrite.

Step 3 - The overall system and TSS response time should be monitored using figure 14-7. Periods of bad response should be checked to ensure the bad responses are not caused by a few (less than 10 percent of all in-range responses) out of range responses. If TSS response is truly poor, a correlation between response and large subsystem usage should be attempted (figure 14-6 and figure 14-7).

Step 4 - Terminals logged onto the system for long periods of time (greater than 75-80 percent of the monitoring session), but having few inputs and/or few inputs/outputs, should be investigated (figure 14-8). Terminals with few inputs, but many outputs, are probably logged onto VIDEO. The number of users on VIDEO should be restricted to one or two per DATANET, due to the buffer load place on the DATANET by VIDEO. If more users are required, monitor output terminals should be used instead of more VIDEO users. Terminals with few inputs and few outputs may be being used just to keep a terminal logged on. This practice causes unnecessary TSS size and processing overhead.

Step 5 - Terminals with a large number of Opcode Rejects (figure 14-9) are an indication of possible line problems. The terminals and lines should be checked for noise and transmission errors. Numerous NCP disconnects per day (figure 14-9) is usually an indication of line or IMP problems. NCP should have fewer than five disconnects per day.

SECTION 15. TIMESHARING EMULATION AND RESPONSE TIME SYSTEM (TEARS)

15.1 Introduction

The Timesharing (TSS) Emulation and Response Time System (TEARS) is a collection of FORTRAN programs which sequentially processes data the Timesharing Subsystem Monitor (TSSM) collected and written to tape. There are three separate and independent phases of reduction: response time analysis, TSS emulation, and optionally, a formatted dump. TEARS produces a number of reports showing the user load supported by TSS, computer resources expended, user response times, and service interruptions. A list of reports and report descriptions is presented in subsection 15.5.

The TSS Monitor and its companion data reduction package, TEARS, provide information to:

- o Identify those time periods in which TSS response to the user was abnormally poor.
- o Identify which resources -- CPU, I/O, memory, or non-TSS competition -- contributed to the poor response.
- o Localize the agent of poor response down to the users command, user or subsystem.

TEARS uses two inputs. The first is the data tape produced by the TSSM in the General Monitor Collector (GMC). The second input is a set of report option control cards used to alter the reports in a way other than the standard default. The various options and their formats are described in subsection 15.6.

Caveat: Section 15 is a general description of TEARS, its reports, and run-time options. It is not intended as a users manual for analysis. Little is provided here as "rules-of-thumb" toward identifying a TSS session showing abnormal behavior or as guides to interpreting the reports. An extension to section 15, in preparation but not currently available, will explain how to interpret the reports and perform an analysis.

15.2 Data Collection Methodology

The TSSM in the GMC uses trace type 74 to collect data about the TSS in execution. TSS transactions are trapped at 104 points within TSS and at one point external to the TSS; each trap produces a trace 74. Section 5 provides a description of the trace mechanism and of the 105 trace subtypes. With the information thus collected, TEARS maintains a list of active users and correlates TSS transactions with a particular user. The interaction between a user terminal and TSS is traced as is the interaction between GCOS and TSS. Processor allocation, memory allocation and user disk I/O are also sources of data.

15.3 Analytical Methodology

TEARS is a multipass program that requires analysis after each pass. The first pass displays user response time (technically, duration of line idle -- from the H6000 side of the DATANET). That display is a collection of charts, each of which shows a demand or response characteristic from minute to minute over the duration of the run. Using the charts* an analyst can determine selected timeframes of poor response for a second pass analysis of the data tape.

The second pass analysis mimics or emulates the TSS process in a limited fashion -- it remembers, dynamically, what each TSS user was doing or attempting to do, how long the attempt was ongoing, the obstacles encountered in the attempt, and some of the computer resources and TSS executive services used in the attempt. Additionally, the second pass identifies events that were of severe import or were catastrophic to the user. Several charts, histograms, and tabular reports provide both summary and snapshot of the mimic phase. Sufficient data is presented for a manual or deductive completion of the analysis.

TEARS will generate additional reports if the CPU Monitor was active.

While the response time and emulation phases are independent, the reports produced are complimentary -- there are several auxiliary reports generated by the response time phase useful for determining the distribution of CPU resources across the entire system and also within TSS.

An optional third pass is an investigative and debug tool for the more technical user. It provides a formatted dump for examining TSS transaction content, sequencing and anomaly.

TEARS produces tabular reports, histograms, and logarithmic bar charts as tools for analysis. Of the three types, the logarithmic bar chart is new to GMF reduction programs - it is a variation of the standard plot seen in other sections of this manual. Most of the bar charts described in following subsections represent, in two ways, the duration of some frequently occurring process. The chart is ordered by clock time and is arranged vertically on the page. The time interval between print lines is under user control with a granularity of one second. For each print line, the left half of the chart is a histogram (or interval-o-gram) showing how many such processes were observed in each of several duration lengths (buckets) during the observation interval. If the process is attributable to a TSS user, the user with the longest duration is identified under the column heading "LINE". The right half of the chart prints the minimum,

* Independently, the CAM Monitor and associated data reduction package can identify periods of poor response.

the average, and the maximum duration as horizontal bars on a logarithmic scale. The usual way of printing a multivariable bar chart is to lay the bars side by side:

```
          minimum      mmm
  time    average      aaaaaa
          maximmum     MMMMMMMMM
```

To make the charts compact, we have overlaid the average bar with the minimum bar and then that composite bar overlays the maximum bar:

```
  time                 mmmaaaMMM
```

If 2 or more bars are of equal length, a digit "2" or "3" is printed in lieu of the characteristic bar symbol. (A "2" is used if 2 bars are of equal length; "3", if all 3 bars are of equal length.) If one or more bars fails to cross the left margin (axis), the axis symbol, "I", is replaced with an asterisk. If one or more bars would break the right margin, an asterisk is printed in lieu of the bar symbol.

For each observation interval (print interval), a process is counted if (1) it terminates during the observation interval, or (2) it is ongoing at the end of the interval. In either case, the duration is measured from the beginning of the process, not from the beginning of the observation interval. Therefore, a process which spans two or more observation intervals will be counted once in each observation interval, and its reported duration may exceed the observation period.

These bar charts are intended as visual displays. It is patterns and trends and order-of-magnitude changes that are important, not the ability to convert a bar length to a numeric with accuracy (the line printing mechanism forbids accuracy anyway). The default method of displaying the y-axis shows numeric values at each order-of-magnitude only:

```
          1.0E01                                          1.0E02
          Y-----------------------------------------------Y
```

By user option (option (logscale)), intermediate values can be displayed:

```
          1.0E01                                          1.0E02
          Y--------------2--------3-----4----5---6--7-8-9-Y
```

15.4 Data Reduction Methodology

TEARS is configured to correlate up to 50 concurrent TSS users (real and deferred) with user-generated traces. If the GMC data tape shows a TSS user load greater than 50, TEARS will terminate the reduction pass and print all accumulated reports to that point. To increase (or decrease) the 50-user limit, edit, then recompile the source file B29IDPX0/SOURCE/TEARS as follows:

```
CASE
B RVS:/USTSUP = 50/;*:/USTSUP = NN/
```

where NN is the new concurrency limit. It is fruitless to set NN beyond 100, because the TSS Monitor will generate no traces with a TSS user load that large.

TEARS, like most other GMF packages, uses the system controller clock to measure the passage of time. Precision of the clock is 1 microsecond, and its capacity is 35 bits. Unlike the other GMF packages, no provision exists in TEARS to account for clock overflow. TEARS is thereby limited to reducing GMC sessions shorter than approximately 9.5 hours. On reaching the time limit, TEARS will terminate the reduction pass gracefully.

15.4.1 Response Pass Reduction Methodology. This pass performs an overview of TSS response during the data collection period. The reports generated give an indication of the periods of time when response, as observed by the user, was poor.

To accomplish this, the response pass keeps track of the following information:

Start and stop of all terminal I/Os
Start and stop of all user subdispatches
Start and stop of all disk I/Os
Requests for core by subsystems
Changes of each user's program stack
Special situations (*JDAC, CONN, etc.*).

To produce the response time logplots, the program measures the amount of time from the completion of a terminal I/O until the initiation of a subsequent output transmission. This value, the line-idle period, is used as a datum for the "Response Time For All Users" logplot, provided none of the following situations have occurred:

If the user's subsystem executes a DRL WAKE, it is assumed that the user knows that the line idle is being extended (e.g. DJST). If the user performed a line-switch to another program, the line will be idle as far as TSS is concerned, but not used in the logplots. If the user performed a non-graceful logoff (e.g. $*$D, or CTRL-C), the UST which was set for reconnect will not contribute to the logplot.

Additionally, each datum used in the All-User logplot will be used in one of the other two response time logplots. The distribution of data between the two logplots is based on the presence of a request-for-core trace (TSS subtypes 77-78) during the line-idle period. This produces a segregation of responses for users in build mode or with program-in-core versus users requiring a program call or normal swap in. Comparison of these reports will indicate when poor response may be due to core contention (provided sufficient users are logged on to produce enough plotting data). This

problem would be indicated by a degradation on the "Response Time For Users With Core Request During Line Idle" logplot without a corresponding one on the "Response Time For Users Not Requesting More Core" logplot.

These logplots can be distorted by a user who performs an extensive amount of disk I/O or CPU work without doing any terminal I/O. The effect would show up as a long period of line idle because the reduction program could not invalidate that time datum. As a run-time option, the amount of time spent doing disk I/O and the amount of processor time can be subtracted from each datum, thus reducing, but not eliminating, the erroneous effect.

To produce the subdispatch logplots, the program measures the amount of time from the point where a user is inserted into the subdispatch queue, to the point where the user is processed as a fault entry after receiving a subdispatch. The logplot "Total Time In Subdispatch Queue" reflects this time value. The logplot "Time in Subdispatch Queue Waiting Service" is the total time value minus the actual processor time used by the subsystem. The latter value has been separated out to produce the logplot "Processor Time in Subdispatch."

15.4.2 Emulation Reduction Methodology. This, the second phase of TEARS, simulates the TSS process in a limited fashion. Each user is placed in one of several defined states (table 15-1) corresponding to the manner in which TSS processes a user. For each user, TEARS maintains a state-stack; pushing, popping, or clearing of the state-stack is driven by a complex algorithm. The nature of the current trace for the user, the user's Time Type (a TSS classification), and his current and previous states interact in deciding which state-stack operation will be performed. Several states are considered vulnerable for a user -- if he remains in one of these states too long, a service for him is being delayed or denied. In a similar fashion, the TSS Executive may be denied CPU access thus delaying all users. These separate events, user delays in vulnerable states, and TSS Executive delays, are examined and reported on together so that the effect of one on the other may be seen.

As each user transits from state to state, several statistics are updated. They indicate possible bottlenecks in CPU processing, I/O, or memory allocation from minute to minute. The statistics are taken across all users and provide a functional representation of TSS bottlenecks. (The CPU reports are created in the Response phase described in section 15.5.1.4, but find application in the Emulation phase).

Concurrently with the simulation, each trace is examined as an indication of TSS trouble. Just by their occurrence, several traces indicate TSS or GCOS trouble. Over 40 conditions, many concerned with derail errors, are taken as exceptions and are items for reporting.

Table 15-1. Emulation Phase User States

#	Mnemonic	Description	Vulnerable?
1	GWAKE	In Gewake	
2	LIDLE	Line is idle	
3	WTSUB	Eligible for subdispatch	
4	RECON	In reconnect mode	
5	SWPOU	Swapping out	Yes
6	SWPIN	Swapping in	Yes
7	DSKIO	Doing disk I/O	Yes
8	SUBDS	In subdispatch queue	Yes
9	WTMEM	Waiting for memory	Yes
10	DRLSR	Waiting for derail service	Yes
11	NOTSS	Non-TSS process	
12	FMSER	Waiting FMS service	Yes
13	WTSMC	Waiting SMC service	Yes
14	LBUSY	Line is busy	
15	IDLME	Idle in memory	
16	UNKWN	Unknown (initial state)	
17	LOGOF	Logged off (final state)	

15.5 Output

Each of the three reduction phases (passes) provides independent but complementary reports. Each report begins with a banner line identifying the TEARS program, the version, and the version date. The banner is usually followed by a line identifying the data source -- this includes system identification, time of day, day of week, date (year-month-day) and reel number of the lead-off GMC data tape. The day of the week and date are continually updated and reflect the correct GMC start time or, when user selected, the time datum of the first collectable trace within a timeframe. Following the data source line is the report title and report number.

The reports generated by the TEARS reduction passes fall into three categories: summary, interval driven, and event driven. Summary reports are generated at the end of a timeframe or at the end of the pass (e.g., Histograms). Interval driven reports are generated at the end of a short time period (defaulting to 60 or 20 seconds) repeatedly during a timeframe (e.g., Logplots). Event driven reports contain data generated when certain traces or conditions are encountered during the processing of the collection tape. The information is organized as a chronological listing.

One report, the Execution Summary (figure 15-1), is common to all three reduction phases; it shows what TEARS found on the data files. Included in this information is the following:

- o A list of the monitors in execution during the GMC data collection.
- o The time, date, tape number and system controller clock at the beginning of collection.
- o An echo of the user input options, and the program's interpretation of them.
- o If the timeframe option is used, a report of when the various timeframes were reached.
- o If the collection period exceeded 9+ hours, an indication that a 35-bit internal clock overflowed (TEARS will then terminate the pass).
- o The time, date, tape number and system controller clock at the end of data collection, but only if a termination trace is processed (not currently implemented).
- o When the NEW option is user selected, the items above are repeated on a new page.

All remaining reports are generated by the individual reduction passes.

The reports that the Response pass produces are:

Trace driven reports:

- o TSS Reduction Event Log (figure 15-2) (no ID #).

```
TSS REDUCTION, VERSION 7.2 - 1.000, 30 SEPTEMBER 82

     ACTIVE MONITORS:  ****  ****  CPUM   ****  ****  ****  ****  ****  ****  **** TSS    ****

     DATA COLLECTION BEGAN WITH TAPE 1469   AT 0818:33.201 FRI 82-05-21,
     RSCRA/Q WERE  000000000031 460323216133

CARD INPUT IS INTERPRETED AS FOLLOWS:

 1:OPTION(-SUBDISPATCH) OPTION(-DISKIO)
     WILL DO -SUBDISPATCH
     WILL DO -DISKIO

 2:CPRINT(20)
     PERIODIC REPORT REPETITION PERIOD IS NOMINALLY 20         SECONDS.

 3:NEW(1469)
     TAPE 1469   WILL LEAD OFF A FOLLOWING REDUCTION.
      *-->  BREAKPOINT AT 0818:54.01



TSS REDUCTION, VERSION 7.2 - 1.000, 30 SEPTEMBER 82

     ACTIVE MONITORS:  ****  ****  CPUM   ****  ****  ****  ****  ****  ****  **** TSS    ****

     DATA COLLECTION BEGAN WITH TAPE 1469   AT 0818:33.201 FRI 82-05-21,
     RSCRA/Q WERE  000000000031 460323216133

CARD INPUT IS INTERPRETED AS FOLLOWS:

 1:ON(FORMATDUMP)
     FORMATTED DUMP TURNED ON .
```

Figure 15-1. Execution Summary Report

TIME SHARING REDUCTION EVENT LOG

```
0818:33.20 ***** MAIN INITIALIZED *****  20.00 SEC. PLOT INTERVAL *****
0818:43.51       0 *** TIMEFRAME ******** BEGIN *****
0818:54.01       1= 90=    -    - = ***** TRACE ON *****
0818:54.01       5= 90=    -    - = *****  4 USERS *****
0823:21.44    7149= 71=    -    - = ** UST AREA INCREASE.    5 USERS.
0823:21.51    7184= 32=<TA>=NEW = IOGON    HIS8010
0824:06.74    8785= 32=<TK>=NEW = IOGON    C433
0826:44.40   16814= 32=<VV>=NEW = IOGON    C722
0826:45.04   17155= 71=    -    - = ** UST AREA INCREASE.    7 USERS.
0828:44.92   21220= 32=<CG>=NEW = IOGON    WSSIB
0834:41.42   38142= 32=<TI>=NEW = IOGON    PRC118
0835:48.38   41838= 44=<DI>=JDAC= LINE SWITCH TO INQUIRY NAME VIDEO
0842:34.15   66305= 12=<TI>=TERM= IOGOFF   PRC118        ON AT  834.41
0843:09.43   67126= 32=<TX>=NEW = IOGON    SSLSM
0843:47.87   68557= 32=<T2>=NEW = IOGON    PRC101
0844:18.47   71175= 12=<T2>=TERM= IOGOFF   PRC101        ON AT  843.48
0844:45.77   72586= 98=    -    - = ** UST AREA INCREASE.    9 USERS.
0844:47.35   72663= 32=<3X>=NEW = IOGON    IRCNT1
0845:15.51   74967= 32=<DN>=NEW = IOGON    HIS
0845:20.80   75598= 32=<T3>=NEW = IOGON    CSC
0845:21.20   75739= 71=    -    - = ** UST AREA INCREASE.   13 USERS.
0845:21.31   75788= 32=<T2>=NEW = IOGON    PRC101
0847:19.78   96856= 22=<T2>=EDTX=   LINE SET FOR RECONNECT
0847:19.79   96863= 12=<DN>=TERM= IOGOFF   HIS            ON AT  845.15
0847:22.77   96886= 22=<TK>=CMDI=   LINE SET FOR RECONNECT
0847:22.78   96889= 22=<TN>=DJST=   LINE SET FOR RECONNECT
0848:38.66  100204= 32=<  >=NEW = IOGON    CSC
0848:59.64  104802= 12=<  >=TERM= IOGOFF   CSC            ON AT  848.35
0849:00.78  105160= 32=<T2>=NEW = IOGON    PRC101
0849:15.98  105889= 58=<T2>=CONN=   CONN COMMAND TO =T2=
0849:16.21  105968= 12=<T2>=TERM= IOGOFF   PRC101        ON AT  849.01
0849:54.15  107675= 32=<CH>=NEW = IOGON    ACMIS
0850:19.96  108868= 12=<TN>=TERM= IOGOFF   PRC104        ON AT  816.23
0850:20.14  108893= 12=<TK>=TERM= IOGOFF   C433           ON AT  824.07
0850:30.26  109360= 32=<TN>=NEW = IOGON    CSC7592
0851:44.14  113369= 12=<CG>=TERM= IOGOFF   WSSIB          ON AT  828.45
0852:21.28  114942= 12=<TA>=TERM= IOGOFF   HIS8010        ON AT  823.21
0852:53.32  116852= 32=<DJ>=NEW = IOGON    C432
0854:03.46  121808= 32=<TK>=NEW = IOGON    PRC118
0854:05.48  122391= 32=<TS>=NEW = IOGON    PRC105
0855:19.36  130701= 32=<TI>=NEW = IOGON    CSC8121
0856:50.36  143038= 22=<TX>=NONE=   LINE SET FOR RECONNECT
0859:48.10  179626= 12=<TX>=TERM= IOGOFF   SSLSM          ON AT  843.09
0859:59.61  180159= 32=<  >=NEW = IOGON    CSC
0900:14.27  184775= 12=<  >=TERM= IOGOFF   CSC            ON AT  859.58
0900:53.82  188034= 12=<DJ>=TERM= IOGOFF   C432           ON AT  852.53
0900:53.82  188035= 99=    -    - = ** UST AREA DECREASE.   12 USERS.
0906:06.37  222036= 22=<DI>=JDAC=   LINE SET FOR RECONNECT
0906:06.52  222107= 71=    -    - = ** UST AREA INCREASE.   13 USERS.
0906:10.12  222287= 32=<DI>=NEW = IOGON    MASB
0906:21.18  224085= 58=<DI>=CONN=   CONN COMMAND TO =DI=
0906:21.43  224170= 12=<DI>=TERM= IOGOFF   MASB           ON AT  906.06
0909:03.30  254008= 12=<TK>=TERM= IOGOFF   PRC118        ON AT  854.03
0912:43.40  272875= 12=<DI>=TERM= IOGOFF   MASB           ON AT  814.41
0913:35.14  277724= 32=<DJ>=NEW = IOGON    HIS
0914:45.30  289489= 12=<TN>=TERM= IOGOFF   CSC7592        ON AT  850.30
0921:31.55  319884= 12=<T2>=TERM= IOGOFF   PRC101        ON AT  845.21
0923:01.71  325571= 32=<TA>=NEW = IOGON    PRC101
0927:02.73  373064= 32=<TK>=NEW = IOGON    HIS
0928:41.54  393158= 12=<TK>=TERM= IOGOFF   HIS            ON AT  927.03
0929:30.73  409020= 32=<2X>=NEW = IOGON    IRCNT1
0929:41.40  412835= 22=<TC>=EDTX=   LINE SET FOR RECONNECT
```


Figure 15-2. Timesharing Reduction Event Log Report

Elapsed-time driven reports:

- o Response Time For All Users (figure 15-3) (ID #4)
- o CPU Times Allocated To TS1-EXEC and TS1-SUBDISPATCH (figure 15-4) (no ID #) (subordinate to the above report)
- o Response Time For Users Not Requesting More Core (figure 15-5) (ID #6)
- o Response Time For Users With Core Request During Line Idle (figure 15-6) (ID #5)
- o Total Time In Subdispatch Queue (figure 15-7) (ID #10)
- o Time In Subdispatch Queue Waiting Service (figure 15-8) (ID #11)
- o Processor Time In Subdispatch (figure 15-9) (ID #12).

Summary reports:

- o TSS Subtraces Encountered (figure 15-10 parts 1 and 2) (ID #1 & #2)
- o Derails Executed By Users (figure 15-11 parts 1 and 2) (ID #7 & #8).

The reports that the Emulation pass produces are:

Trace driven reports:

- o Exception Message Report (Tabular) (figure 15-12) (no ID #) and
- o Users Map (Tabular) (figure 15-13) (ID #3) (subordinate to above report)
- o TSS Delay & User Delay Report (Tabular) (figure 15-14) (no ID #) and
- o TSS Core Map (Tabular) (figure 15-15) (no ID #) (subordinate to above report)
- o Error Message Report (Tabular) (figure 15-16) (no ID #).

Elapsed-time driven reports:

- o Intertrace Gap - Duration (Tabular and Logplot) (figure 15-17) (ID #13)
- o Memory Activity (Tabular) (figure 15-18) (no ID #) (subordinate to above report)
- o Eligibles For Subdispatch - Duration (Tabular and Logplot) (figure 15-19) (ID #14)
- o In Subdispatch - Duration (Tabular and Logplot) (figure 15-20) (ID #15)
- o Eligible For & In Subdispatch - Duration (Tabular and Logplot) (figure 15-21) (ID #16)
- o User Swaps - Swap Rate (Logplot) (figure 15-22) (ID #17)
- o User Swaps - Swap Amount (Tabular and Logplot) (figure 15-23) (ID #18)
- o User Swaps - Duration (Tabular and Logplot) (figure 15-24) (ID #19)
- o User I/Os - Duration (Tabular and Logplot) (figure 15-25) (ID #20)
- o In System Master Catalog Wait - Duration (Tabular and Logplot) (figure 15-26) (ID #21).

Summary reports:

- o Intertrace Time (Adjusted for TSS GEWAKES) (Histogram) (figure 15-27) (ID #9).

...S REDUCTION, VERSION 7.2 - 1.000, 30 SEPTEMBER 82

DISTRIBUTION COLLECTED ON SYSTEM DSCC1 AT 0818:43.515 FRI 82-05-21. INITIAL GMC TAPE #1469

RESPONSE TIME FOR ALL USERS PLOT- 1

-= MIN /= AVG += MAX

USERS	TIME	WAITS IN INTERVAL (SEC) <1	<2	<4	<8	<16	<30	>30	LINE
4	819:03.52	1	.	.	.	3			<DI>
4	819:23.52	6	.	.	.	.	2		<DI>
4	819:43.52	7	.	.	.	.	.	2	<DI>
4	820:03.52	6	.	.	.	.	.	2	<DI>
4	820:23.52	8	.	.	.	.	.	2	<DI>
4	820:43.52	6	.	.	.	.	.	2	<DI>
4	821:03.52	10	.	.	.	.	.	2	<DI>
4	821:23.52	16	.	.	.	.	.	1	<DI>
4	821:43.52	12	.	.	.	.	.	1	<DI>
4	822:03.52	10	1	.	.	.	.	1	<DI>
4	822:23.52	5	.	.	.	.	.	1	<DI>
4	822:43.52	2	.	.	.	.	.	1	<DI>
4	823:03.52	5	.	.	.	.	.	1	<DI>
5	823:23.52	8	1	.	.	.	.	1	<DI>
5	823:43.52	16	.	.	.	.	.	1	<DI>
5	824:03.52	7	.	.	.	.	.	1	<DI>
6	824:23.52	6	.	.	.	.	.	1	<DI>
6	824:43.52	7	1	.	.	.	.	1	<DI>
6	825:03.52	5	1	.	.	.	.	1	<DI>
6	825:23.52	11	.	.	.	.	.	1	<DI>
6	825:43.52	8	.	.	.	.	.	1	<DI>
6	826:03.52	9	1	.	.	.	.	1	<DI>
6	826:23.52	9	.	.	.	.	.	1	<DI>
6	826:43.52	7	.	.	.	.	.	1	<DI>
7	827:03.52	23	1	.	.	.	.	1	<DI>
7	827:23.52	10	.	.	.	.	.	1	<DI>
7	827:43.52	9	.	.	.	.	.	1	<DI>
7	828:03.52	4	.	.	.	.	.	1	<DI>
7	828:23.52	8	.	.	.	.	.	1	<DI>
7	828:43.52	9	.	.	.	.	.	1	<DI>
8	829:03.52	16	.	.	.	.	.	1	<DI>
8	829:23.52	11	.	.	.	.	.	1	<DI>
8	829:43.52	14	1	.	.	.	.	1	<DI>
8	830:03.52	13	.	.	.	.	.	1	<DI>
8	830:23.52	10	.	.	.	.	.	1	<DI>
8	830:43.52	8	1	.	.	.	.	1	<DI>
8	831:03.52	13	.	.	.	.	.	1	<DI>
8	831:23.52	18	.	.	.	.	.	1	<DI>
8	831:43.52	14	.	1	.	.	.	1	<DI>
8	832:03.52	20	.	.	.	.	.	1	<DI>
8	832:23.52	18	.	.	.	.	.	1	<DI>
8	832:43.52	17	.	.	.	.	.	1	<DI>
8	833:03.52	26	.	.	.	.	.	1	<DI>

WAIT (SECONDS)

1.0E 00 1.0E 01 1.0E 02

CURVE	MINIMUM	MAXIMUM
MIN	-2.6488E-01	4.3092E-01
AVG	2.1931E-02	1.2712E 02
MAX	1.0316E-01	1.0095E 03

Figure 15-3. Response Time For All Users Report

TSS REDUCTION, VERSION 7.2 - 1.000, 30 SEPTEMBER 82

DISTRIBUTION COLLECTED ON SYSTEM DSCC1 AT 0818:43.515 FRI 82-05-21. INITIAL GMC TAPE #1469

CPU TIMES ALLOCATED TO TS1-EXEC AND TS1-SUBDISPATCH REPORT- 1

PLOT TIME REFERENCE	FIRST T/O TIME	LAST T/O TIME	NUMBER OF T/OS	**** CPU SECONDS ATTRIBUTABLE TO THE FOLLOWING **** TS1-EXEC	TS1-SDSP	IDLE-TIME	OTHER-CPU	TOTAL	
0819:03.5	0818:43.5	0818:43.5	1	0.25	0.04	0.	0.01	0.30	----I
0819:23.5	0818:43.5	0819:17.0	2	0.73	0.04	65.73	0.70	67.21	I++++++
0819:43.5	0819:17.0	0819:23.7	1	0.40	0.01	12.58	0.34	13.33	------I
0820:03.5	0819:23.7	0819:48.3	1	0.35	0.02	48.37	0.38	49.12	I++
0820:23.5	0819:48.3	0820:12.7	1	0.36	0.02	48.27	0.36	49.01	I++
0820:43.5	0820:12.7	0820:30.7	1	0.19	0.01	35.33	0.47	35.99	-I
0821:03.5	0820:30.7	0820:55.5	1	0.21	0.01	48.92	0.42	49.56	I++
0821:23.5	0820:55.5	0821:19.2	2	0.74	0.24	45.72	0.66	47.35	I+
0821:43.5	0821:19.2	0821:43.1	2	0.65	0.09	46.53	0.66	47.93	I+
0822:03.5	0821:43.1	0822:03.2	11	0.99	0.07	32.74	6.52	40.33	I
0822:23.5	0822:03.2	0822:12.1	4	0.46	0.02	13.86	3.55	17.89	-----I
0822:43.5	0822:12.1	0822:38.1	1	0.34	0.02	50.69	0.38	51.43	I+++
0823:03.5									
0823:23.5	0822:38.1	0823:21.4	4	0.73	0.21	84.80	1.47	87.22	I**********
0823:43.5	0823:21.4	0823:42.1	3	1.17	0.24	39.09	1.07	41.57	I
0824:03.5	0823:42.1	0824:02.6	3	0.38	0.02	39.33	1.17	40.88	I
0824:23.5	0824:02.6	0824:07.0	1	0.16	0.00	8.28	0.36	8.81	-------I
0824:43.5	0824:07.0	0824:32.1	2	0.71	0.67	48.24	0.62	50.24	I++
0825:03.5	0824:32.1	0825:02.9	3	1.11	1.05	58.51	0.96	61.63	I+++++
0825:23.5	0825:02.9	0825:20.1	2	0.55	0.45	32.88	0.66	34.54	-I
0825:43.5	0825:20.1	0825:39.4	4	0.36	0.04	36.88	1.42	38.70	I
0826:03.5	0825:39.4	0826:01.4	8	1.74	5.53	34.47	2.30	44.04	I
0826:23.5	0826:01.4	0826:20.8	4	0.34	0.21	36.87	1.41	38.83	I
0826:43.5	0826:20.8	0826:42.2	6	1.01	2.52	37.46	1.87	42.86	I
0827:03.5	0826:42.2	0827:00.4	6	2.27	3.32	27.20	1.69	34.48	I
0827:23.5	0827:00.4	0827:14.8	2	0.66	0.94	28.53	0.61	30.75	--I
0827:43.5	0827:14.8	0827:40.4	1	0.24	0.01	50.59	0.41	51.26	I++
0828:03.5									
0828:23.5	0827:40.4	0828:06.5	1	0.25	0.19	51.47	0.38	52.29	I+++
0828:43.5	0828:06.5	0828:39.1	4	0.63	0.27	62.86	1.47	65.23	I++++++
0829:03.5	0828:39.1	0829:01.1	3	0.83	0.07	42.11	1.04	44.05	I
0829:23.5	0829:01.1	0829:23.1	2	0.81	0.05	42.59	0.65	44.08	I+
0829:43.5	0829:23.1	0829:40.8	5	1.19	0.27	31.66	2.23	35.35	-I
0830:03.5	0829:40.8	0830:02.4	15	0.83	0.03	32.85	9.81	43.52	I
0830:23.5	0830:02.4	0830:23.1	3	0.76	0.03	39.29	1.35	41.42	I
0830:43.5	0830:23.1	0830:41.2	5	0.73	0.96	32.05	2.43	36.17	I
0831:03.5	0830:41.2	0831:03.3	10	1.16	0.05	37.82	5.30	44.32	I+
0831:23.5	0831:03.3	0831:21.0	9	1.54	0.30	27.22	6.55	35.61	-I
0831:43.5	0831:21.0	0831:44.9	3	0.56	0.03	45.74	1.40	47.73	I+
0832:03.5	0831:44.9	0832:01.5	3	0.87	0.29	31.10	1.11	33.37	-I

Figure 15-4. CPU Times Allocated to Exec and Subdispatch Report

TSS REDUCTION, VERSION 7.2 - 1.000, 30 SEPTEMBER 82

DISTRIBUTION COLLECTED ON SYSTEM DSCC1 AT 0818:43.515 FRI 82-05-21. INITIAL GMC TAPE #1469

RESPONSE TIME FOR USERS NOT REQUESTING MORE CORE PLOT- 2

-= MIN /= AVG += MAX

USERS	TIME	WAITS IN INTERVAL (SEC) <1	<2	<4	<8	<16	<30	>30	LINE	WAIT (SECONDS) 1.0E 00	1.0E 01	1.0E 02
4	819:03.52	.	.	.	.	3			<DI>	X3333333333333333333333		
4	819:23.52	1	.	.	.	.	2		<DI>	*////////////////////////+++++		
4	819:43.52	4	.	.	.	.	.	2	<DI>	*////////////////////////++++++++++		
4	820:03.52	3	.	.	.	.	.	2	<DI>	*////////////////////////+++++++++		
4	820:23.52	6	.	.	.	.	.	2	<DI>	*////////////////////////++++++++++++		
4	820:43.52	5	.	.	.	.	.	2	<DI>	*////////////////////////+++++++++++		
4	821:03.52	7	.	.	.	.	.	2	<DI>	*////////////////////////++++++++++++++		
4	821:23.52	9	.	.	.	.	.	1	<DI>	*////////////////////////++++++++++++++++++		
4	821:43.52	10	.	.	.	.	.	1	<DI>	*////////////////////////++++++++++++++++++++		
4	822:03.52	8	.	.	.	.	.	1	<DI>	*////////////////////////++++++++++++++++++		
4	822:23.52	3	.	.	.	.	.	1	<DI>	*////////////////////////+++++++++++		
4	822:43.52	2	.	.	.	.	.	1	<DI>	*////////////////////////++++++++++		
4	823:03.52	4	.	.	.	.	.	1	<DI>	*////////////////////////+++++++++++++		
5	823:23.52	3	.	.	.	.	.	1	<DI>	*////////////////////////++++++++++		
5	823:43.52	7	.	.	.	.	.	1	<DI>	*////////////////////////++++++++++++++++		
5	824:03.52	4	.	.	.	.	.	1	<DI>	*////////////////////////+++++++++++++		
6	824:23.52	3	.	.	.	.	.	1	<DI>	*////////////////////////++++++++++		
6	824:43.52	5	.	.	.	.	.	1	<DI>	*////////////////////////+++++++++++++		
6	825:03.52	4	.	.	.	.	.	1	<DI>	*////////////////////////++++++++++++		
6	825:23.52	8	.	.	.	.	.	1	<DI>	*////////////////////////++++++++++++++++		
6	825:43.52	7	.	.	.	.	.	1	<DI>	*////////////////////////+++++++++++++++		
6	826:03.52	8	.	.	.	.	.	1	<DI>	*////////////////////////++++++++++++++++		
6	826:23.52	7	.	.	.	.	.	1	<DI>	*////////////////////////+++++++++++++++		
6	826:43.52	5	.	.	.	.	.	1	<DI>	*////////////////////////++++++++++++++		
7	827:03.52	12	.	.	.	.	.	1	<DI>	*////////////////////////+++++++++++++++++		
7	827:23.52	8	.	.	.	.	.	1	<DI>	*////////////////////////++++++++++++++++		
7	827:43.52	9	.	.	.	.	.	1	<DI>	*////////////////////////++++++++++++++++		
7	828:03.52	4	.	.	.	.	.	1	<DI>	*////////////////////////+++++++++++++		
7	828:23.52	6	.	.	.	.	.	1	<DI>	*////////////////////////++++++++++++++		
7	828:43.52	7	.	.	.	.	.	1	<DI>	*////////////////////////++++++++++++++++		
8	829:03.52	11	.	.	.	.	.	1	<DI>	*////////////////////////+++++++++++++++++++		
8	829:23.52	6	.	.	.	.	.	1	<DI>	*////////////////////////++++++++++++++		
8	829:43.52	11	.	.	.	.	.	1	<DI>	*////////////////////////+++++++++++++++++++		
8	830:03.52	10	.	.	.	.	.	1	<DI>	*////////////////////////++++++++++++++++++		
8	830:23.52	5	.	.	.	.	.	1	<DI>	*////////////////////////+++++++++++++		
8	830:43.52	5	.	.	.	.	.	1	<DI>	*////////////////////////++++++++++++++		
8	831:03.52	6	.	.	.	.	.	1	<DI>	*////////////////////////++++++++++++++		
8	831:23.52	10	.	.	.	.	.	1	<DI>	*////////////////////////+++++++++++++++++++		
8	831:43.52	9	.	.	.	.	.	1	<DI>	*////////////////////////++++++++++++++++++		
8	832:03.52	11	.	.	.	.	.	1	<DI>	*////////////////////////+++++++++++++++++++		
8	832:23.52	13	.	.	.	.	.	1	<DI>	*////////////////////////++++++++++++++++++		
8	832:43.52	9	.	.	.	.	.	1	<DI>	*////////////////////////++++++++++++++++++		
8	833:03.52	19	.	.	.	.	.	1	<DI>	*////////////////////////++++++++++++++++++++++++		

CURVE	MINIMUM	MAXIMUM
MIN	-2.6488E-01	9.5021E 00
AVG	5.6924E-03	2.9651E 02
MAX	1.5157E-02	1.0095E 03

Figure 15-5. Response Time For Users Not Requesting More Core Plot

TSS REDUCTION, VERSION 7.2 - 1.000, 30 SEPTEMBER 82

DISTRIBUTION COLLECTED ON SYSTEM DSCC1 AT 0818:43.515 FRI 82-05-21. INITIAL GMC TAPE #1469

RESPONSE TIME FOR USERS WITH CORE REQUEST DURING LINE IDLE PLOT- 3

-= MIN /= AVG += MAX

USERS	TIME	WAITS IN INTERVAL (SEC) <1	<2	<4	<8	<16	<30	>30	LINE	WAIT (SECONDS) 1.0E 00	1.0E 01	1.0E 02
4	819:03.52	1							<TN>	*		
4	819:23.52	5							<TN>	*		
4	819:43.52	3							<TN>	*		
4	820:03.52	3							<TN>	*		
4	820:23.52	2							<TN>	*		
4	820:43.52	1							<TN>	*		
4	821:03.52	3							<TN>	*		
4	821:23.52	7							<TC>	*		
4	821:43.52	2							<TN>	*		
4	822:03.52	2	1						<TN>	*++		
4	822:23.52	2							<TN>	*		
4	822:43.52									I		
4	823:03.52	1							<TN>	*		
5	823:23.52	5	1						<TN>	*+		
5	823:43.52	9							<TC>	*		
5	824:03.52	3							<TC>	*		
6	824:23.52	3							<TK>	*		
6	824:43.52	2	1						<TK>	*++		
6	825:03.52	1	1						<TC>	*		
6	825:23.52	3							<TC>	*		
6	825:43.52	1							<TA>	*		
6	826:03.52	1	1						<TN>	*+		
6	826:23.52	2							<TN>	*		
6	826:43.52	2							<TN>	*		
7	827:03.52	11	1						<TN>	*++++		
7	827:23.52	2							<TK>	*		
7	827:43.52									I		
7	828:03.52									I		
7	828:23.52	2							<TK>	*		
7	828:43.52	2							<TK>	*		
8	829:03.52	5							<VV>	*		
8	829:23.52	5							<TK>	*		
8	829:43.52	3	1						<VV>	*		
8	830:03.52	3							<VV>	*		
8	830:23.52	5							<VV>	*		
8	830:43.52	3	1						<TK>	*+		
8	831:03.52	7							<CG>	*		
8	831:23.52	8							<TN>	*		
8	831:43.52	5	.	1					<TN>	*+++++++++		
8	832:03.52	9							<TN>	*		
8	832:23.52	5							<TC>	*		
8	832:43.52	8							<TK>	*		
8	833:03.52	7							<TK>	*		

CURVE	MINIMUM	MAXIMUM
MIN	-7.3665E-02	6.1448E-01
AVG	6.4500E-02	4.5227E 00
MAX	6.4500E-02	3.8571E 01

Figure 15-6. Response Time For Users With Core Requests During Line Idle Plot

TSS REDUCTION, VERSION 7.2 - 1.000, 30 SEPTEMBER 82

DISTRIBUTION COLLECTED ON SYSTEM DSCC1 AT 0818:43.515 FRI 82-05-21. INITIAL GMC TAPE #1469

TOTAL TIME IN SUBDISPATCH QUEUE PLOT- 4

-= MIN /= AVG += MAX

USERS	TIME	QUEUE STATISTICS #SD	CPU	TROUTS	I/O-BSY	LINE	Q-TIME (MILLISEC) 1.0E-01 1.0E 00 1.0E 01 1.0E 02 1.0E 03 1.0E 04
4	819:03.52	9	19.31%	0.00%	0.00%	<TN>	X----------///+++
4	819:23.52	64	8.03%	0.00%	37.50%	<CV>	I2222222222222222+
4	819:43.52	21	10.21%	0.00%	28.57%	<CV>	I2222222222222222++
4	820:03.52	63	7.74%	0.00%	76.19%	<CV>	I2222222222222222+++
4	820:23.52	15	10.35%	0.00%	26.67%	<TN>	I2222222222222222+
4	820:43.52	6	14.73%	0.00%	33.33%	<TN>	I2222222222222222++
4	821:03.52	18	15.37%	0.00%	33.33%	<TN>	I----------/++
4	821:23.52	96	29.50%	7.29%	28.13%	<TN>	I----------/////////+++++++++++++++++
4	821:43.52	48	28.35%	0.00%	14.58%	<TN>	I----------/+++++
4	822:03.52	103	11.42%	0.00%	21.36%	<TN>	I----------////++
4	822:23.52	51	13.32%	0.00%	5.88%	<TN>	X----------///++
4	822:43.52	24	6.38%	0.00%	95.83%	<CV>	I----------////+
4	823:03.52	7	6.75%	0.00%	14.29%	<TN>	I----------2222
5	823:23.52	80	44.57%	0.00%	0.00%	<TN>	I----------/////+++
5	823:43.52	83	42.55%	9.64%	22.89%	<TC>	I----------/////++++++++
5	824:03.52	26	12.74%	0.00%	50.00%	<TC>	I----------/////+++
6	824:23.52	34	6.48%	0.00%	26.47%	<TK>	I2222222222222222+
6	824:43.52	116	71.73%	33.62%	32.76%	<TK>	I----------////////++++++
6	825:03.52	113	77.05%	59.82%	11.50%	<TK>	I----------////////+++++
6	825:23.52	50	27.25%	4.00%	72.00%	<TK>	I----------/+++++++
6	825:43.52	55	16.59%	0.00%	83.64%	<TA>	X----------///++++
6	826:03.52	348	85.86%	75.00%	6.03%	<TN>	I----------//////////++++
6	826:23.52	31	16.03%	0.00%	74.19%	<TA>	I----------////++
6	826:43.52	219	84.74%	70.32%	0.91%	<TK>	I----------//////////+++++
7	827:03.52	305	71.99%	31.80%	32.79%	<TN>	I----------////////++++++
7	827:23.52	90	78.76%	46.67%	27.78%	<TK>	I----------/////////+++
7	827:43.52	8	7.02%	0.00%	12.50%	<TK>	I----------///++
7	828:03.52	0	0.%	0.%	0.%		I
7	828:23.52	44	62.35%	20.45%	22.73%	<TK>	I----------///////++++
7	828:43.52	42	61.57%	16.67%	50.00%	<TK>	I----------///////++++
8	829:03.52	99	14.53%	0.00%	58.59%	<TA>	X----------////+++
8	829:23.52	96	38.64%	0.00%	67.71%	<TN>	I----------/////++++++
8	829:43.52	104	22.69%	0.00%	22.12%	<TN>	I----------/////+++++++
8	830:03.52	58	9.26%	0.00%	6.90%	<TN>	I----------////++
8	830:23.52	62	28.55%	4.84%	16.13%	<TK>	I----------/////++++++++
8	830:43.52	80	77.83%	43.75%	1.25%	<TK>	I----------//////////+++++
8	831:03.52	109	9.48%	0.00%	15.60%	<CG>	I----------////++
8	831:23.52	166	31.30%	3.61%	19.88%	<TK>	I----------/////+++++++
8	831:43.52	41	14.28%	0.00%	34.15%	<TA>	I----------/////+++
8	832:03.52	130	43.24%	1.54%	43.08%	<TC>	I----------/////++++++
8	832:23.52	39	8.15%	0.00%	28.21%	<TC>	X----------/////++++
8	832:43.52	116	76.34%	41.38%	20.69%	<TN>	I----------/////////+++++
8	833:03.52	82	23.79%	0.00%	24.39%	<TC>	I----------/////++++

CURVE	MINIMUM	MAXIMUM
MIN	1.1250E 00	4.0980E 00
AVG	3.2497E 00	8.3807E 01
MAX	4.1530E 00	9.2997E 03

Figure 15-7. Total Time In Subdispatch Queue Plot

TSS REDUCTION, VERSION 7.2 - 1.000, 30 SEPTEMBER 82

DISTRIBUTION COLLECTED ON SYSTEM DSCC1 AT 0818:43.515 FRI 82-05-21. INITIAL GMC TAPE #1469

TIME IN SUBDISPATCH QUEUE WAITING SERVICE PLOT- 5

-= MIN /= AVG += MAX

USERS	TIME	QUEUE STATISTICS #SD	CPU	TROUTS	I/O-BSY	LINE	WAIT (MILLISEC) 1.0E-01 1.0E 00 1.0E 01 1.0E 02 1.0E 03 1.0E 04
4	819:03.52	9	19.31%	0.00%	0.00%	<TN>	X----------////+
4	819:23.52	64	8.03%	0.00%	37.50%	<TN>	I----------22
4	819:43.52	21	10.21%	0.00%	28.57%	<CV>	I----------//+
4	820:03.52	63	7.74%	0.00%	76.19%	<CV>	I----------//+++
4	820:23.52	15	10.35%	0.00%	26.67%	<TN>	I----------22
4	820:43.52	6	14.73%	0.00%	33.33%	<TN>	I----------22
4	821:03.52	18	15.37%	0.00%	33.33%	<TN>	I----------//+
4	821:23.52	96	29.50%	7.29%	28.13%	<TN>	I----------/////+++++++++++++++++++
4	821:43.52	48	28.35%	0.00%	14.58%	<TC>	I----------+
4	822:03.52	103	11.42%	0.00%	21.36%	<TN>	I----------////++
4	822:23.52	51	13.32%	0.00%	5.88%	<TN>	X----------//++
4	822:43.52	24	6.38%	0.00%	95.83%	<CV>	I----------////+
4	823:03.52	7	6.75%	0.00%	14.29%	<TN>	I----------////+
5	823:23.52	80	44.57%	0.00%	0.00%	<TN>	I----------////++
5	823:43.52	83	42.55%	9.64%	22.89%	<TC>	I----------////++++++++
5	824:03.52	26	12.74%	0.00%	50.00%	<TA>	I----------////+
6	824:23.52	34	6.48%	0.00%	26.47%	<TK>	I----------2
6	824:43.52	116	71.73%	33.62%	32.76%	<TA>	I----------///+++++
6	825:03.52	113	77.05%	39.82%	11.50%	<TK>	I----------///+++
6	825:23.52	50	27.25%	4.00%	72.00%	<TK>	I----------//+++
6	825:43.52	55	16.59%	0.00%	83.64%	<TA>	X----------///++++
6	826:03.52	348	85.86%	75.00%	6.03%	<TN>	I----------//+++++++
6	826:23.52	31	16.03%	0.00%	74.19%	<CV>	I----------////++
6	826:43.52	219	84.74%	70.32%	0.91%	<TN>	I----------//++++++
7	827:03.52	305	71.99%	31.80%	32.79%	<TN>	I----------///+++++
7	827:23.52	90	78.76%	46.67%	27.78%	<TK>	I----------///+++
7	827:43.52	8	7.02%	0.00%	12.50%	<TK>	I----------////+
7	828:03.52	0	0. %	0. %	0. %		I
7	828:23.52	44	62.35%	20.45%	22.73%	<TK>	I----------////+
7	828:43.52	42	61.57%	16.67%	50.00%	<TK>	I----------///+++++
8	829:03.52	99	14.53%	0.00%	58.59%	<CG>	X----------///++++
8	829:23.52	96	38.64%	0.00%	67.71%	<TN>	I----------////++
8	9:43.52	104	22.69%	0.00%	22.12%	<VV>	I----------////++
8	0:03.52	58	9.26%	0.00%	6.90%	<TN>	I----------////++
8	830:23.52	62	28.55%	4.84%	16.13%	<TK>	I----------////+
8	830:43.52	80	77.83%	43.75%	1.25%	<TK>	I----------////+++++++
8	831:03.52	109	9.48%	0.00%	15.60%	<CG>	I----------////++
8	831:23.52	166	31.30%	3.61%	19.88%	<TN>	I----------////+++++
8	831:43.52	41	14.28%	0.00%	34.15%	<TN>	I----------////++
8	832:03.52	130	43.24%	1.54%	43.08%	<TA>	I----------////++
8	832:23.52	39	8.15%	0.00%	28.21%	<TC>	X----------////++++
8	832:43.52	116	76.34%	41.38%	20.69%	<TN>	I----------////+++
8	833:03.52	82	23.79%	0.00%	24.39%	<TK>	I----------///++++

CURVE	MINIMUM	MAXIMUM
MIN	9.5400E-01	3.9000E 00
AVG	2.7024E 00	8.0279E 01
MAX	3.8150E 00	9.2758E 03

Figure 15-8. Total Time In Subdispatch Queue Waiting Service Plot

TSS REDUCTION, VERSION 7.2 - 1.000, 30 SEPTEMBER 82

DISTRIBUTION COLLECTED ON SYSTEM DSCC1 AT 0818:43.515 FRI 82-05-21. INITIAL GMC TAPE #1469

PROCESSOR TIME IN SUBDISPATCH PLOT- 6

-= MIN /= AVG += MAX

USERS	TIME	QUEUE STATISTICS #SD	CPU	TROUTS	I/O-BSY	LINE	PROCTM (MILLISEC) 1.0E-01	1.0E 00	1.0E 01
4	819:03.52	9	19.31%	0.00%	0.00%	<TN>	X-----/////////////+++++++++++++++		
4	819:23.52	64	8.03%	0.00%	37.50%	<CV>	I-----/////+++++++++++++++++++		
4	819:43.52	21	10.21%	0.00%	28.57%	<TN>	I-----/////+++++++++++++++++++++		
4	820:03.52	63	7.74%	0.00%	76.19%	<TN>	I-----/////+++++++++++++++++++++++++++		
4	820:23.52	15	10.35%	0.00%	26.67%	<TN>	I------/////+++++++++++++++		
4	820:43.52	6	14.73%	0.00%	33.33%	<TN>	I------/////////++++++++++++		
4	821:03.52	18	15.37%	0.00%	33.33%	<TN>	I------/////////+++++++++++++++++++		
4	821:23.52	96	29.50%	7.29%	28.13%	<TC>	I-----/////////////////////+++++++++++++++++++++++++		
4	821:43.52	48	28.35%	0.00%	14.58%	<TN>	I-----/////////////////////+++++++++++++++++++++		
4	822:03.52	103	11.42%	0.00%	21.36%	<TN>	I-----/////////++++		
4	822:23.52	51	13.32%	0.00%	5.88%	<TN>	X-----/////////+++++++++++++++++		
4	822:43.52	24	6.38%	0.00%	95.83%	<CV>	I-----///++++++++++		
4	823:03.52	7	6.75%	0.00%	14.29%	<TN>	I-----///++++++		
5	823:23.52	80	44.57%	0.00%	0.00%	<TN>	I-----/////////////////////////+++++++++		
5	823:43.52	83	42.55%	9.64%	22.89%	<TC>	I-----/////////////////////////+++++++++++++++++++++		
5	824:03.52	26	12.74%	0.00%	50.00%	<TC>	I------/////////+++++++++++++++++++		
6	824:23.52	34	6.48%	0.00%	26.47%	<TK>	I-----///++++++++++		
6	824:43.52	116	71.73%	33.62%	32.76%	<TK>	I-----/////////////////////////////////+++++++++++++++		
6	825:03.52	113	77.05%	39.82%	11.50%	<TK>	I-----/////////////////////////////////+++++++++++++++		
6	825:23.52	50	27.25%	4.00%	72.00%	<TK>	I-----/////////////////////////+++++++++++++++++++++++		
6	825:43.52	55	16.59%	0.00%	83.64%	<TA>	X-----/////////+++++++++		
6	826:03.52	348	85.86%	75.00%	6.03%	<TN>	I-----///+++++++++		
6	826:23.52	31	16.03%	0.00%	74.19%	<TA>	I-----/////////+++++++++		
6	826:43.52	219	84.74%	70.32%	0.91%	<TK>	I-----///+++++++++		
7	827:03.52	305	71.99%	31.80%	32.79%	<TN>	I-----///+++++++++++++++		
7	827:23.52	90	78.76%	46.67%	27.78%	<TK>	I-----///++++++++		
7	827:43.52	8	7.02%	0.00%	12.50%	<TK>	I-----///++++		
7	828:03.52	0	0. %	0. %	0. %		I		
7	828:23.52	44	62.35%	20.45%	22.73%	<TK>	I------/////////////////////////////////+++++++++++++++		
7	828:43.52	42	61.57%	16.67%	50.00%	<TK>	I-----/////////////////////////////////+++++++++++++++		
8	829:03.52	99	14.53%	0.00%	58.59%	<VV>	X-----/////////////+++++++++++++++		
8	829:23.52	96	38.64%	0.00%	67.71%	<TN>	I-----/////////////////////+++++++++++++++++++++		
8	829:43.52	104	22.69%	0.00%	22.12%	<TN>	I-----/////////////////////+++++++++++++++++++++++++++		
8	830:03.52	58	9.26%	0.00%	6.90%	<TN>	I------/////////++++++++++++++++++++++		
8	830:23.52	62	28.55%	4.84%	16.13%	<TK>	I-----/////////////////////////////+++++++++++++++++++++++++++		
8	830:43.52	80	77.83%	43.75%	1.25%	<TK>	I-----///+++++++++		
8	831:03.52	109	9.48%	0.00%	15.60%	<TN>	I-----/////////++++++		
8	831:23.52	166	31.30%	3.61%	19.88%	<CV>	I-----/////////////////////+++++++++++++++++++++++		
8	831:43.52	41	14.26%	0.00%	34.15%	<TA>	I-----/////////+++++++++++++++		
8	832:03.52	130	43.24%	1.54%	43.08%	<TC>	I-----/////////////////////+++++++++++++++++++++		
8	832:23.52	39	8.15%	0.00%	28.21%	<CV>	X-----///++++++++++++++++++++++++		
8	832:43.52	116	76.34%	41.38%	20.69%	<TN>	I-----/////////////////////////////////////+++++++++++++		
8	833:03.52	82	23.79%	0.00%	24.39%	<TC>	I-----/////////////////////++++++++++++++ ++++++		

CURVE	MINIMUM	MAXIMUM
MIN	1.5600E-01	1.8700E-01
AVG	2.2800E-01	1.6408E 01
MAX	3.4300E-01	4.0156E 01

Figure 15-9. Processor Time In Subdispatch Plot

TSS REDUCTION, VERSION 7.2 - 1.000, 30 SEPTEMBER 82

DISTRIBUTION COLLECTED ON SYSTEM DSCC1 AT 0818:43.515 FRI 82-05-21. INITIAL GMC TAPE #1469

TSS SUBTRACES ENCOUNTERED (IOW HALF)

INDIV. NUMBER	CUMUL. NUMBER	CUMUL. PRC	INDIV. PRC	TSS SUBTRC	PERCENT OF OCCURRENCE 00 05 10 15 20 25 30 35 40 45 50	HISTOGRAM 1
					I....I....I....I....I....I....I....I....I....I....I	
70	70	0.006	0.006	1- 1	I	LOG-ON COMPLETE
0	70	0.006	0.	2- 2	I	SWAP SPACE DENIED
187	257	0.024	0.017	3- 3	I	PAT DENIAL
719	976	0.089	0.066	4- 4	I	GEFSYE DEALLOCATE
719	1695	0.155	0.066	5- 5	I	DEALLOC COMPLETE
7	1702	0.156	0.001	6- 6	I	ERROR MESSAGE
1564	3266	0.299	0.143	7- 7	I	POPUP PRIMITIVE
2207	5473	0.500	0.202	8- 8	I	CALIP PRIMITIVE
928	6401	0.585	0.085	9- 9	I	ENTER BUILD MODE
1801	8202	0.750	0.165	10- 10	I	EXEC PRIMITIVE
117	8319	0.761	0.011	11- 11	I	SYSTM PRIMITIVE
53	8372	0.766	0.005	12- 12	I	LOG-OFF
1237	9609	0.879	0.113	13- 13	I	COMMAND RECEIVED
9705	19314	1.766	0.887	14- 14	IX	PERIODIC CHECK
0	19314	1.766	0.	15- 15	I	GEWAKE NO USERS
53	19367	1.771	0.005	16- 16	I	BREAK/DISCONNECT
5160	24527	2.243	0.472	17- 17	I	GEWAKE S/D DONE
156686	181213	16.570	14.327	18- 18	IXXXXXXXXXXXXXX	GEWAKE S/D BUSY
307	181520	16.598	0.028	19- 19	I	APB RIO C/C
1083	182603	16.697	0.099	20- 20	I	BUILD INPUT REC
58	182661	16.703	0.005	21- 21	I	SY** I/O COMPLETE
14	182675	16.704	0.001	22- 22	I	RECONNECT MODE
95117	277792	25.401	8.698	23- 23	IXXXXXXXXX	PROCESS DRI
31902	309694	28.319	2.917	24- 24	IXXX	REQUEST FILE I/O
29583	339277	31.024	2.705	25- 25	IXXX	DISK I/O COMPLETE
0	339277	31.024	0.	26- 26	I	DRI DEFIL DENIAL
1755	341032	31.184	0.160	27- 27	I	ISSUE MME GEFSYE
187	341219	31.201	0.017	28- 28	I	BAD FILACT PARAM
1755	342974	31.362	0.160	29- 29	I	GEFSYE/27 COMPLETE
52	343026	31.367	0.005	30- 30	I	DELAY 200 MS
36	343062	31.370	0.003	31- 31	I	DELAY 2 SEC
70	343132	31.376	0.006	32- 32	I	STORE USERID
20	343152	31.378	0.002	33- 33	I	PASUST -1 COMP
0	343152	31.378	0.	34- 34	I	DRI MORINK ERROR
91	343243	31.386	0.008	35- 35	I	SPAWN I/O COMPLETE
91	343334	31.395	0.008	36- 36	I	BATCH JOB
0	343334	31.395	0.	37- 37	I	DRI SPAWN ERROR
7	343341	31.395	0.001	38- 38	I	WRITE ABRT
4	343345	31.396	0.000	39- 39	I	ABRT I/O COMPLETE
91	343436	31.404	0.008	40- 40	I	EXEC+ - DRI RESTOR
91	343527	31.412	0.008	41- 41	I	IOAD DRI RESTOR
376	343903	31.447	0.034	42- 42	I	RESTOR IOAD COMPI
928	344831	31.532	0.085	43- 43	I	PERM I/O RESTOR
4	344835	31.532	0.000	44- 44	I	START LINE SWITCH
15	344850	31.533	0.001	45- 45	I	FILE GROW .MFS19
15	344865	31.535	0.001	46- 46	I	GROW COMPLETE
0	344865	31.535	0.	47- 47	I	CONSOIE I/O
34	344899	31.538	0.003	48- 48	I	DRI JOUT RETURN
0	344899	31.538	0.	49- 49	I	JOUT OUTPUT BUSY
2728	347627	31.787	0.249	50- 50	I	START GWAKE
0	347627	31.787	0.	51- 51	I	*J WRITE - TASK
0	347627	31.787	0.	52- 52	I	START DRI TASK
0	347627	31.787	0.	53- 53	I	*J READ - TASK
3	347630	31.788	0.000	54- 54	I	DRI SAVE COMPLETE

1093606 ENTRIES TOTAL 745976 OUT OF RANGE

Figure 15-10. TSS Subtraces Encountered Histogram
(Part 1 of 2)

TSS REDUCTION, VERSION 7.2 - 1.000, 30 SEPTEMBER 82

DISTRIBUTION COLLECTED ON SYSTEM DSCC1 AT 0818:43.515 FRI 82-05-21. INITIAL GMC TAPE #1469

TSS SUBTRACES ENCOUNTERED (HIGH HALF)

INDIV. NUMBER	CUMUL. NUMBER	CUMUL. PRC	INDIV. PRC	TSS SUBTRC	PERCENT OF OCCURRENCE 00 05 10 15 20 25 30 35 40 45 50	HISTOGRAM 2
					I....I....I....I....I....I....I....I....I....I....I	
347630	347630	31.787	31.787	1- 54	IXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX	DRI SAVE COMPLETE
15	347645	31.789	0.001	55- 55	I	IDS ATTRI .MFS13
15	347660	31.790	0.001	56- 56	I	IDS ATTRI STATUS
0	347660	31.790	0.	57- 57	I	DRI T.SYOT DENIAL
3	347663	31.791	0.000	58- 58	I	UST SWITCH T.CONN
0	347663	31.791	0.	59- 59	I	TAP* I/O COMPLETE
68443	416106	38.049	6.258	60- 60	IXXXXXX	ALLOCATOR SERVICES
259546	675652	61.782	23.733	61- 61	IXXXXXXXXXXXXXXXXXXXXXXXXX	ENTER PROC ALLOC
34996	710648	64.982	3.200	62- 62	IXXX	ENTER MEMORY ALLOC
1206	711854	65.092	0.110	63- 63	I	ENTER SWAP DECIS
9200	721054	65.934	0.841	64- 64	IX	NEXT PRIMITIVE IS
5	721059	65.934	0.000	65- 65	I	CONSIDER TSS INCR
3	721062	65.934	0.000	66- 66	I	INITIATE SIZE INCR
9	721071	65.935	0.001	67- 67	I	SET UP FENCE
1	721072	65.935	0.000	68- 68	I	MEMORY RESERVE
12	721084	65.936	0.001	69- 69	I	FORCE SWAP
0	721084	65.936	0.	70- 70	I	TERMINATE SWAP
10	721094	65.937	0.001	71- 71	I	TSSI UST AREA +1K
7	721101	65.938	0.001	72- 72	I	GEMORE MEMORY
5	721106	65.938	0.000	73- 73	I	GEMORE SUCCESSFUL
3	721109	65.939	0.000	74- 74	I	MEMORY RELEASE
137	721246	65.951	0.013	75- 75	I	GEMORE REFUSED
12	721258	65.952	0.001	76- 76	I	CANCEL CRUN
9709	730967	66.840	0.888	77- 77	IX	ATTEMPT MEM ALLOC
23932	754899	69.028	2.188	78- 78	IXX	MEM MAP CHANGE
7218	762117	69.688	0.660	79- 79	IX	SWAP PROGRAM
48418	810535	74.116	4.427	80- 80	IXXXX	TIME TYPES
2646	813181	74.358	0.242	81- 81	I	REQUEST EBM
896	814077	74.440	0.082	82- 82	I	SCAN COMMAND LIST
62	814139	74.445	0.006	83- 83	I	EBM REFUSED
3124	817263	74.731	0.286	84- 84	I	INPUT COMPLETE
632	817895	74.789	0.058	85- 85	I	POPUP TO CALLSS
17	817912	74.790	0.002	86- 86	I	APB COMPLETE
11195	829107	75.814	1.024	87- 87	IX	TERM I/O COMPLETE
97904	927011	84.766	8.952	88- 88	IXXXXXXXXX	S/D FAULT QUEUE
11909	938920	85.855	1.089	89- 89	IX	PROCESS .CRTSQ
5	938925	85.856	0.000	90- 90	I	TRACE ON/OFF
97905	1036830	94.808	8.952	91- 91	IXXXXXXXXX	MAKE S/D ENTRY
70	1036900	94.815	0.006	92- 92	I	LOG-ON REQUEST
0	1036900	94.815	0.	93- 93	I	BAD LINE STATUS
67	1036967	94.821	0.006	94- 94	I	CHECK VIP TERM TYP
3	1036970	94.821	0.000	95- 95	I	UST WAIT TIME
0	1036970	94.821	0.	96- 96	I	REJECT USER
73	1037043	94.828	0.007	97- 97	I	UST COMPRESSION
4	1037047	94.828	0.000	98- 98	I	TSSN UST AREA +1K
5	1037052	94.829	0.000	99- 99	I	UST AREA -1K
18549	1055601	96.525	1.696	100- 100	IXX	TERMINAL I/O
93	1055694	96.533	0.009	101- 101	I	CRUN $*$ FUNC
67	1055761	96.539	0.006	102- 102	I	REMOTE INQUIRY
268	1056029	96.564	0.025	103- 103	I	VIP INPUT COMPLETE
37577	1093606	100.000	3.436	104- 104	IXXX	DRI COMPLETE

1093606 ENTRIES TOTAL 0 OUT OF RANGE

Figure 15-10. (Part 2 of 2)

...ECTION, VERSION 7.2 - 1.000, 30 SEPTEMBER 82

DISTRIBUTION COLLECTED ON SYSTEM DSCC1 AT 0818:43.515 FRI 82-05-21. INITIAL GMC TAPE #1469

DERAILS (LOW HALF)

INDIV. NUMBER	CUMUL. NUMBER	CUMUL. PRC	INDIV. PRC	TSS DRI #		PERCENT OF OCCURRENCE 00 05 10 15 20 25 30 35 40 45 50	HISTOGRAM 3
11	11	0.012	0.012	-4-	-4	I	W.PSZ
210	221	0.232	0.221	-3-	-3	I	W.STOR
0	221	0.232	0.	-2-	-2	I	W.ERCH
208	429	0.451	0.219	-1-	-1	I	W.USER
0	429	0.451	0.	0-	0	I	ZERO
28939	29368	30.876	30.425	1-	1	IXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX	DIO
12469	41837	43.985	13.109	2-	2	IXXXXXXXXXXXXX	KOUT
3136	44973	47.282	3.297	3-	3	IXXX	KOUTN
4578	49551	52.095	4.813	4-	4	IXXXXX	KIN
1565	51116	53.740	1.645	5-	5	IXX	RETURN
1882	52998	55.719	1.979	6-	6	IXX	DEFIL
0	52998	55.719	0.	7-	7	I	ABORT
1383	54381	57.173	1.454	8-	8	IX	SETSWH
294	54675	57.482	0.309	9-	9	I	RSTSWH
3604	58279	61.271	3.789	10-	10	IXXXX	REW
443	58722	61.737	0.466	11-	11	I	FILSP
1616	60338	63.436	1.699	12-	12	IXX	RETFIL
179	60517	63.624	0.188	13-	13	I	RELMEM
216	60733	63.851	0.227	14-	14	I	ADDMEM
128	60861	63.985	0.135	15-	15	I	CORFIL
98	60959	64.088	0.103	16-	16	I	SNUMB
13594	74553	78.380	14.292	17-	17	IXXXXXXXXXXXXXX	TIME
74	74627	78.458	0.078	18-	18	I	PASAFT
419	75046	78.899	0.441	19-	19	I	TERMTP
774	75820	79.712	0.814	20-	20	IX	PDIO
376	76196	80.108	0.395	21-	21	I	RESTOR
96	76292	80.209	0.101	22-	22	I	SPAWN
0	76292	80.209	0.	23-	23	I	TAPEIN
763	77055	81.011	0.802	24-	24	IX	CALLSS
70	77125	81.084	0.074	25-	25	I	USERID
54	77179	81.141	0.057	26-	26	I	TERMPG
2198	79377	83.452	2.311	27-	27	IXX	PASUST
5	79382	83.457	0.005	28-	28	I	MORINK

95117 ENTRIES TOTAL 15735 OUT OF RANGE

Figure 15-11. Derails Executed By User Histogram
(Part 1 of 2)

TSS REDUCTION, VERSION 7.2 - 1.000, 30 SEPTEMBER 82

DISTRIBUTION COLLECTED ON SYSTEM DSCC1 AT 0818:43.515 FRI 82-05-21. INITIAL GMC TAPE #1469

DERAILS (HIGH HALF)

INDIV. NUMBER	CUMUL. NUMBER	CUMUL. PRC	INDIV. PRC	TSS DRI #	PERCENT OF OCCURRENCE 00 10 20 30 40 50 60 70 80 90 100	HISTOGRAM 4
					I....I....I....I....I....I....I....I....I....I....I	
79382	79382	83.457	83.457	-4- 28	IXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX	MORLNK
0	79382	83.457	0.	29- 29	I	NEWUSR
1948	81330	85.505	2.048	30- 30	IX	FILACT
5	81335	85.510	0.005	31- 31	I	SETLNO
0	81335	85.510	0.	32- 32	I	SYSRET
0	81335	85.510	0.	33- 33	I	STPSYS
0	81335	85.510	0.	34- 34	I	STATUS
81	81416	85.596	0.085	35- 35	I	DRLDSC
157	81573	85.761	0.165	36- 36	I	PASDES
3054	84627	88.971	3.211	37- 37	IXX	JSTS
71	84698	89.046	0.075	38- 38	I	CGROUT
14	84712	89.061	0.015	39- 39	I	PART
43	84755	89.106	0.045	40- 40	I	GROW
0	84755	89.106	0.	41- 41	I	ABTJOB
0	84755	89.106	0.	42- 42	I	CONSOL
401	85156	89.528	0.422	43- 43	I	SWITCH
0	85156	89.528	0.	44- 44	I	DRLIMT
463	85619	90.014	0.487	45- 45	I	JOUT
1103	86722	91.174	1.160	46- 46	IX	KOTNOW
233	86955	91.419	0.245	47- 47	I	OBJTIM
55	87010	91.477	0.058	48- 48	I	PASFIR
0	87010	91.477	0.	49- 49	I	STOPPT
1	87011	91.478	0.001	50- 50	I	DRLSAV
0	87011	91.478	0.	51- 51	I	TASK
4038	91049	95.723	4.245	52- 52	IXX	PSEUDO
313	91362	96.052	0.329	53- 53	I	PRGDES
2980	94342	99.185	3.133	54- 54	IXX	GWAKE
7	94349	99.193	0.007	55- 55	I	T.CRYP
15	94364	99.208	0.016	56- 56	I	ATTRI
8	94372	99.217	0.008	57- 57	I	T.STAT
238	94610	99.467	0.250	58- 58	I	T.GOTO
233	94843	99.712	0.245	59- 59	I	T.CMOV
5	94848	99.717	0.005	60- 60	I	T.LINK
0	94848	99.717	0.	61- 61	I	T.SYOT
3	94851	99.720	0.003	62- 62	I	T.CONN
125	94976	99.852	0.131	63- 63	I	T.CFIO
0	94976	99.852	0.	64- 64	I	T.EXEC
0	94976	99.852	0.	65- 65	I	T.RSCC
0	94976	99.852	0.	66- 66	I	T.MAIL
141	95117	100.000	0.148	67- 67	I	T.ERR

95117 ENTRIES TOTAL 0 OUT OF RANGE

Figure 15-11. (Part 2 of 2)

TSS REDUCTION, VERSION 7.2 - 1.000, 30 SEPTEMBER 82

DISTRIBUTION COLLECTED ON SYSTEM DSCC1 AT 0818:33.201 FRI 82-05-21. INITIAL GMC TAPE #1469

EXCEPTION MESSAGE REPORT

TOD & RECNO		EXCEPTION MESSAGE	
0818:56.604*	43	LINE <TN>, USERID PRC104	:: SWAP SPACE DENIED.
0818:56.605	45	LINE <TN>, USERID PRC104	:: JOUT OUTPUT BUSY.
0818:56.610*	49	LINE <TN>, USERID PRC104	:: CANNOT SWAP.
0818:56.693*	55	LINE <TN>, USERID PRC104	:: I/O ERROR ON DEVICE.
0818:56.700	61	LINE <TN>, USERID PRC104	:: UNDEFINED DEVICE.
0818:56.770	73	LINE <TN>, USERID PRC104	:: SEEK ERROR.
0818:56.790	85	LINE <TN>, USERID PRC104	:: FAILURE IN NAME SCAN.
0818:56.854	105	LINE <TN>, USERID PRC104	:: BAD SPACE INVENTORY.
0818:56.926	111	LINE <TN>, USERID PRC104	:: BAD SPACE TABLE.
0818:56.933*	117	LINE <TN>, USERID PRC104	:: UNACCOUNTABLE ERROR 1.
0818:56.934*	121	LINE <TN>, USERID PRC104	:: CHECKSUM ERROR.
0818:56.935	123	LINE <TN>, USERID PRC104	:: PAGE REQUEST DEADLOCK.
0818:56.937*	127	LINE <TN>, USERID PRC104	:: UNACCOUNTABLE ERROR 2.
0818:57.053	133	LINE <TN>, USERID PRC104	:: BAD PRIMITIVE.
0818:57.060*	139	LINE <TN>, USERID PRC104	:: BAD PROGRAM DESCRIPTOR.
0819:08.916*	198	LINE <TN>, USERID PRC104	:: LOOP IN PRIMITIVES.
0819:08.917	200	LINE <TN>, USERID PRC104	:: PROGRAM TOO LARGE TO SWAP.
0819:08.918*	202	LINE <TN>, USERID PRC104	:: BAD STATUS - SWAP OUT.
0819:09.045*	204	LINE <TN>, USERID PRC104	:: BAD STATUS - SWAP IN.
0819:09.049*	208	LINE <TN>, USERID PRC104	:: BAD STATUS - LOAD.
0819:09.110	214	LINE <TN>, USERID PRC104	:: ERROR IN WRITING TAP*.
0819:09.116*	220	LINE <TN>, USERID PRC104	:: NOT ENOUGH CORE TO RUN PROGRAM.
[illegible]	222	LINE <TN>, USERID PRC104	:: OUT OF SWAP SPACE.
0819:09.264*	226	LINE <TN>, USERID PRC104	:: SY** I/O ERROR.
0819:09.305	238	LINE <TN>, USERID PRC104	:: DRI DEFIL - NO FILE SPACE.
0819:09.358*	250	LINE <TN>, USERID PRC104	:: DRI DEFIL - NO PAT SPACE.
0819:09.360	252	LINE <TN>, USERID PRC104	:: DRI MORINK - NO FILE SPACE.
0819:09.411	254	LINE <TN>, USERID PRC104	:: DRI MORINK - PAT FULL.
0819:09.412	256	LINE <TN>, USERID PRC104	:: DRI JOUT - BATCH SYSTEM FULL.
0819:09.416*	260	LINE <TN>, USERID PRC104	:: DRI JOUT - LOST PAT.
0819:12.996*	314	LINE <CV>, USERID PRC105	:: DRI SPAWN - NO PAT.
0819:12.997*	316	LINE <CV>, USERID PRC105	:: DRI SPAWN - NO PROGRAM NO.
0819:12.998*	318	LINE <CV>, USERID PRC105	:: DRI SPAWN - SCHEDULER QUEUE FULL.
0819:13.114*	320	LINE <CV>, USERID PRC105	:: DRI TASK - NO PROGRAM NO.
0819:13.119	324	LINE <CV>, USERID PRC105	:: DRI TASK - BAD *J STATUS.
0819:13.225*	334	LINE <CV>, USERID PRC105	:: DRI TASK - NO PUSHDOWN FILE.
0819:13.232	340	LINE <CV>, USERID PRC105	:: DRI T.SYOT - SYSTEM LOADED.
0819:13.440	342	LINE <CV>, USERID PRC105	:: DRI T.SYOT - BACKDOOR FILE NOT CONFIGD.
0819:13.444*	346		:: GEMORE DENIED.
0819:13.528	360	LINE <CV>, USERID PRC105	:: CPU HOG? CPS SINCE SWAPIN =

Figure 15-12. Exception Message Report

```
TSS REDUCTION, VERSION 7.2 - 1.000, 30 SEPTEMBER 82

                    DISTRIBUTION COLLECTED ON SYSTEM DSCC1  AT 0818:33.201 FRI 82-05-21.  INITIAL GMC TAPE #1469

                                  USERS MAP

**********************************************REASON FOR FOLLOWING MAP**********************************************
0818:56.604*        43 LINE <TN>, USERID PRC104       :: SWAP SPACE DENIED.
----------------------------------------------------------------------------------------------------------------------

 USR USTOCT  LID USER-ID----- TIME-IN:STATE TRMIO LAST-TRACE----------&DERAIL--- SIZE  PROGRAM-STACK----------
 --- ------ --- ------------ ------- ----- ----- -------------------- -------- ----- ------------------------
   1 065060 <TN> PRC104              0 WTSUB       88 S/D FAULT QUEUE     0 ZERO     2KT NONE CMDI
   2 063530 <DI> MASB             2591 UNKWN       90 TRACE ON/OFF        0 ZERO     1K  NONE
   3 063230 <CV> PRC105           2591 UNKWN       90 TRACE ON/OFF        0 ZERO     1K  NONE
   4 062730 <TC> C721             2591 UNKWN       90 TRACE ON/OFF        0 ZERO     1K  CARD

**********************************************REASON FOR FOLLOWING MAP**********************************************
0818:56.610*        49 LINE <TN>, USERID PRC104       :: CANNOT SWAP.
----------------------------------------------------------------------------------------------------------------------

 USR USTOCT  LID USER-ID----- TIME-IN:STATE TRMIO LAST-TRACE----------&DERAIL--- SIZE  PROGRAM-STACK----------
 --- ------ --- ------------ ------- ----- ----- -------------------- -------- ----- ------------------------
   1 065060 <TN> PRC104              0 WTSUB       88 S/D FAULT QUEUE    27 PASUST   2KT NONE CMDI
   2 063530 <DI> MASB             2597 UNKWN       90 TRACE ON/OFF        0 ZERO     1K  NONE
   3 063230 <CV> PRC105           2597 UNKWN       90 TRACE ON/OFF        0 ZERO     1K  NONE
   4 062730 <TC> C721             2597 UNKWN       90 TRACE ON/OFF        0 ZERO     1K  CARD

**********************************************REASON FOR FOLLOWING MAP**********************************************
0818:56.693*        55 LINE <TN>, USERID PRC104       :: I/O ERROR ON DEVICE.
----------------------------------------------------------------------------------------------------------------------

 USR USTOCT  LID USER-ID----- TIME-IN:STATE TRMIO LAST-TRACE----------&DERAIL--- SIZE  PROGRAM-STACK----------
 --- ------ --- ------------ ------- ----- ----- -------------------- -------- ----- ------------------------
   1 065060 <TN> PRC104              0 WTSUB       29 GEFSYE/27 COMPLETE 30 FIIACT   2KT NONE CMDI
   2 063530 <DI> MASB             2680 UNKWN       90 TRACE ON/OFF        0 ZERO     1K  NONE
   3 063230 <CV> PRC105           2680 UNKWN       90 TRACE ON/OFF        0 ZERO     1K  NONE
   4 062730 <TC> C721             2680 UNKWN       90 TRACE ON/OFF        0 ZERO     1K  CARD

**********************************************REASON FOR FOLLOWING MAP**********************************************
0818:56.933*       117 LINE <TN>, USERID PRC104       :: UNACCOUNTABLE ERROR 1.
----------------------------------------------------------------------------------------------------------------------

 USR USTOCT  LID USER-ID----- TIME-IN:STATE TRMIO LAST-TRACE----------&DERAIL--- SIZE  PROGRAM-STACK----------
 --- ------ --- ------------ ------- ----- ----- -------------------- -------- ----- ------------------------
   1 065060 <TN> PRC104              0 WTSUB       88 S/D FAULT QUEUE    12 RETFIL   2KT NONE CMDI
   2 063530 <DI> MASB             2920 UNKWN       90 TRACE ON/OFF        0 ZERO     1K  NONE
   3 063230 <CV> PRC105           2920 UNKWN       90 TRACE ON/OFF        0 ZERO     1K  NONE
   4 062730 <TC> C721             2920 UNKWN       90 TRACE ON/OFF        0 ZERO     1K  CARD

**********************************************REASON FOR FOLLOWING MAP**********************************************
0818:56.934*       121 LINE <TN>, USERID PRC104       :: CHECKSUM ERROR.
----------------------------------------------------------------------------------------------------------------------

 USR USTOCT  LID USER-ID----- TIME-IN:STATE TRMIO LAST-TRACE----------&DERAIL--- SIZE  PROGRAM-STACK----------
 --- ------ --- ------------ ------- ----- ----- -------------------- -------- ----- ------------------------
   1 065060 <TN> PRC104              0 WTSUB      104 DRI COMPLETE       13 RELMEM   2KT NONE CMDI
   2 063530 <DI> MASB             2921 UNKWN       90 TRACE ON/OFF        0 ZERO     1K  NONE
   3 063230 <CV> PRC105           2921 UNKWN       90 TRACE ON/OFF        0 ZERO     1K  NONE
   4 062730 <TC> C721             2921 UNKWN       90 TRACE ON/OFF        0 ZERO     1K  CARD
```

Figure 15-13. Users Map

TSS REDUCTION, VERSION 7.2 - 1.000, 30 SEPTEMBER 82

DISTRIBUTION COLLECTED ON SYSTEM DSCC1 AT 0818:33.201 FRI 82-05-21. INITIAL GMC TAPE #1469

TSS DELAY REPORT --- USER DELAY REPORT

TOD	RCNT	*-TSS DELAYED-- DLAY	PRV,CUR	--POSSIBLE REASON & TRACE-*:*--USER DELAYED IN VULNERABLE STATE (:) OR USERS SNAPSHOT (>)-----* XEC	DRI	--FMS--	USR-I/O	: UST	LINE	USERID------	DLAY	STATE	T	PRV	DERATI	SIZE	PROGRAM	STACK--
0823:05.297	6371		61					: 1	<TN>	PRC104	1	DRISR		60	JOUT	8K	NONE	JOUT
0823:05.299	6375		91					: 1	<TN>	PRC104	1	DRISR+		60	JOUT	8K	NONE	JOUT
0831:32.618	30344		61					: 1	<TN>	PRC104	3	DRISR		60	JOUT	8K	NONE	JOUT
0831:32.620	30348		91					: 1	<TN>	PRC104	3	DRISR+		60	JOUT	8K	NONE	JOUT
0845:03.892	73475	1	18, 61	18														
								> 10	<3X>	IRCNTI	1	IDLME	*	80	T.GOTO	3K	TSEC	
								> 9	<TX>	SSISM	1	IBUSY	*	78	KOUTN	34K	RUNY	
								> 8	<CG>	WSSIB	41	IBUSY	*	100	RETURN	1K	NONE*	
								> 7	<VV>	C722	33	IBUSY	*	78	KOUTN	26K	NONE	CMDI
								> 6	<TK>	C433	1	WTMEM		100	GWAKE	1K	NONE	CMDI
								> 5	<TA>	HIS8010	12	IBUSY	*	100	KOUTN	5K	NONE	CMDI
								> 1	<TN>	PRC104	2	IBUSY	*	78	KOUTN	8K	NONE	JOUT
								> 2	<DI>	MASB	556	NOTSS		100	CGROUT	12K	NONE	JDAC
								> 3	<CV>	PRC105	5	IBUSY	*	100	KOUTN	11K	CARD	EDTX
								> 4	<TC>	C721	43	IBUSY	*	100	KOUTN	11K	NONE	EDTX
0846:03.890	84574		29*					: 12	<T3>	CSC	1	FMSER		60	FILACT	5K	NONE	CMDI
0847:07.558	95739		61					: 11	<DN>	HIS	2	DRISR		60	JOUT	8K	NONE	GRUN JOUT
0847:07.567	95749		91					: 11	<DN>	HIS	2	DRISR+		100	JOUT	8K	NONE	GRUN JOUT
0848:44.024*	100696		77					: 11	9595	CSC	4	WTMEM		78	KOUTN	3K	TSEC	
0848:44.024*	100698		80					: 11	9595	CSC	4	WTMEM+		78	KOUTN	3K	TSEC	
0852:37.491	115212	6	23,104					: 10	<3X>	IRCNTI	6	DRISR		87	DEFIL	38K	MDQ	CMDI
								> 13	<T2>	PRC101	63	GWAKE		100	GWAKE	1K	EDTX	FRN RUNY
								> 12	<T5>	CSC	12	IBUSY	*	78	KOUTN	5K	NONE	CMDI
								> 10	<3X>	IRCNTI	0	WTSUB		104	DEFIL	38K	MDQ	CMDI
								> 9	<TX>	SSISM	383	IBUSY	*	100	RETURN	2K	NONE*	
								> 11	<CH>	ACMIS	2	WTMEM		78	GWAKE	1K	NONE	EDTX JRN
								> 7	<VV>	C722	95	IBUSY	*	78	KOUTN	26K	NONE	CMDI
								> 6	<TN>	CSC7592	7	IBUSY	*	100	RETURN	1K	NONE	LIST
								> 2	<DI>	MASB	1009	NOTSS		100	CGROUT	12K	NONE	JDAC
								> 3	<CV>	PRC105	2	WTMEM		80	KOUTN	11K	CARD	EDTX
								> 4	<TC>	C721	28	IBUSY	*	100	KOUTN	11K	NONE	EDTX
0854:03.303	121735		29*					: 10	<3X>	IRCNTI	1	FMSER	*	60	FILACT	38K	MDQ	CMDI
0910:12.780	258364		100					: 12	<T3>	CSC	2	DRISR		60	JOUT	8K	NONE	CMDI JOUT
0910:12.781	258366		91					: 12	<T3>	CSC	2	DRISR+		100	JOUT	8K	NONE	CMDI JOUT
0918:20.932	301132	2	23,104					: 10	<3X>	IRCNTI	2	DRISR		87	DEFIL	38K	MDQ	CMDI
								> 10	<3X>	IRCNTI	0	WTSUB		104	DEFIL	38K	MDQ	CMDI
								> 11	<CH>	ACMIS	139	IBUSY	*	100	KOUTN	38K	RUNY	

Figure 15-14. TSS Delay and Users Delay Report

```
TSS REDUCTION, VERSION 7.2 - 1.000, 30 SEPTEMBER 82
                DISTRIBUTION COLLECTED ON SYSTEM DSCC1  AT 0818:33.201 FRI 82-05-21.  INITIAL GMC TAPE #1469

                                TSS CORE MAP

**********************************************REASON FOR FOLLOWING MAP**********************************************
0848:44.024* 100696             77                              :  11 9595 CSC          4 WTMEM    78 KOUTN    3K TSEC
-------------------------------------------------------------------------------------------------------------------------

 USR BASE SIZE FREE  LID USER-ID----- TSTAT  TCORE TCPU DIOS ISTDRI PROGRAM-STACK---------- USTOCT STATE TRMIO LAST-TRACE-------
--- ---- ---- ----  --- ------------ -----  ----- ---- ---- ------ ---------------------- ------ ----- ----- ----------------
HOLE  31K       58K

**********************************************REASON FOR FOLLOWING MAP**********************************************
0848:44.024* 100698             80                              :  11 9595 CSC          4 WTMEM+   78 KOUTN    3K TSEC
-------------------------------------------------------------------------------------------------------------------------

 USR BASE SIZE FREE  LID USER-ID----- TSTAT  TCORE TCPU DIOS ISTDRI PROGRAM-STACK---------- USTOCT STATE TRMIO LAST-TRACE-------
--- ---- ---- ----  --- ------------ -----  ----- ---- ---- ------ ---------------------- ------ ----- ----- ----------------
  11  31K   3K  55K 9595 CSC              0   4392    0    0 KOUTN  TSEC                   071610 SWPIN   F   TIME TYPES

**********************************************REASON FOR FOLLOWING MAP**********************************************
1031:42.147*1047558             77                              :   5 <TN> PRC104      12 WTMEM   100 GWAKE    2K CARD EDTX DJST
-------------------------------------------------------------------------------------------------------------------------

 USR BASE SIZE FREE  LID USER-ID----- TSTAT  TCORE TCPU DIOS ISTDRI PROGRAM-STACK---------- USTOCT STATE TRMIO LAST-TRACE-------
--- ---- ---- ----  --- ------------ -----  ----- ---- ---- ------ ---------------------- ------ ----- ----- ----------------
HOLE  36K       5K
  25  41K   4K  45K <DI> MASA             0  16759   28    0 FILACT NONE MAS2              064030 SUBDS   F   MAKE S/D ENTRY

**********************************************REASON FOR FOLLOWING MAP**********************************************
1031:42.148*1047560             80                              :   5 <TN> PRC104      12 WTMEM+  78 GWAKE    2K CARD EDTX DJST
-------------------------------------------------------------------------------------------------------------------------

 USR BASE SIZE FREE  LID USER-ID----- TSTAT  TCORE TCPU DIOS ISTDRI PROGRAM-STACK---------- USTOCT STATE TRMIO LAST-TRACE-------
--- ---- ---- ----  --- ------------ -----  ----- ---- ---- ------ ---------------------- ------ ----- ----- ----------------
   5  36K   2K   3K <TN> PRC104            0  17219    0    0 GWAKE  CARD EDTX DJST         070010 SWPIN   F   TIME TYPES
  25  41K   4K  45K <DI> MASA             2  16760   28    0 FILACT NONE MAS2              064030 SUBDS   F   MAKE S/D ENTRY

**********************************************REASON FOR FOLLOWING MAP**********************************************
1031:42.152*1047564             80                              :  12 <TM> WSSIB       10 WTMEM+  78 GWAKE    2K DJST
-------------------------------------------------------------------------------------------------------------------------

 USR BASE SIZE FREE  LID USER-ID----- TSTAT  TCORE TCPU DIOS ISTDRI PROGRAM-STACK---------- USTOCT STATE TRMIO LAST-TRACE-------
--- ---- ---- ----  --- ------------ -----  ----- ---- ---- ------ ---------------------- ------ ----- ----- ----------------
   5  36K   2K   0K <TN> PRC104            3  17222    0    0 GWAKE  CARD EDTX DJST         070010 SWPIN   F   TIME TYPES
  12  38K   2K   1K <TM> WSSIB            0  14541   12   20 GWAKE  DJST                   077570 SWPIN   F   TIME TYPES
  25  41K   4K  45K <DI> MASA             5  16763   28    0 FILACT NONE MAS2              064030 SUBDS   F   MAKE S/D ENTRY

**********************************************REASON FOR FOLLOWING MAP**********************************************
1031:46.303*1047827             77                              :   1 <TS> PRC105       4 WTMEM    77 KOUTN   11K CARD EDTX
-------------------------------------------------------------------------------------------------------------------------

 USR BASE SIZE FREE  LID USER-ID----- TSTAT  TCORE TCPU DIOS ISTDRI PROGRAM-STACK---------- USTOCT STATE TRMIO LAST-TRACE-------
--- ---- ---- ----  --- ------------ -----  ----- ---- ---- ------ ---------------------- ------ ----- ----- ----------------
HOLE  36K       2K
  12  38K   2K   0K <TM> WSSIB            1    928    1    2 DIO    DJST                   077570 SUBDS   F   MAKE S/D ENTRY
  18  40K   2K   0K <TK> USERID UNKWN     2    646    0    0 CGROUT NEW                    073440 SUBDS   F   MAKE S/D ENTRY
  25  42K   2K   4K <DI> MASA             0    144    0    0 RETURN NONE CMDI              064030 SUBDS   F   MAKE S/D ENTRY
```

Figure 15-15. TSS Core Map

TSS REDUCTION, VERSION 7.2 - 1.000, 30 SEPTEMBER 82

DISTRIBUTION COLLECTED ON SYSTEM DSCC1 AT 0818:33.201 FRI 82-05-21. INITIAL GMC TAPE #1469

ERROR MESSAGE REPORT

TOD & RECNO		ERROR MESSAGE	
0850:19.630	106727	LINE <TK>, USERID C433	:: LAST KNOWN - WAIT NORM. JUST REPORTED - OUT OF MEM.
0920:36.929	3148308	LINE 7412, USERID DJ8XC344AM	:: LAST KNOWN - WAIT FORCD. JUST REPORTED - OUT OF MEM.

Figure 15-16. Error Message Report

TSS REDUCTION, VERSION 7.2 - 1.000, 30 SEPTEMBER 82

DISTRIBUTION COLLECTED ON SYSTEM DSCC1 AT 0818:33.201 FRI 82-05-21. INITIAL GMC TAPE #1469

INTERTRACE GAP (DURATION) - ADJUSTED FOR TSS GEWAKES PLOT- 7

-= MIN /= AVG += MAX

USERS	TIME	DURATION IN INTERVAL (MS) <4	<8	<16	<32	<64	<128	>128	GAP TIMES (MILLISEC) 1.0E-01 1.0E 00 1.0E 01 1.0E 02 1.0E 03
0	0818:54.01	DATA NOT COLLECTED IN THE INTERVAL.							X
4	0819:14.01	365	0	2	1	5	2	2	*/////////////////////+++++++++++++++++++++++++
4	0819:34.01	607	0	0	1	2			*/////////+++++++++++++++++++++++++++
4	0819:54.01	486							*/////////++++++++++
4	0820:14.01	151							*/////////+++++++++
4	0820:34.01	261	0	0	0	0	0	1	*/////////////+++++++++++++++++++++++++++++
4	0820:54.01	93							*/////////++++++++++
4	0821:14.01	455	0	0	2				*/////////++++++++++++++++++++++++
4	0821:34.01	674	4	4	0	0	0	3	*/////////////////+++++++++++++++++++++++++++++++
4	0821:54.01	758	0	0	0	0	0	2	*/////////////////+++++++++++++++++++++++++++++++++++
4	0822:14.01	1451	3	2	1				*/////////++++++++++++++++++++++++
4	0822:34.01	175							*/////////++++++++
4	0822:54.01	16							*/////////+++++
4	0823:14.01	380							*/////////++++++++++
5	0823:34.01	1155	8	7	2	0	2	1	*/////////////////+++++++++++++++++++++++++++++++++
5	0823:54.01	581	0	5	2	0	1	1	*/////////////////+++++++++++++++++++++++++++++
6	0824:14.01	446	2	2	1				*/////////++++++++++++++++++++++
6	0824:34.01	864	18	23	1	0	1	2	*/////////////////+++++++++++++++++++++++++++++++++
6	0824:54.01	374	2	9	0	0	2	1	*/////////////////+++++++++++++++++++++++++++++++
6	0825:14.01	889	19	24					*/////////////////++++++++++++++++
6	0825:34.01	513							*/////////++++++++++
6	0825:54.01	588	17	25	2				*/////////////////++++++++++++++++++
6	0826:14.01	1631	73	137	2				*/////////////////++++++++++++++
6	0826:34.01	276							*/////////++++++++++
7	0826:54.01	3145	101	150	4				*/////////////////++++++++++++++++++++
7	0827:14.01	1126	20	26	2				*/////////////////++++++++++++++++++++
7	0827:34.01	184							*/////////++++++++++
7	0827:54.01	117							*/////////++++++++++
7	0828:14.01	311	5	7	0	1			*/////////////////+++++++++++++++++++++++++
7	0828:34.01	409	3	6	1				*/////////////////++++++++++++++++++++
8	0828:54.01	806	3	1	1				*/////////++++++++++++++++
8	0829:14.01	323	0	2	0	1	0	1	*/////////////////+++++++++++++++++++++++++++++
8	0829:34.01	913	1	8	1	1			*/////////////////++++++++++++++++++++++++++
8	0829:54.01	1220	0	2	2	1	2	2	*/////////////////+++++++++++++++++++++++++++++++
8	0830:14.01	666	4	4	4	0	1		*/////////////////++++++++++++++++++++++++
8	0830:34.01	867	21	23	2	0	2	3	*/////////////////////+++++++++++++++++++++++++++++++
8	0830:54.01	860	1	1	0	1	3		*/////////////////+++++++++++++++++++++++++++++
8	0831:14.01	1478	2	3					*/////////////////++++++++++++++++++
8	0831:34.01	1236	4	1	0	2	0	1	*/////////////////+++++++++++++++++++++++++++++++++
8	0831:54.01	368	1	2					*/////////////////++++++++++++++++++
8	0832:14.01	1250	0	0	0	0	0	1	*/////////////////+++++++++++++++++++++++++++++++
8	0832:34.01	882	24	25	0	1	1	2	*/////////////////////+++++++++++++++++++++++++++++++++
8	0832:54.01	705	0	0	0	0	0	1	*/////////////////+++++++++++++++++++++++++++++++

CURVE	MINIMUM	MAXIMUM
MIN	0.	0.
AVG	4.3866E-01	1.0582E 01
MAX	1.0940E 00	1.2234E 04

Figure 15-17. Intertrace Gap Duration Plot

TSS REDUCTION, VERSION 7.2 - 1.000, 30 SEPTEMBER 82

DISTRIBUTION COLLECTED ON SYSTEM DSCC1 AT 0818:33.201 FRI 82-05-21. INITIAL GMC TAPE #1469

MEMORY ACTIVITY (TRACES 63, 65-75) REPORT- 1

TOD	TRACE 63	TRACE 65	TRACE 66	TRACE 67	TRACE 68	TRACE 69	TRACE 70	TRACE 71	TRACE 72	TRACE 73	TRACE 74	TRACE 75
0818:54.012												
0819:14.012												
0819:34.012												
0819:54.012												
0820:14.012												
0820:34.012												
0820:54.012												
0821:14.012												
0821:34.012												
0821:54.012												
0822:14.012												
0822:34.012												
0822:54.012												
0823:14.012												
0823:34.012	-	-	-	-	-	-	-	1				
0823:54.012												
0824:14.012												
0824:34.012												
0824:54.012												
0825:14.012												
0825:34.012												
0825:54.012	1	-	-	-	-	1	-	1				
0826:14.012	1											
0826:34.012												
0826:54.012												
0827:14.012												
0827:34.012												
0827:54.012												
0828:14.012												
0828:34.012												
0828:54.012												
0829:14.012												
0829:34.012												
0829:54.012												
0830:14.012	1											
0830:34.012												
0830:54.012												
0831:14.012												
0831:34.012												
0831:54.012												
0832:14.012												
0832:34.012												
0832:54.012												

63=ENTER SWAP DECIS
65=CONSIDER TSS INCR
66=INITIATE SIZE INCR
67=SET UP FENCE
68=MEMORY RESERVE
69=FORCE SWAP
70=TERMINATE SWAP
71=TSSI UST AREA +1K
72=GEMORE MEMORY
73=GEMORE SUCCESSFUL
74=MEMORY RELEASE
75=GEMORE REFUSED

Figure 15-18. Memory Activity Report

TSS REDUCTION, VERSION 7.2 - 1.000, 30 SEPTEMBER 82

DISTRIBUTION COLLECTED ON SYSTEM DSCC1 AT 0818:33.201 FRI 82-05-21. INITIAL GMC TAPE #1469

ELIGIBIES FOR SUBDISPATCH (DURATION) PLOT- 8

-= MIN /= AVG .+= MAX

USERS	TIME	DURATION IN INTERVAL (MS) <1	<2	<4	<8	<16	<32	>32	LINE	ELIGIBIE TIMES (MILLISEC) 1.0E-01	1.0E 00	1.0E 01	1.0E 02
23	1025:14.01	260	5	109	5				<TS>	I---------///////	+++++++++++++++		
23	1025:34.01	407	12	90	28	4	2	1	<TP>	X---------///////	+++++++++++++++++++++++++++++++		
23	1025:54.01	319	9	66	26	6			<TN>	I---------///////	+++++++++++++++		
22	1026:14.01	508	13	157	42	0	0	5	<TO>	I---------///////	+++++++++++++++++++++++++++++++		
23	1026:34.01	402	10	108	32	9	2	7	<DB>	I---------///////	+++++++++++++++++++++++++++++++		
23	1026:54.01	563	12	150	23	3	1	4	<CH>	I---------///////	+++++++++++++++++++++++++++++++		
24	1027:14.01	482	9	115	12	2	0	1	<DB>	I---------///////	+++++++++++++++++++++++++++++++		
24	1027:34.01	613	10	124	36	7	0	8	<DJ>	I---------///////	+++++++++++++++++++++++++++++++		
24	1027:54.01	482	14	140	30	5	0	1	<DJ>	I---------///////	+++++++++++++++++++++++++++++++		
23	1028:14.01	428	18	109	17	7	0	2	<TO>	I---------///////	+++++++++++++++++++++++++++++++		
23	1028:34.01	640	27	144	23	1	0	1	<TM>	I---------///////	+++++++++++++++++++++++++++++++		
23	1028:54.01	353	11	113	6	3	0	1	<CH>	X---------///////	+++++++++++++++++++++++++++++++		
24	1029:14.01	395	12	135	28	10	0	1	<DB>	I---------///////	###############################		
25	1029:34.01	493	9	142	14	3	1		<TS>	I---------///////	+++++++++++++++++++++		
24	1029:54.01	604	18	191	56	16	3	4	<TN>	I---------///////	+++++++++++++++++++++++++++++++		
23	1030:14.01	511	16	192	39	5	0	2	<TS>	I---------///////	+++++++++++++++++++++++++++++++		
23	1030:34.01	543	14	182	43	10	3	1	<TM>	I---------///////	+++++++++++++++++++++++++		
21	1030:54.01	549	15	148	41	15	2	3	<TM>	I---------///////	+++++++++++++++++++++++++++++++		
21	1031:14.01	347	8	122	8	1	1	2	<TS>	I---------///////	+++++++++++++++++++++++++++		
21	1031:34.01	392	5	157	26	7	1	1	<TS>	I---------///////	+++++++++++++++++++++++++++++++		
22	1031:54.01	236	4	45	12	11	3	6	<TN>	I---------///////	###############################		
22	1032:14.01	392	5	158	16	4	0	2	<TC>	X---------///////	+++++++++++++++++++++++++++++++		
22	1032:34.01	1939	19	184	60	12	0	2	<TX>	I---------///////	+++++++++++++++++++++++++++++++		
22	1032:54.01	681	26	204	81	18	2	6	<TN>	I---------///////	+++++++++++++++++++++++++++++++		
22	1033:14.01	444	5	133	18	4			<DJ>	I---------///////	++++++++++++++++		
22	1033:34.01	488	6	166	17	2	0	1	<CH>	I---------///////	+++++++++++++++++++++++++++++++		
22	1033:54.01	375	4	140	10	0	1		<CH>	I---------///////	+++++++++++++++++++++		
22	1034:14.01	457	15	123	16	2			<TK>	I---------///////	+++++++++++++++++++++		
22	1034:34.01	1117	22	142	39	6	0	1	<TM>	I---------///////	+++++++++++++++++++++++++++++++		
21	1034:48.96	1093	7	119	17	2	0	1	<TC>	I---------///////	+++++++++++++++++++++++++++++++		

CURVE	MINIMUM	MAXIMUM
MIN	7.2000E-02	4.1700E-01
AVG	5.5007E-01	4.8654E 01
MAX	2.1680E 00	7.9746E 03

Figure 15-19. Eligibles For Subdispatch Plot

TSS REDUCTION, VERSION 7.2 - 1.000, 30 SEPTEMBER 82

DISTRIBUTION COLLECTED ON SYSTEM DSCC1 AT 0818:33.201 FRI 82-05-21. INITIAL GMC TAPE #1469

IN SUBDISPATCH (DURATION) PLOT- 9

-= MIN /= AVG += MAX

USERS	TIME	DURATION IN INTERVAL (MS) <1	<2	<4	<8	<16	<32	>32	LINE	SUBDISPATCH TIMES (MILLISEC) 1.0E 00	1.0E 01	1.0E 02	1.0E 03
0	0818:54.01	DATA NOT COLLECTED IN THIS INTERVAL.								X			
4	0819:14.01	0	1	18	6				<TN>	I---///////+++			
4	0819:34.01	0	0	55	8				<TN>	I----------/+			
4	0819:54.01	0	0	36	14				<TN>	I----------/++++			
4	0820:14.01	0	0	19	1				<CV>	I----------/++++			
4	0820:34.01	0	0	13	8				<TN>	I----------/++			
4	0820:54.01	0	0	2	4				<TN>	I----------/+++			
4	0821:14.01	0	1	27	11				<TN>	I----///////++			
4	0821:34.01	0	0	35	24	0	7	3	<TN>	I---------////////++++++++++++++++++++++++			
4	0821:54.01	0	1	41	27	2			<TN>	I----///////+++++++			
4	0822:14.01	0	0	81	50				<TN>	X-----////+++			
4	0822:34.01	0	1	11	12				<CV>	I---///////+			
4	0822:54.01									I			
4	0823:14.01	0	1	11	14	6			<TN>	I----///////+++++			
5	0823:34.01	0	1	46	28	23	9	1	<TC>	I----/////////++++++++++++			
5	0823:54.01	0	2	36	13				<TC>	I---///////+++			
6	0824:14.01	0	0	29	10				<TK>	I----------/+			
6	0824:34.01	0	2	31	30	11	40	2	<TK>	I---////////////////+++++++++			
6	0824:54.01	0	2	12	12	5	7		<TC>	I---////////////////++++++++			
6	0825:14.01	0	1	36	34	3	42	1	<TK>	I----////////////////++++++++++			
6	0825:34.01	0	3	42	17	1			<TA>	X---///////+++++			
6	0825:54.01	0	1	16	18	3	43		<TK>	I---////////////////+++++			
6	0826:14.01	0	0	6	17	28	219	2	<TN>	I--------////////////////+++++			
6	0826:34.01	0	1	20	11				<TA>	I---///////+++			
7	0826:54.01	0	5	100	60	39	255	6	<TK>	I---////////////////+++++++			
7	0827:14.01	0	1	58	23	4	43		<TK>	I---////////////////+++++++			
7	0827:34.01	0	0	9	5				<TA>	I----------/++			
7	0827:54.01	0	2	3	3				<TK>	I---////+++			
7	0828:14.01	0	1	13	6	3	9		<TK>	I---////////////++++++			
7	0828:34.01	0	1	20	15	1	10		<TK>	I----////////////++++++++			
8	0828:54.01	0	2	58	29	1			<TA>	X----///////+++++			
8	0829:14.01	0	3	14	5	2			<TK>	I---///////++++++++			
8	0829:34.01	0	0	48	48	7	6		<TN>	I---------///++++++++++			

CURVE	MINIMUM	MAXIMUM
MIN	5.9000E-02	4.0850E 00
AVG	3.2497E 00	1.0896E 02
MAX	4.1540E 00	9.2997E 03

Figure 15-20. In Subdispatch Plot

TSS REDUCTION, VERSION 7.2 - 1.000, 30 SEPTEMBER 82

DISTRIBUTION COLLECTED ON SYSTEM DSCC1 AT 0818:33.201 FRI 82-05-21. INITIAL GMC TAPE #1469

ELIGIBLE FOR AND IN SUBDISPATCH (DURATION) PLOT-10

-= MIN /= AVG += MAX

USERS	TIME	DURATION IN INTERVAL (MS) <1	<2	<4	<8	<16	<32	>32	LINE	SUM TIMES (MILLISEC) 1.0E-01 1.0E 00 1.0E 01 1.0E 02 1.0E 03
0	0818:54.01	DATA NOT COLLECTED IN THIS INTERVAL.								X
4	0819:14.01	33	1	2	22	1			<TN>	I-------/////////+++++++
4	0819:34.01	63	0	0	62	1			<TN>	I-------/////////+++++++
4	0819:54.01	48	2	0	50				<CV>	I-------/////////++++++
4	0820:14.01	20	0	1	18	1			<CV>	I-------/////////+++++++++
4	0820:34.01	22	0	0	21				<TN>	I-------/////////+++++
4	0820:54.01	5	1	0	6				<TN>	I-------/////////++++++
4	0821:14.01	40	0	0	39				<TN>	I-------/////////++++++
4	0821:34.01	67	1	0	57	2	7	3	<TN>	I-------/////////////////++++++++++++++++++++++
4	0821:54.01	71	1	1	64	6			<TN>	I-------/////////++++++++++
4	0822:14.01	133	1	2	124	6			<TN>	X-------/////////++++++
4	0822:34.01	24	0	1	23				<CV>	I-------/////////+++++++
4	0822:54.01									I
4	0823:14.01	32	0	0	26	6			<TN>	I-------/////////+++++++
5	0823:34.01	113	1	1	71	26	9	1	<TC>	I-------/////////++++++++++++
5	0823:54.01	52	1	2	47	2			<TC>	I-------/////////+++++++
6	0824:14.01	46	0	0	39				<TK>	I-------/////////++++++
6	0824:34.01	75	0	4	57	13	41	2	<TK>	I-------/////////////////++++++++++
6	0824:54.01	32	0	1	23	6	8		<TC>	I-------/////////////////+++++++++
6	0825:14.01	75	1	1	69	4	42	1	<TK>	I-------/////////////////++++++++++
6	0825:34.01	63	0	2	60	1			<TA>	X-------/////////////////++++++++
6	0825:54.01	40	1	2	31	4	44		<TK>	I-------/////////////////////++++++
6	0826:14.01	52	0	0	23	26	221	2	<TN>	I-------/////////////////////////++++++
6	0826:34.01	32	0	1	31				<TA>	I-------/////////////////++++++
7	0826:54.01	221	2	19	134	51	256	6	<TN>	I-------/////////////////////++++++++++
7	0827:14.01	87	2	4	78	3	44		<TK>	I-------/////////////////////+++++++++
7	0827:34.01	14	0	0	14				<TA>	I-------/////////////////+++++
7	0827:54.01	8	0	2	6				<TK>	I-------/////////////////++++++
7	0828:14.01	22	0	1	19	3	9		<TK>	I-------/////////////////////++++++++
7	0828:34.01	40	1	2	33	2	10		<TK>	I-------/////////////////////+++++++++
8	0828:54.01	97	0	2	86	2			<TA>	X-------/////////////////++++++++
8	0829:14.01	28	1	2	19	3			<TK>	I-------/////////////////++++++++++
8	0829:34.01	108	0	0	85	18	6		<TN>	I-------/////////////////++++++++++++

CURVE	MINIMUM	MAXIMUM
MIN	7.2000E-02	4.1700E-01
AVG	2.3760E 00	5.9207E 01
MAX	6.0190E 00	9.3007E 03

Figure 15-21. Eligibles For and In Subdispatch Plot

TSS REDUCTION, VERSION 7.2 - 1.000, 30 SEPTEMBER 82

DISTRIBUTION COLLECTED ON SYSTEM DSCC1 AT 0818:33.201 FRI 82-05-21. INITIAL GMC TAPE #1469

USER SWAPS (SWAP RATE) PLOT-11

-= RATE

K SWAPPED / SEC

	USERS	TIME	1.0E 00	1.0E 01	1.0E 02
DATA NOT COLLECTED IN THIS INTERVAL.	0	0818:54.01	I		
	4	0819:14.01	*		
	4	0819:34.01	I----------		
	4	0819:54.01	I----------		
	4	0820:14.01	I		
	4	0820:34.01	I----------		
	4	0820:54.01	-		
	4	0821:14.01	I----------		
	4	0821:34.01	-		
	4	0821:54.01	I----------		
	4	0822:14.01	X------		
	4	0822:34.01	I		
	4	0822:54.01	I		
	4	0823:14.01	I----------		
	5	0823:34.01	I----------		
	5	0823:54.01	I------		
	6	0824:14.01	I------		
	6	0824:34.01	I------		
	6	0824:54.01	*		
	6	0825:14.01	-		
	6	0825:34.01	-		
	6	0825:54.01	*		
	6	0826:14.01	-		
	6	0826:34.01	*		
	7	0826:54.01	I----------		
	7	0827:14.01	I----------		
	7	0827:34.01	*		
	7	0827:54.01	*		
	7	0828:14.01	-		
	7	0828:34.01	*		
	8	0828:54.01	X------		
	8	0829:14.01	-		
	8	0829:34.01	I------		

CURVE	MINIMUM	MAXIMUM
RATE	0.	7.3000E 01

Figure 15-22. User Swap Rate Plot

TSS REDUCTION, VERSION 7.2 - 1.000, 30 SEPTEMBER 82

DISTRIBUTION COLLECTED ON SYSTEM DSCC4 AT 0818:33.201 FRI 82-05-21. INITIAL GMC TAPE #1469

USER SWAPS (SWAP AMOUNT) PLOT-12

-= AVG /= TOT

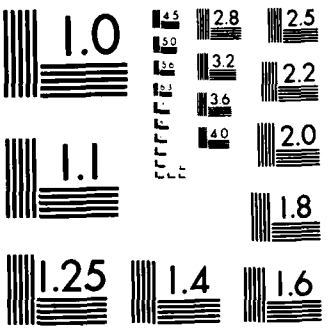
USERS	TIME	CORE IN INTERVAL (K) <2	<4	<8	<16	<32	<64	>64	LINE	SWAP SIZE (K) 1.0E 00	1.0E 01	1.0E 02	1.0E 03
0	0818:54.01	DATA NOT COLLECTED IN THIS INTERVAL.								X			
4	0819:14.01	3	1	2					<TN>	I------///////////			
4	0819:34.01	2	0	3	4				<TN>	I-------------///////////////			
4	0819:54.01	2	0	2	4				<TN>	I-------------///////////////			
4	0820:14.01									I			
4	0820:34.01	0	0	0	6				<TN>	I-------------///////////////			
4	0820:54.01	0	0	0	2				<TN>	I-------------/////			
4	0821:14.01	1	0	1	8				<TN>	I----------//////////////////			
4	0821:34.01	4	1	2	1				<TN>	I--------///////////			
4	0821:54.01	3	1	3	4				<TN>	I-------------///////////////			
4	0822:14.01	3	10	0	2				<TN>	X--------////////////////////			
4	0822:34.01									I			
4	0822:54.01									I			
4	0823:14.01	1	0	0	6				<TN>	I-------------///////////////			
5	0823:34.01	2	10	3	7				<TN>	I-----------////////////////////			
5	0823:54.01	2	1	7					<TC>	I----------///////////			
6	0824:14.01	1	5	4					<TA>	I--------///////////			
6	0824:34.01	1	0	4	1				<TK>	I------------/////////			
6	0824:54.01	1	0	0	1				<TC>	I----------/////			
6	0825:14.01	0	1	1	2				<TC>	I-------------///////////			
6	0825:34.01	1	0	4					<TA>	X----------///////////			
6	0825:54.01	0	1	0	1				<TN>	I----------///////			
6	0826:14.01	1	0	0	3				<TN>	I-------------///////////			
6	0826:34.01	1	0	2					<TA>	I---------////////			
7	0826:54.01	2	10	2	5	2			<VV>	I-------------////////////////////			
7	0827:14.01	3	3	0	1	7			<VV>	I-----------------//////////////////////			
7	0827:34.01	0	0	2					<TA>	I-----------//////			
7	0827:54.01	1							<TK>	2			
7	0828:14.01	1	0	1	1				<TK>	I------------///////			
7	0828:34.01	0	1	1					<TA>	I---------/////			
8	0828:54.01	2	5	4					<TA>	X--------//////////////////			
8	0829:14.01	1	1	1	1				<VV>	I-----------///////			
8	0829:34.01	2	0	3	3				<VV>	I-----------//////////////////			

CURVE	MINIMUM	MAXIMUM
AVG	1.0000E 00	2.0667E 01
TOT	1.0000E 00	1.1770E 03

Figure 15-23. User Swap Amount Plot

AD-A131 954 GENERALIZED MONITORING FACILITY USERS MANUAL CHANGE 6 8/8
(U) COMMAND AND CONTROL TECHNICAL CENTER WASHINGTON DC
B M WALLACK ET AL. 01 JUL 83 CCTC-CSM-UM-246-82-CHG-6
UNCLASSIFIED DCA/DF-83/001B F/G 9/2 NL





MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

TSS REDUCTION, VERSION 7.2 - 1.000, 30 SEPTEMBER 82

DISTRIBUTION COLLECTED ON SYSTEM DSCC1 AT 0818:33.201 FRI 82-05-21. INITIAL GMC TAPE #1469

USER SWAPS (DURATION) PLOT-13

-= MIN /= AVG += MAX

USERS	TIME	DURATION IN INTERVAL (HS) <1	<2	<4	<8	<16	<32	>32	LINE	SWAP TIMES (HUNDREDTHS SEC) 1.0E-01	1.0E 00	1.0E 01	1.0E 02
0	0818:54.01	DATA NOT COLLECTED IN THIS INTERVAL.								X			
4	0819:14.01	0	0	3					<CV>	I----------2			
4	0819:34.01	0	0	2	3	4			<TN>	I----------/////+++++			
4	0819:54.01	0	0	1	1	4			<TN>	I----------/////+++++			
4	0820:14.01									I			
4	0820:34.01	0	0	0	0	6			<TN>	I----------2			
4	0820:54.01	0	0	0	0	2			<TN>	I3333333333333333333333333333333			
4	0821:14.01	0	0	1	1	8			<TN>	I----------/////////+			
4	0821:34.01	0	0	4	1	1			<TN>	I----------///+++++			
4	0821:54.01	0	0	4	1	3	0	1	<TN>	I----------/////////++++++++++			
4	0822:14.01	0	1	6	2	2			<TN>	X----------//////+++++++			
4	0822:34.01									I			
4	0822:54.01									I			
4	0823:14.01	0	0	0	5	1			<TN>	I----------///++++			
5	0823:34.01	0	1	7	9	1			<TA>	I----------////////+++++			
5	0823:54.01	0	0	2	7				<TA>	I----------///+			
6	0824:14.01	0	0	3	4				<TA>	I----------////++			
6	0824:34.01	0	0	0	4	1			<TK>	I----------//++			
6	0824:54.01	0	0	0	0	1			<TC>	I3333333333333333333333333333333			
6	0825:14.01	0	0	1	1	2			<TC>	I----------///////+++++			
6	0825:34.01	0	0	0	4				<TA>	X22222222222222222222222222+			
6	0825:54.01	0	0	2					<TN>	I----------/+			
6	0826:14.01	0	0	0	2	2			<TN>	I----------/+			
6	0826:34.01	0	0	0	2				<TA>	I22222222222222222222222222+			
7	0826:54.01	0	0	6	8	2	1		<VV>	I----------////////+++++++++			
7	0827:14.01	0	0	3	2	1	7		<VV>	I----------///////////+++			
7	0827:34.01	0	0	0	2				<TA>	I22222222222222222222222222+			
7	0827:54.01									I			
7	0828:14.01	1	0	0	0	1			<TK>	I----------/////////+++++			
7	0828:34.01	0	0	1	2				<TA>	I----------////++			
8	0828:54.01	0	0	1	6				<TA>	X----------////+			
8	0829:14.01	0	1	0	1	1			<VV>	I----------/////+++++			
8	0829:34.01	0	0	1	4	3			<VV>	I----------/////+++			

CURVE	MINIMUM	MAXIMUM
MIN	7.5500E-02	9.4358E 00
AVG	3.0151E 00	4.4564E 01
MAX	3.4088E 00	3.2226E 02

Figure 15-24. User Swap Duration Plot

TSS REDUCTION, VERSION 7.2 - 1.000, 30 SEPTEMBER 82

DISTRIBUTION COLLECTED ON SYSTEM DSCC1 AT 0818:33.201 FRI 82-05-21. INITIAL GMC TAPE #1469

USER I/OS (DURATION) PLOT-14

-= MIN /= AVG += MAX

USERS	TIME	DURATION IN INTERVAL (HS) <1	<2	<4	<8	<16	<32	>32	LINE	I/O TIMES (HUNDREDTHS SEC) 1.0E-01	1.0E 00	1.0E 01	1.0E 02
0	0818:54.01	DATA NOT COLLECTED IN THIS INTERVAL.								X			
4	0819:14.01	1	3	4	3				<CV>	I----------/////////+++++			
4	0819:34.01	0	3	2	4				<CV>	I----------////+++			
4	0819:54.01	0	2	2	2				<CV>	I----------////+++++			
4	0820:14.01	0	3						<CV>	I3333333333333333333			
4	0820:34.01	0	1						<TM>	I333333333333333			
4	0820:54.01									I			
4	0821:14.01	0	2	2	2				<TM>	I----------/////+++++			
4	0821:34.01	0	3	12	4				<TC>	I----------/////+++++			
4	0821:54.01	1	5	10	4				<TC>	I----------/////+++++			
4	0822:14.01	34	5	24	41				<TM>	I----------/////+++++			
4	0822:34.01	0	2	0	1				<CV>	I----------/////+++++			
4	0822:54.01									I			
4	0823:14.01	0	0	4	6	3			<TM>	I----------/////+++++			
5	0823:34.01	1	2	3	28	4			<TM>	I----------/////////+++++			
5	0823:54.01	1	3	2	5	1			<TC>	I----------/////////+++++			
6	0824:14.01	3	0	0	2	1			<TK>	I----------/////+++++			
6	0824:34.01	2	4	12	5				<TK>	I----------/////+++++			
6	0824:54.01	0	5	6	3				<TC>	I----------/////+++++			
6	0825:14.01	1	6	6	4	1			<TK>	I----------/////+++++			
6	0825:34.01	1	2	1	1	1			<TK>	I----------/////+++++			
6	0825:54.01	1	5	5	1	1			<TK>	I----------/////+++++			
6	0826:14.01	0	0	0	26	4			<TM>	I----------//++++			
6	0826:34.01	0	1	1	1				<TK>	I----------/////+++++			
7	0826:54.01	4	9	6	30	6	1		<VV>	I----------/////////+++++++++			
7	0827:14.01	2	6	6	4	2			<TK>	I----------/////+++++			
7	0827:34.01	0	0	2					<TA>	I3333333333333333333333			
7	0827:54.01	0	0	0	1				<TK>	I333333333333333333333333333			
7	0828:14.01	0	2	3	1				<TK>	I----------/////+++++			
7	0828:34.01	1	3	2	1	2			<TK>	I----------/////+++++			
8	0828:54.01	1	2	2	2	1			<TK>	I----------/////+++++			
8	0829:14.01	1	3	1	2	1			<VV>	I----------/////+++++			
8	0829:34.01	0	3	2	2				<CV>	I----------/////+++++			

CURVE	MINIMUM	MAXIMUM
MIN	1.6000E-02	6.5128E 00
AVG	1.0655E 00	2.3179E 01
MAX	1.3016E 00	2.0628E 02

Figure 15-25. User I/OS Duration Plot

TSS REDUCTION, VERSION 7.2 - 1.000, 30 SEPTEMBER 82

DISTRIBUTION COLLECTED ON SYSTEM DSCC4 AT 0818:33.201 FRI 82-05-21. INITIAL GMC TAPE #1469

IN SMC WAIT (DURATION) PLOT-15

-= MIN /= AVG += MAX

USERS	TIME	WAITS IN INTERVAL (HS) <1	<2	<4	<8	<16	<32	>32	LINE	WAIT TIMES (HUNDREDTHS SEC) 1.0E 00 1.0E 01 1.0E 02 1.0E 03 1.0E 04
19	1016:34.01									
19	1016:54.01	1	0	0	0	0	1		<CW>	*/////////////++++
19	1017:14.01									
19	1017:34.01									
19	1017:54.01									
21	1018:14.01									
21	1018:34.01									
21	1018:54.01									
19	1019:14.01									
19	1019:34.01									
20	1019:54.01									
20	1020:14.01									
19	1020:34.01									
19	1020:54.01	0	0	0	0	0	0	1	<CW>	I333333333333333333333333
19	1021:14.01	1	0	0	0	0	0	1	<CW>	*/////////////////////+++++
19	1021:34.01									
18	1021:54.01	1							<CW>	*
19	1022:14.01	0	0	0	0	0	0	1	<CW>	I33333333333333333333333333333333
19	1022:34.01	0	0	0	0	0	1		<TX>	I3333333333333333
21	1022:54.01									
21	1023:14.01									
24	1023:34.01									
22	1023:54.01									
22	1024:14.01									
22	1024:34.01	0	0	0	0	0	0	1	<CW>	I333333333333333333333333
22	1024:54.01									
23	1025:14.01									
23	1025:34.01									
23	1025:54.01									
22	1026:14.01									
23	1026:34.01	3	0	0	0	1	3		<CH>	*/////////////////++++
23	1026:54.01									
24	1027:14.01									

CURVE	MINIMUM	MAXIMUM
MIN	2.9210E-01	1.4286E 02
AVG	2.9440E-01	1.4286E 02
MAX	2.9440E-01	1.4286E 02

Figure 15-26. In System Master Catalog (SMC) Wait Plot

TSS REDUCTION, VERSION 7.2 - 1.000, 30 SEPTEMBER 82

DISTRIBUTION COLLECTED ON SYSTEM DSCC1 AT 0818:33.201 FRI 82-05-21. INITIAL GMC TAPE #1469

INTERTRACE TIME (ADJUSTED FOR TSS GEWAKES)

INDIV. NUMBER	CUMUL. NUMBER	CUMUL. PRC	INDIV. PRC	MILLI-SEC	PERCENT OF OCCURRENCE 00 10 20 30 40 50 60 70 80 90 100	HISTOGRAM 5
972506	972506	99.881	99.881	0- 63	XXXX

The reports that the Formatted Dump pass produces are:

- o Formatted Dump (figure 15-28) (no ID #)
- o Intertrace Time (Histogram) (figure 15-29) (ID #9).

15.5.1 TEARS Response Pass Reports

15.5.1.1 TSS Reduction Event Log. The report is a chronological listing produced as the tape is processed (figure 15-2). It indicates control changes of the Response routine, actions of the TSS Executive, and major events for individual users. For developmental use, the report also shows reduction anomalies; they are not of concern to the user. The report gives an overview of the physical makeup of the data collection period.

Each event record is formatted on one line with a time-stamp at the leftmost margin. Next to the time-stamp, for events which are related to the processing of a GMF trace type 74, the number of the trace (counting from the start of the tape) is recorded. These values can be used to locate an event on other reports or formulate timeframes for subsequent reductions. The third field is the Trace 74 subtype number. The fourth is the two-character line ID of the terminal or the four-character DRUN ID (without the suffix letter "D"). (At logon, DRUN identification is not available; the field is left blank). The next field is the name of the subsystem in control of the user at the time of the event. Following this information is a textual explanation of the event.

Examples of items on the report are:

Response control messages:
- Start of reduction tape
- Start of timeframe
- End of timeframe
- End of reduction tape.

TSS executive actions:
- Grow/reduce TSS core size
- Grow/reduce UST area
- Monitor TRACE ON/OFF.

Major user events:
- Logon/logoff
- Abnormal disconnect
- Connect command
- JDAC command.

Reduction anomalies (for developmental use):
- Program stack synchronization problems
- I/O conflicts.

```
TSS REDUCTION, VERSION 7.2 - 1.000, 30 SEPTEMBER 82

                    DISTRIBUTION COLLECTED ON SYSTEM DSCC1  AT 0818:43.515 FRI 82-05-21.  INITIAL GMC TAPE #1469

TOD          REC-CNT    MNEMONIC  &      #  UST#  LINE  LITERAL          USERID          NOTES
-----------  ---------  --------------  ---  ----  ----  ---------------  --------------  ----------------------------------
0818:54.014  LINE <TC> IS "C721      "  -- LOGON AT 0809:30.861
             LINE <CV> IS "PRC105    "  -- LOGON AT 0814:02.739
             LINE <DI> IS "MASB      "  -- LOGON AT 0814:41.176
             LINE <TN> IS "PRC104    "  -- LOGON AT 0816:22.622
0818:54.014          6  PROCESS .CRTSQ    89
0818:54.015          7  ENTER PROC ALLOC   61
0818:54.015          8  ENTER MEMORY ALL   62
0818:54.015          9  ENTER PROC ALLOC   61
0818:54.016         10  GEWAKE S/D BUSY    18            CPS        640000
0818:55.586         11  ENTER PROC ALLOC   61
0818:55.587         12  PERIODIC CHECK     14
0818:55.587         13  PROCESS .CRTSQ     89
0818:55.587         14  ENTER PROC ALLOC   61
0818:55.588         15  ENTER MEMORY ALL   62
0818:55.588         16  ENTER PROC ALLOC   61
0818:55.588         17  GEWAKE S/D BUSY    18            CPS        640000
0818:56.558  *      18  BUILD INPUT REC    20    1 <TN>  GEROUT STAT 000004          LFLG2 0000001000000000000
0818:56.558  *      19  SCAN COMMAND LIS   82    1 <TN>  REMC
0818:56.561         20  ENTER PROC ALLOC   61
0818:56.561         21  NEXT PRIMITIVE I   64    1 <TN>  XCALL      1 LEVELS
0818:56.562         22  NEXT PRIMITIVE I   64    1 <TN>  CALLP      1 LEVELS
0818:56.562         23  CALLP PRIMITIVE     8    1 <TN>  CMDL       2 LEVELS
0818:56.562         24  NEXT PRIMITIVE I   64    1 <TN>  EXEC       2 LEVELS
0818:56.563         25  EXEC PRIMITIVE     10    1 <TN>  CMDL       2 LEVELS
0818:56.563         26  ALLOCATOR SERVIC   60    1 <TN>  NEW TASK IN CALC          LFLAG 0000000000000000000
0818:56.563         27  TIME TYPES         80    1 <TN>                            NEW-  WAIT NORM.  OLD-  OUT OF MEM.
0818:56.564         28  ENTER PROC ALLOC   61
0818:56.564         29  ENTER MEMORY ALL   62
0818:56.565         30  ATTEMPT MEM ALLO   77    1 <TN>  PGM SIZ-      2042
0818:56.565         31  MEM MAP CHANGE     78    1 <TN>                            SYSLAL 066  SYSSIZ 004
0818:56.566         32  TIME TYPES         80    1 <TN>                            NEW-  SWAP/LOAD.  OLD-  UNKNOWN
0818:56.568         33  ENTER MEMORY ALL   62
0818:56.568         34  ENTER PROC ALLOC   61
0818:56.569         35  PERIODIC CHECK     14
0818:56.569         36  PROCESS .CRTSQ     89
0818:56.570         37  GEWAKE S/D BUSY    18            CPS        640000
0818:56.598  *      38  TIME TYPES         80    1 <TN>                            NEW-  USEFUL MEM.  OLD-  SWAP/LOAD.
0818:56.600         39  ENTER PROC ALLOC   61
0818:56.600         40  MAKE S/D ENTRY     91    1 <TN>
0818:56.601         41  GEWAKE S/D BUSY    18            CPS          640
0818:56.604         42  ENTER PROC ALLOC   61
0818:56.604         43  S/D FAULT QUEUE    88    1 <TN>  CPS           12          DRI         FAULT
0818:56.605         44  PROCESS DRI        23    1 <TN>  PASUST, IC= 000147
0818:56.605         45  DRI COMPLETE      104    1 <TN>
0818:56.606         46  MAKE S/D ENTRY     91    1 <TN>
0818:56.606         47  GEWAKE S/D BUSY    18            CPS          640
0818:56.609         48  ENTER PROC ALLOC   61
0818:56.610         49  S/D FAULT QUEUE    88    1 <TN>  CPS           58          DRI         FAULT
0818:56.610         50  PROCESS DRI        23    1 <TN>  FIIACT, IC= 002017
0818:56.611         51  ISSUE MME GEFSYE   27    1 <TN>                            GEFSY FUNC- FILE ALLOCATE
0818:56.611         52  ALLOCATOR SERVIC   60    1 <TN>  TURN ON I/O RDBLK         LFLAG 0000100000000001000
0818:56.661         53  ENTER PROC ALLOC   61
0818:56.661         54  GEWAKE S/D BUSY    18            CPS        640000
0818:56.693  *      55  GEFSYE/27 COMPLE   29    1 <TN>
0818:56.694  *      56  ALLOCATOR SERVIC   60    1 <TN>  TURN OFF I/O RDBLK        LFLAG 1000100000000001000
0818:56.696         57  ENTER PROC ALLOC   61
0818:56.696         58  MAKE S/D ENTRY     91    1 <TN>
```

Figure 15-28. Formatted Dump Report

TSS REDUCTION, VERSION 7.2 - 1.000, 30 SEPTEMBER 82

DISTRIBUTION COLLECTED ON SYSTEM DSCC1 AT 0818:33.201 FRI 82-05-21. INITIAL GMC TAPE #1469

INTERTRACE TIME

INDIV. NUMBER	CUMUL. NUMBER	CUMUL. PRC	INDIV. PRC	MILLI-SEC	PERCENT OF OCCURRENCE 00 10 20 30 40 50 60 70 80 90 100	HISTOGRAM
					I....I....I....I....I....I....I....I....I....I....I	
1077956	1077956	98.569	98.569	0- 63	IXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXXX	
7802	1085758	99.283	0.713	64- 127	I	
1743	1087501	99.442	0.159	128- 191	I	
1304	1088805	99.562	0.119	192- 255	I	
877	1089682	99.642	0.080	256- 319	I	
461	1090143	99.684	0.042	320- 383	I	
333	1090476	99.714	0.030	384- 447	I	
313	1090789	99.743	0.029	448- 511	I	
271	1091060	99.768	0.025	512- 575	I	
229	1091289	99.789	0.021	576- 639	I	
256	1091545	99.812	0.023	640- 703	I	
196	1091741	99.830	0.018	704- 767	I	
148	1091889	99.844	0.014	768- 831	I	
124	1092013	99.855	0.011	832- 895	I	
124	1092137	99.866	0.011	896- 959	I	
183	1092320	99.883	0.017	960-1023	I	
96	1092416	99.892	0.009	1024-1087	I	
88	1092504	99.900	0.008	1088-1151	I	
73	1092577	99.906	0.007	1152-1215	I	
70	1092647	99.913	0.006	1216-1279	I	
56	1092703	99.918	0.005	1280-1343	I	
75	1092778	99.925	0.007	1344-1407	I	
46	1092824	99.929	0.004	1408-1471	I	
43	1092867	99.933	0.004	1472-1535	I	
44	1092911	99.937	0.004	1536-1599	I	
43	1092954	99.941	0.004	1600-1663	I	
53	1093007	99.946	0.005	1664-1727	I	
39	1093046	99.949	0.004	1728-1791	I	
39	1093085	99.953	0.004	1792-1855	I	
27	1093112	99.955	0.002	1856-1919	I	
32	1093144	99.958	0.003	1920-1983	I	
26	1093170	99.961	0.002	1984-2047	I	
22	1093192	99.963	0.002	2048-2111	I	
22	1093214	99.965	0.002	2112-2175	I	
20	1093234	99.967	0.002	2176-2239	I	
22	1093256	99.969	0.002	2240-2303	I	
26	1093282	99.971	0.002	2304-2367	I	
13	1093295	99.972	0.001	2368-2431	I	
15	1093310	99.973	0.001	2432-2495	I	
15	1093325	99.975	0.001	2496-2559	I	
15	1093340	99.976	0.001	2560-2623	I	
7	1093347	99.977	0.001	2624-2687	I	
11	1093358	99.978	0.001	2688-2751	I	
13	1093371	99.979	0.001	2752-2815	I	
10	1093381	99.980	0.001	2816-2879	I	
4	1093385	99.980	0.000	2880-2943	I	
7	1093392	99.981	0.001	2944-3007	I	
9	1093401	99.982	0.001	3008-3071	I	
11	1093412	99.983	0.001	3072-3135	I	
10	1093422	99.984	0.001	3136-3199	I	
2	1093424	99.984	0.000	3200-3263	I	
7	1093431	99.985	0.001	3264-3327	I	
7	1093438	99.985	0.001	3328-3391	I	
4	1093442	99.986	0.000	3392-3455	I	

1093600 ENTRIES TOTAL AVERAGE = 6.97210 VARIANCE = 8474.988 STANDARD DEVIATION = 92.060

158(0%) OUT OF RANGE AVERAGE FOR THESE = 4968.54 IN RANGE AVERAGE = 6.26

Figure 15-29. tertrace Time Histogram

15.5.1.2 Response Time Logplots. The response time logplots show delay time that users experienced waiting for a response from TSS. They are:

- o Response Time For All Users (figure 15-3)
- o Response Time For Users Not Requesting More Core (figure 15-5)
- o Response Time For Users With Core Request During Line Idle (figure 15-6).

(Logplots are described in more detail in section 15.3).

The All-Users logplot shows, from time interval to time interval, the minimum, average, and maximum amount of time a user waited for a response. The user who waited longest is identified under the column LINE. An individual wait may extend over more than one interval. In this case the time value used on a plot line will be the total amount of time since the start of the line-idle state (not just the portion of time in that plot line period). This logplot should indicate periods of time when users were dissatisfied with TSS performance.

The information used on the All-User logplot is also used on two other logplots. These logplots help to indicate times when the source of response time degradation is due to swap area competition. The wait times are distributed on the second and third response time logplots on the basis of a request for core during the line-idle period (e.g., call to subsystem, or swap in). Differences in these plots (high response times on figure 15-6 and lower response times on figure 15-5) would be explained by Swap Area Competition. The three figures in this document would indicate that swap areas were not a problem. In fact, the only terminal experiencing difficulties appears to be terminal ID whose response time keeps on increasing during each interval.

15.5.1.3 Subdispatch Time Logplots. The subdispatch time logplots show delay time for subsystems in obtaining service by GCOS and TSS Executive. They are:

- o Total Time in Subdispatch Queue (figure 15-7)
- o Time in Subdispatch Queue Waiting Service (figure 15-8)
- o Processor Time in Subdispatch (figure 15-9).

The left-hand side of the three reports are the same from the "USERS" column to the "I/O-BSY" column. The "USERS" column is the count of users in the time-interval ending at "TIME". During each interval, the total number of subdispatches performed is contained in the column "#SD". The value in the "CPU" column is the percentage of total amount of time in subdispatch queue during which the processor was executing TSS user subsystem code (the remainder being queue overhead time). The percentage of these subdispatches which ended as a timer runout fault (vs. a DRL fault) are indicated in the column "TROUTS". The percentage of subdispatches executed for subsystems which, concurrently, had an outstanding I/O to the terminal is reported in

the "I/O-BSY" column. The "LINE" field indicates the line-ID of the user who contributed the datum for that line of that plot with the greatest magnitude (see below for an explanation of the data that is plotted on each of the reports).

The right-hand side of the "Total Time" logplot shows, from time interval to time interval, the minimum, average, and maximum duration of all subsystem subdispatches. The time value used as data for this logplot is measured from the event at which the TSS executive creates a "ready entry" in the subdispatch table, to the event at which the executive recogizes that GCOS has completed processing by changing the entry to a "fault entry".

The two constituents of each time value contributing to the "Total Time" logplot are segregated to produce the right-hand sides of the "Waiting Service" and "Processor Time" logplots. The first of the two values is the sum of the time that the "ready entry" waited for allocation by the CPU allocator, and the time that the "fault entry" waited to be detected by the TSS executive. The second value is the actual processor time used by the subsystem (between "ready entry" and "fault entry" states).

15.5.1.4 CPU Times Allocated to TS1-Exec and TS1-Subdispatch Report. The CPU time allocation chart (figure 15-4) shows the distribution of processor time attributable to the TS1 Executive, to the TSS subdispatch mechanism, to all non-TSS customers, and to the idle state. This report is produced only if the CPU Monitor (CPUM) was in execution when the data tape was written. Data collection for each print line is independent of other TEARS procedures, however the trigger to print a line is synchronized with the response time logplots. Thus, the accumulation of CPU data into this report will "get ahead" of other TEARS reports when a large gap exists between TSS traces. Here's why:

> Time driven reports, including this one, are based on TSSM trace times alone, not on CPUM trace times. If a large gap exists between TSSM traces, one or more successive "times-to-print" may expire without any printing. During the TSSM trace gap, however, CPUM traces continue to be generated and accumulated into this report. When TEARS eventually finds the long delayed TSSM trace time stamp, it triggers plot lines at successive plot intervals until it catches up with the time stamp. At the first such trigger, all the accumulated CPU data is included in the print line; successive triggers will, of course, find no additional CPU data to print. The mechanism just described explains (1) the blank lines in the CPU Time Allocation report, and (2) the greater-than-expected CPU times immediately preceeding a blank line. The mechanism was deliberate so that gaps in TSS service are obvious.

Because this mechanism differs from the other TEARS time-driven reports, and to fully describe the scope of the data for each print line, the first three columns of the report are devoted to time data. In Figure 15-4, our sample report, the first item shows the expected "time-to-print" (PLOT TIME

REFERENCE). The next two columns show the time of the last CPU trace processed in the previous print line (FIRST T70 TIME) and the last CPU trace processed in the current print line (LAST T70 TIME) -- that interval, first to last, expresses the actual elapsed time over which the CPU data pertains. The next item is the number of CPU traces on which the data is accumulated (NUMBER OF T70S). The right half of the report is a tabular chart showing the CPU time, in seconds, attributed to the TS1 Executive (TS1-EXEC), to the TSS subdispatch mechanism (TS1-SDSP), to the idle state (IDLE-TIME), and to all non-TSS consumers (OTHER-CPU). The last numeric column is a total (TOTAL) of the preceeding four columns. The graphic to the extreme right shows the deviation of the total CPU time from that expected based on the plot interval. Each "+" to the right of the axis ("I"), shows a 10% deviation "greater than expected". An asterisk replaces the "+" when the deviation is 100% or greater. Deviations in the other direction are shown as "-" symbols to the left of the axis.

For a normally functioning system under moderate load, the number of CPU traces approximates 1 for each 2 CPU-seconds. If the plot interval is so small or CPU activity so light that fewer than about five CPU traces are collected for each print line, the report will show false abnormality. Thus, an "abnormal" report may result from (1) low CPU activity, (2) too short a plot interval, or (3) gaps in TSS service.

15.5.1.5 TSS Subtraces Encountered, and Derails Executed by Users. The two summary histograms "TSS Subtraces Encountered" (figure 15-10) and "Derails Executed by Users" (figure 15-11) are incorporated in the pass to show the presence and volume of those traces/events within a timeframe.

15.5.2 TEARS Emulation Reports. The TEARS Emulation phase is both trace driven and elapsed-time driven; the distinction also separates two categories of reports for the user. Reports driven by trace events correspond with mimicking TSS - they give evidence of user difficulties and delays or of TSS Executive difficulties and delays. The reduction program remembers, on a short term basis, what each TSS user was attempting to do, the obstacles encountered in the attempt, how long the attempt was ongoing, and some of the computer resources used in the attempt. Report entries are made when some threshold, usually based on time delay, is exceeded or when a service denial, failure, or error thwarts the user. The program also measures the frequency with which the TSS Executive did its work - a report entry follows a delay threshold.

Reports driven by elapsed time measure several TSS functions without regard to any particular user -- they report measures of the subdispatch mechanism, memory activity, and the user disk I/O functions among others.

Several reports participate in a superior-subordinate arrangement; only selected entries in the superior report cause an entry in the subordinate report. The Users' Map is subordinate to the Exception Message Report and the TSS Core Map is subordinate to the TSS Delay-User Delay Report. They

are event driven reports. One elapsed-time driven report, the Memory Activity report, is subordinate to the logplot of Intertrace Gap but only for synchronization convenience.

15.5.2.1 Exception Message Report. When a request for some system service, TSS or GCOS, is denied or terminates with an error, an entry is made in this Exception Message Report (figure 15-12). This report also controls (triggers) item entries in the Users' Map, section 15.5.2.2, for several of the errors/denials. The entries are categorized by function/origin into four classes: TSS, FMS, Derails, and the inevitable Miscellaneous.

An additional entry identifies those users who consume excessive CPU resources (no denigration of the user is intended, nor should you infer it). The intent of the entry is (1) to correlate an apparent excessive response time pattern with a user engaged in CPU intensive work (with occassional messages to his terminal), and (2) to identify magnitude of such work.

A contrived example of the Exception Message Report is given in figure 15-12. It provides one example of each possible error message. Reading from the left, you can see the time (TOD) and sequence number (RECNO) of the trace which captured the denial/error. Time of day is represented as HHMM:SS.sss. The current trace sequence number is relative to the first TSS trace following GMC start. This is followed by the message proper, shown under the heading EXCEPTION MESSAGE. We list the user's terminal ID and his USERID, if applicable, followed by the error message. Error messages are taken directly from Honeywell publications, and some are cryptic. For excessive CPU usage, an exception is declared when the accumulated CPU time, since swapping in, exceeds 64,000 clock pulses; the CPU time is checked each time a user is removed from the subdispatch fault queue.

To reduce repetitious entries for the same user, such entries are inhibited so as to print not more than once per second. The graphic to the right of an entry shows one dot, up to 30, for each inhibited "print" of the previous entry for that user. An asterisk following the time of day indicates that an accompanying entry in the Users' Map report has been made.

15.5.2.2 Users' Map. This report (figure 15-13) is driven by the Exception Message Report for several exception entries. Recall that in the parent report (section 15.5.2.1), the triggering entry is marked with an asterisk following the time of day. The same entry from the parent report is again printed on the Users' Map as the reason for the map. Following the reason is the users' map proper. It displays, in inverse core location order, all current TSS users and their status. We show the user sequence number (USR), his full address (USTOCT) in octal, his line/terminal identification (LID) and his user identification (USER-ID). The user sequence number is a TEARS-generated subscript into a set of program tables, while LID shows the line ID written either as id| for terminal sessions, or as NNNN for DRUN sessions (where NNNN is the DRUN number). Following the user identification

is user state data -- how long the user has dwelled in his current state (TIME-IN) in milliseconds and his current state (STATE) in mnemonic form (viz, table 15-1). If terminal I/O is in progress, a "T" (for true) is printed in the next column (TRMIO). We then show the number and mnemonic of the last recorded trace (LAST-TRACE) and the last recorded derail (DERAIL) for the user. The following items report the last recorded program size (SIZE) and an indicator that core is currently assigned ("T" for true). The last item shows the entire program stack (PROGRAM-STACK); an asterisk following the top of stack denotes a user in build mode.

15.5.2.3 TSS Delay - User Delay Report. This is the most complex TEARS report (figure 15-14). It is a time-ordered merge of two separate event histories so that user delays rooted in TSS delays become obvious. It also reveals some TSS delays caused by user service requests, the most abundant being security package requirements (.FS049 during DRL DEFIL). This report also controls (triggers) item entries in the TSS Core Map, section 15.5.2.4, for users waiting core too long.

From the left in figure 15-14, the columns TOD, RCNT, and CUR are common to both TSS delay and user delay entries; they show the current time of day, the sequence number of the current TSS trace, and the subtype of the current TSS trace. Time of day is represented as HHMM:SS.sss. The current trace sequence number is relative to the first TSS trace following GMC start, and the subtype tells which of the 105 traps caused the trace. An asterisk following the current trace number identifies a trace performed in courtesy call. The remainder of the report describes a TSS delay, a user delay, or a concurrency of both types of delay. In figure 15-14, examples of all three-type entries appear. To highlight specific entries in figure 15-14 for further discussion, those entries are enclosed in braces and marked with a key letter (A, B, ...). Reference to these specific entries is made by a phrase such as "Key Z...".

Key A is an example of TSS delay reporting - it shows a 1-second delay (DLAY) between trace subtypes 18 and 61 (PRV, CUR), where the delay is printed as rounded seconds. Subtype 18 is a "Gewake with subdispatch busy" and subtype 61 is an "Enter processor allocation." TEARS program logic identified the subtype 18 trace as the "cause" by executive delay (XEC); this means that the TS1 executive gave up the processor with a GEWAKE and was delayed at least 1 second longer than the GEWAKE time. Thus the delay was external to TSS, and cannot be further localized. The indicators of external-to-TSS delays are the following subtypes (PRV):

15, "Gewake - no users"
17, "Gewake until subdispatch done"
18, "Gewake with subdispatch busy"
72, "Issue GEMORE for memory"
102, "Issue remote inquiry GEROUT".

and (CUR):

61, "Enter processor allocation"
89, "Enter routine to process queue".

The central portion of the report entry (columns DRL through USR-I/O) is a tableau that links functional areas with last known trace subtypes for any user. The functional areas are the column headers and represent derail service (DRL), File Management Supervisor service (FMS), and disk service (USR-I/O). The functional areas are ones in which TSS delays have been observed or are possible; column entries are the trace subtypes associated with such a TSS delay. If a linking is made, it is only a suggestive or circumstantial cause of delay. Several disparate linkings are possible, as is a failure to make any linking at all. The possible linkings are as follows:

For derails (DRL), the unique entry is trace subtype 23, "Process derail."
File Management Supervisor delays (FMS) are associated with two trace subtypes - subtype 27, "Issue MME GEFSYE" and subtype 45, "Call .MFS19 for file grow."
Users requesting disk I/O (USR-I/O) have caused TSS delays with trace subtype 24, "Request file I/O", and subtype 41, "Load subsystem with DRL RESTOR."

The right half of the report (Key A remains the referent) is a snapshot of all current users. The items reported for each user are "as of" just after the current trace was processed. Each user is identified by user sequence number (UST), terminal identification (LINE), and user identification (USERID). Following this, the time the user has dwelled in his current state (DLAY) is shown as rounded seconds and the current state (STATE) is printed. Following the state, an asterisk in column T marks a user engaged in terminal I/O. The next two columns display the last recorded trace subtype (PRV) and the last recorded derail type (DERAIL) for the user. Next, the last known program size (SIZE) is dimensioned as words-of-memory/1024. The last item shows the top three (of a possible five) members of the user's program stack (PROGRAM STACK). An asterisk following the top of stack denotes a user in build mode.

Notes:

(1) In this example, Key A, the total time between the listed traces was 1 second (rounded); program logic detected the delay as excessive because the time between listed traces was at least 1 second (not rounded) longer than the GEWAKE time.

(2) In single processor systems, the delay would not have been detected, since GEWAKE time is not available. (Program logic would have taken the GEWAKE time to be very large. The logic could as well have taken the GEWAKE time to be very small. We chose the former alternative to eliminate false alarms).

(3) It is possible that TSS traces performed in courtesy call appeared within the interval bracketed by "PRV,CUR"; they are ignored within (the TSS delay part of) this report.

Key B is an example of user delay reporting. This type entry is made only for users who are delayed in a vulnerable state (refer to table 15-1). Further, an entry is made for a user only when a trace attributable to that user is the current trace. Except for the state "waiting for memory" (mnemonic WTMEM), a delay is detected when state occupancy exceeds 1 second. For the state "waiting for memory", a delay is detected when state occupancy is 1 second longer than the wait time computed on program size.* The items reported for a delayed user are "as of" just before the current trace was processed. The data items are the same as described above for the user snapshot portion of TSS delay entry.

Key C is an example of user delay reporting while waiting memory. The asterisk following the time-of-day indicates that an accompanying entry in the TSS Core Map Report has been made.

Key D is an example of a concurrent TSS delay and user delay where TEARS logic failed to identify the cause (.FS049 in DRL DEFIL) but inspection can. The trace subtype which caused the problem was a "derail" (trace subtype 23) and the derail was a DEFIL. After the DRL DEFIL began, a "terminal I/O complete" (subtype 87) in courtesy call was trapped, and that trap caused TEARS logic to "forget about" the derail although the state was preserved. The clues to the interpretation are the values under user state, DRLSR, the last known derail, DEFIL, and the fact that the user delay and TSS delay are the same. (Logic error may be corrected in delivered version.)

Key E shows two user delay entries for the same user in the same state, but with a user option in force. To preclude repetitious entries for a user in the same state, a user option (option(onchange)) was selected here. It inhibited all but the first print following delay detection and the last print just prior to state change. The plus sign following the state indicates (a) the user's state will change as a result of the current trace and (b) the delay in this state for this user has been printed once before (but for a lesser delay time).

15.5.2.4 TSS Core Map. This report is driven by the TSS Delay - User Delay Report logic when it detects a user waiting memory too long. Recall that in the parent report (section 15.5.2.3) the triggering entry is marked with an asterisk following the time-of-day. The same entry from the parent report is again printed on the TSS Core Map as the reason for the map. Following the reason is the core map proper. It displays, in core location order, the

* TS1 computes an alarm time as a function of each user's program size. It will take no extraordinary action to obtain memory for a user until the time has expired. TEARS uses the same function (plus 1 second) to detect a delay when waiting for memory.

allocation of all the user reserved swap area. The allocation printout is condensed so that each line describes one user in core along with any free memory contiguously above him.

Reading from the left in figure 15-15, find the user sequence number (USR), and the starting address of his swap area (BASE) relative to TS1. The amount of core assigned (SIZE) is followed by the unused memory (FREE) contiguously above it. If no user occupies the lowest core slot, the first print line shows that as a HOLE in core. The remaining items of the print line describe the user and what he is doing.

The user is identified by his line identification (LID) and user identification (USER-ID). The next three items are timers dimensioned in milliseconds: the users's time in current state (TSTAT); time since last completed "swapping in" to core (TCORE); and accumulated CPU time (TCPU) since last completed swapping in to core. The next item is a measure of disk resource use -- since disk connect time is not available from within TSSM, we collect and print instead the number of disk I/Os (DIOS). This is followed by the mnemonic for the last recorded derail (LSTDRL), the entire program stack (PROGRAM STACK), the complete address of the user's UST (USTOCT) in octal, and the user's current state (STATE). If terminal I/O is in transmission, a "T" (for true) is printed in the next column (TRMIO). The last item is the mnemonic of the last recorded trace (LAST TRACE) for the user.

Notes:

(1) Core allocated to a user as Extended Buffer Memory is indicated as such by printing "---ebm---" in place of the program stack (PROGRAM STACK).

(2) The items TCORE, TCPU, and DIOS for a user just swapping in may be misleading -- because a user swapping in has core assigned as shown in the core map, but he is not considered in core. The three items mentioned are "as of" the completion of his previous swap into core.

The Core Map must be interpreted in conjunction with its parent to avoid wrong inference. In particular, a long delay waiting memory may be rooted in a TSS delay and thus is not a sign of memory allocation trouble. Also, not every excessive wait for memory (shown in the parent) results in a core map entry; to avoid redundancy, the core map print is completely inhibited if TEARS logic shows no change of core allocation since the last printed map. Further, the option described in section 15.5.2.3 that limits printing to "first and last occurrence" also limits the core map.

15.5.2.5 Error Message Report. The Error Message Report identifies conflicts between how the TSS works and how the TEARS designers expected it to work. Errors reported here are not fatal to the reduction pass and will not invalidate analysis. Nevertheless, the errors should be reported to

C751 for program correction. The report format is identical with the Exception Message Report described in subsection 15.5.1.1. Figure 15-16 is both a sample report and a list of the known errors as of the version number and compilation date shown.

15.5.2.6 Intertrace Gap - Duration. This logplot is an attempt to graph the time between TSSM traces while subtracting out the expected duration of any TSS Gewakes. The goal is a graph in which large gaps represent periods of delayed service and not periods in which TSS was idling by design. To avoid contamination of the data by traces performed in courtesy call during the periods TSS was idling in Gewake, all courtesy call traces are excluded. Figure 15-17 is an example of this report. It plots the minimum, average and maximum time between traces that are not performed in courtesy call. Those gaps which include a Gewake are corrected by subtracting out the expected Gewake time. The plot also shows the number of users, the time of the plot line, and seven counts indicating the number of gaps in seven ranges of time. For example, the 4 column lists the number of gaps that were less than 4 milliseconds; the 8 gap, between 4 and 8 milliseconds, and so on. The same data set used here is also source for the histogram Intertrace Time described in subsection 15.5.2.16.

For single processor systems, the expected Gewake time is not available - the gap time following a Gewake is taken as zero.

15.5.2.7 Memory Activity Report. The TSSM subtraces 63 and 65 through 75 form the bulk of exceptional TSS Executive service in support of memory requirements for the user. Figure 15-18 is an example of this report. It is a time-driven tabular report; it is printed in synchrony with the logplot of intertrace times for convenience alone. Each print line shows the count of each of the traces 63 and 65 through 75 that were trapped during the plot interval. The traces are:

63, "Enter swap decision processing"
65, "Consider TSS size increase"
66, "Initiate size increase for urgent user"
67, "Set up fence for urgent user"
68, "Set up urgent user class memory reserve"
69, "Force swap"
70, "Terminate force swap"
71, "UST area increase by 1K"
72, "Issue Gemore for memory"
73, "Gemore successful"
74, "Memory release"
75, "Gemore refused or reduction not possible".

Together with the two logplots described in sections 15.5.2.11 and 15.5.2.12, the Memory Activity report provides a provisional and incomplete picture of TSS memory allocation service.

15.5.2.8 Eligibles for Subdispatch - Duration. This report plots the minimum, average and maximum wait-subdispatch-time (WTSUB state) in milliseconds for all users. The WTSUB state begins when TEARS logic determines that the user is eligible for subdispatch; it ends when the user is placed in the subdispatch queue. Each line of the logplot is a window of a length determined by user option (default is 20 seconds). The logplot shows the number of users, the time of the plot line, and the line ID of the user with the maximum time. It also shows seven counts indicating the number of waits in seven ranges of time. For example, the 1 column indicates the number of waits that were less than 1 millisecond. The 2 column indicates the number of waits that were between 1 and 2 milliseconds, etc. Figure 15-19 is an example of this report.

15.5.2.9 In Subdispatch - Duration. This report plots the minimum, average and maximum subdispatch time (SUBDS state) in milliseconds for all users. The SUBDS state begins when the user is placed in the subdispatch queue; it ends when TS1 finds the user in the fault queue. Each line of the logplot is a window of a length determined by user option (default is 20 seconds). The logplot shows the number of users, the time of the plot line, and the line ID of the user with the maximum subdispatch time. It also shows seven counts indicating the number of subdispatches in seven ranges of time. Figure 15-20 is an example of this report.

15.5.2.10 Eligible for and in Subdispatch - Duration. This report plots the minimum, average and maximum combined time of wait for subdispatch (WTSUB state) and in subdispatch (SUBDS state) for all users. An example of this report is given in figure 15-21. This logplot defines the length of a combination event in three ways. When a wait-subdispatch is followed by a subdispatch, the length of the event is from the start of the wait to the end of the subdispatch. If a wait-subdispatch is followed by something other than a subdispatch, the length of the event is from the start of the wait to the end of the wait. If a subdispatch is preceded by something other than a wait-subdispatch, the length of the event is from the start of the subdispatch to the end of the subdispatch. Each plot line is identical in format to both the Eligibles for Subdispatch, and the In Subdispatch logplots.

15.5.2.11 User Swaps - Swap Rate. This report plots the rate of core swapped, in K-words per second, for all users. Core to core swaps are not included. Each plot line is a window of a length of time determined by user option (default is 20 seconds). The logplot shows the number of users and the time of the plot line. The rate is defined as the total swapped core divided by the window interval. This logplot is not keyed on users' states but on swap events. Figure 15-22 is an example of this report.

15.5.2.12 User Swaps - Swap Amount. This report plots the average and total core swapped, in K, for all users. Each plot line shows the number of users, the plot line time, the line ID of the user with the maximum core swapped, and seven counts indicating the number of swaps in seven block

sizes of core. This logplot is not keyed on user states but on swap events. Figure 15-23 is an example of this report.

15.5.2.13 User Swaps - Duration. This report plots the minimum, average and maximum swapping time in hundredths of a second and does not differentiate between swap in (SWPIN) states and swap out (SWPOU) states. The state begins when TS1 initiates the swap; it ends when the swap is complete. Along with the number of users, the time of the plot line, and the line ID of the user with the maximum swapped time, this logplot shows seven counts indicating the number of swaps in seven ranges of time. Figure 15-24 is an example of this report.

15.5.2.14 User I/Os - Duration. This report plots the minimum, average and maximum time all users were in disk I/O (DSKIO state) in hundredths of a second. Each plot line shows the number of users, the plot line time, the line ID of the user with the maximum I/O time and seven counts indicating the number of I/Os in seven ranges of time. Figure 15-25 is an example of this report.

15.5.2.15 In System Master Catalog Wait - Duration. This report plots the minimum, average and maximum time all users were in SMC wait (WTSMC state) in hundredths of a second. Each plot line shows the number of users, the plot line time, the line ID of the user with the maximum SMC wait time and seven counts indicating the number of SMC waits in seven ranges of time. Figure 15-26 is an example of this report.

15.5.2.16 Intertrace Time (Adjusted for TSS GEWAKES). This report is a histogram of the time between TSSM traces that are not performed in courtesy call. The intertrace time is corrected for Gewakes as explained in section 15.5.2.6. Data for the report is presented in rounded milliseconds. A sample report is given as Figure 15-27.

15.5.3 Formatted Dump Reports

15.5.3.1 Formatted Dump. The formatted dump is a debug or investigative tool for the technical user. It decodes most of the data fields which comprise each trace, correlates UST-bearing traces with the corresponding line ID and prints a one-line expansion of the trace. Were the report unlimited, one GMC data tape would generate enormous output. To filter this mass of data, the user options TIMEFRAME, TRACES, and LINES may be used as described in subsection 15.6.

Since the data content of the 105 trace 74 subtypes varies so widely, a generalized column descriptor heads the report. From left to right, each trace expansion shows: the current time-of-day (TOD) as HHMM:SS.sss; the sequence number of the current TSS trace (REC-CNT) beginning with the first TSS trace after GMC start; the name or mnemonic and subtype number (MNEMONIC & #) identifying the current TSS trace. If the trace bears UST information, the two columns, UST # and LINE, uniquely identify the user to which the

trace is attributable -- the UST # is an internal subscript into a set of program tables, while LINE shows the line ID written either as id| for terminal sessions, or as NNNN for DRUN sessions (where NNNN is the DRUN number). The columns headed LITERAL and NOTES may contain decoded interpretations of various data fields within the trace -- in general, 12-character numerals and numerals with leading zeros are octal except that the data item lflg2 (i.e., .LFLG2) is binary. Any trace performed in courtesy call shows an asterisk following the current time-of-day. The subtype 90 trace is not decoded directly -- after a "90 sequence" does occur (and also at the start of a formatted dump), a list of current users is shown in line-item format:

current-time LINE id IS "user id" --LOGON AT logon-time

The subtype 105 trace contains too much data for a 1-line expansion. Currently, it is only cursorily expanded - LITERAL shows the trace A-register.

15.5.3.2 Intertrace Time. This report is a histogram of the time between TSSM traces. Data for the report is presented in rounded milliseconds. The report is of identical format to that described in section 15.5.2.16; a sample report is not repeated here.

15.6 Reduction Run-Time Options

The general card format for an option request is as follows:

COMMAND(PARAMETER-LIST)

where COMMAND is a keyword specifying what action is to be taken and PARAMETER-LIST, if required, provides data to accomplish the action. The COMMAND is interpreted through its sixth character, so the input must match the first six characters of a valid command. Shorter commands must match exactly.

Fourteen commands are recognized by the card input interpreter. The list below shows each command, the action to be taken, and in parentheses, the default implications:

- o CPRINT — Print all time driven reports at a specified repetition period. (Print at 60-second intervals for response pass reports; 20 seconds for emulation pass reports).
- o DEBUG — Activate the debug statements in one or more modules. (All debug statements are inactive).
- o DUMP — Reserved for development.

- o HISTOGRAM — Modify a histogram report. (Histogram report parameters are as shown in table 15-2).

- o LINES — Restrict reduction to include or exclude specified line/DRUN ids. (No restrictions).

- o NEW — Prepare to accept another, or repeat, set of GMC data tapes and another set of option alteration commands for a following reduction run. (Reduction is complete when the current set is reduced).

- o ON — (1) Activate the reduction routines for a specified pass over the GMC data tapes. (Response time pass is active; emulation and formatted dump passes are inactive).

  (2) Turn on one or more specified reports. (All reports are on).

- o OFF — Turn on one or more specified reports. (All reports are on).

- o OPTION — Enforce a special option. (All special options are inactive).

- o PLOT — Modify a logplot. (Plot parameters are as shown in table 15-3).

- o STOP — Stop reduction after processing a specified number of physical tape records. (No limit).

- o TIMEFRAME

  (1) Accept one to five time windows for overall data reduction or for specified reports. (Data reduction and report output are not time limited).

  (2) Accept one to five time windows for debug output. (Debug output is not time limited).

- o TITLE — Accept a new title to be printed on all reports. ("TSS REDUCTION, VERSION 7.2 - 1.000, 1 JULY 82" is printed as the banner).

- o TRACES — Restrict reduction to include or exclude specified TSSM subtrace types. (No restriction).

15.6.1 Parameter List. The parameter list is a sequence of data items separated one from the next by a field of spaces. Any one space in the field may optionally be replaced by a comma or colon. A data item may be an alphanumeric, a quoted string, an integer, a real number or a null. Strings

Table 15-2. Histogram Default Parameters

REPORT ID	TRAILER FLAG	# INTERVALS	LOW VALUE	INTERVAL SIZE
1	OFF	54	1	1
2	OFF	54	54	1
7	OFF	39	-10	1
8	OFF	40	28	1
9	ON	54	0	64

Table 15-3. Plot Default Parameters

REPORT ID	# POINTS TO PLOT	Y-MINIMUM	ORDERS-MAGNITUDE
4	UNLIMITED	1.0	3.0
5	UNLIMITED	1.0	3.0
6	UNLIMITED	1.0	3.0
10	UNLIMITED	0.1	6.0
11	UNLIMITED	0.1	6.0
12	UNLIMITED	0.1	3.0
13	UNLIMITED	0.1	5.0
14	UNLIMITED	0.1	4.0
15	UNLIMITED	1.0	4.0
16	UNLIMITED	0.1	5.0
17	UNLIMITED	1.0	3.0
18	UNLIMITED	1.0	4.0
19	UNLIMITED	0.1	4.0
20	UNLIMITED	0.1	4.0
21	UNLIMITED	1.0	5.0

are enclosed in single or double quotes. Real numbers include the standard FORTRAN forms: for example, 12.34, 1234E-2, 1.234+1 are all accepted as the same number. A null is two consecutive commas or colons; it stands for a missing data item and for the separator fields on either side of the missing item.

The data acronym RID stands for a one- or two-digit integer specifying the report identification number of the report to be affected. Report identification numbers are shown in section 15.5. The specifications "integer-W", "alphanumeric-X", "real-Y", and "string-Z" require user input of an integer number in a field at most W wide, an alphanumeric in a field at most X wide, a real number in a field at most Y wide, and a quoted string in a field at most Z wide (not including the quote delimiters), respectively. Alphanumerics and strings exceeding the specified field width are accepted but are truncated to the right. Integers and reals exceeding the specified field width are rejected as illegal data. Integers may not contain a decimal point. Reals must contain a decimal point or an exponent. Reals are limited, for all data input, to a maximum of 14 characters including sign, decimal point and exponent. Reals without exponent are further limited to a maximum of 11 characters, including leading sign and decimal point. Numeric data items shorter than the specified field width are acceptable in all commands, but a null item is acceptable only where provided for.

15.6.2 Command Syntax. Each command with its parameter list must be complete on a single card, but more than one command may be placed on each card if spaces separate them. The parameter list may either be contiguous with the command keyword or spaces may intervene. The syntax for each command is given below.

TEARS permits reduction of the GMC data tapes in one of three phases: response time analysis, emulation analysis, or formatted dump analysis. The response time analysis routines are on, and the other phase routines are off, by default. The following commands may be used to exercise control over the reduction:

- o DEBUG Activate all debug; if emulation phase, load special debug link.

- o DEBUG(MODULE-NAME-1,...,MODULE-NAME-N)
  Activates the debug statements in the explicitly named modules, where MODULE-NAME-i, alphanumeric-6, is the FORTRAN name given to the subroutine, function or link. The modules currently nameable and the intent of the debug action, in parentheses, are shown below.

  | BUGDMP - additional debug link for emulation phase. (Performs an interlineated formatted dump and finite-state-machine dump of the 1000 traces preceeding a fatal failure of the emulation phase logic).

| CHANGER - decodes trace data for formatted dump. (Allows comparison of TEARS decoding with trace data).
| CHKPOINT - saves state of the emulation phase finite-state-machine to random file. (Verifies action of routine).
| DOTRAON - intermediate control routine for formatted dump. (Shows correlation of user sequence number (a major TEARS data structure) with TSS UST address).
| DOTRALIMIT - not a module but a control. (Limits formatted dump to one sample of each TSSM trace subtype).
| GETOKE - lexical analysis of card input. (Verifies action of routine).
| LOGPLOT - formats and prints logarithmic bar chart. (Prints arithmetic values in lieu of bar chart).
| NXTRECRD - tape input handler. (Verifies action of routine).
| SAMSON - controls emulation phase when within a reduction timeframe. (Performs formatted dump during emulation phase).
| SETSTATE - the finite state machine of the emulation phase. (Provides additional data in case of non-fatal errors).
| STATECHART - not a module but a control in the emulation phase. (Prints one table, at each plot interval, of the state transitions of the finite state machine during the interval).
| TIMESET - sets up array of start-stop times for user-specified reduction windows. (Verifies action of the routine).
| T90 - initializes reduction variables in accordance with TSSM trace subtype 90. (Verifies action of the routine).

The one shot subroutine GMFEXC, which reads the initial tape record, is executed before card input. It may be debugged by setting the program switch word, bit 11.

- HISTOGRAM(RID, TRAILER-FLAG, NUMBER-OF-INTERVALS, LOW-VALUE, INTERVAL-SIZE)

  Turns on the histogram with report identification number RID, integer-2. Prints/does not print the histogram summary if TRAILER-FLAG is ON/OFF, alphanumeric-3. The new number of histogram intervals is NUMBER-OF-INTERVALS, integer-3. The new low value and interval size are LOW-VALUE and INTERVAL-SIZE, both integer-10. Note: In the current version, the number of intervals may not exceed 54.

The following two options are as above, but without change to the current specifications of those parameters not appearing in the list.

- HISTOGRAM(RID , , NUMBER-OF-INTERVALS, LOW-VALUE, INTERVAL-SIZE)
- HISTOGRAM(RID, TRAILER-FLAG)

- LINES(SENSE-OF-RESTRICTION, "ID-1",...,"ID-N")
  Excludes from, or includes in, the formatted dump which only those lines/DRUNs listed. Exclusion is obtained when the SENSE-OF-RESTRICTION is DONTDO, alphanumeric-6; inclusion, when the SENSE-OF-RESTRICTION is DOONLY, alphanumeric-6. Up to 10 lines or DRUNs may be listed in 1 or several commands, but if the SENSE-OF-RESTRICTION changes, previous commands are discarded. Terminals/lines are specified by a string "ID", string-2; DRUNs, by a string "NNNN", string-4. String lengths must be exact.

- NEW("TAPE-NUMBER")
  Flags TEARS to stop reading card input options and to proceed with data reduction. It also informs the program that another data reduction pass is to follow, and that the following pass is to lead off with the tape number specified by "TAPE-NUMBER", string-6. For that following reduction, the default options will prevail unless modified by option alteration cards following the NEW command. This command must be the last command on a card; commands following on the same card are discarded. Note: TAPE-NUMBER may optionally be of the form alphanumeric-6, NEW(AJ283) for example.

- OFF Turns off all switchable reports.

- OFF(RID-1,...,RID-N)
  Turns off the reports specified by report identification numbers RID-i, integer-2.

- OFF(CATEGORY)
  Turns off all logplots or histograms depending on CATEGORY: PLOTS or HISTOGRAMS, both alphanumeric-2.

- ON ON options use the same parameters as those shown above for the OFF command, with appropriate reversal of intent.

- ON(PHASE-OF-REDUCTION)
  Turns on, or engages, exactly one of the three phases of reduction depending on the value of PHASE-OF-REDUCTION: RESPONSE, EMULATE, or FORMATDUMP, all alphanumeric-6. The response phase is on by default; all others are off. Only one phase may be accomplished on a single pass over the data tapes.

- OPTION(CHOICE)
  Modifies action of the reduction program depending on the following values of CHOICE, each alphanumeric-6:

| ALLCORE
Expands the two emulation phase plots that depict User Swap amount and rate so as to include core-to-core movement. The expansion is imperfectly implemented and the option is not recommended.

| LOGSCALE
Inserts intermediate numeric values on the y-axis for all logplots.

| MINUSDISKIO or -DISKIO
Adjusts individual response times of the response phase by subtracting the time spent accomplishing each disk I/O in the line-idle period.

| MINUSSUBDISPATCH or -SUBDISPATCH
Adjusts individual response times of the response phase by subtracting the CPU time in subdispatch expended during the line-idle period.

| NOABORT
Reserved for development.

| ONCHANGE
Limits the trace driven report "TSS Delay - User Delay" of the emulation phase to print user delay items not more than twice: first, when the delay is initially recognized, and second when the delay is about to end.

| TK4014
Reserved for development.

- PLOT(RID, NUMBER-OF-POINTS, Y-MINUMUM, ORDERS-OF-MAGNITUDE)
  Turns on the logarithmic plot with report identification number RID, integer-2. The new number of points to plot is NUMBER-OF-POINTS, integer-10, maximum. The y-axis is bounded below by Y-MINIMUM, real-14, and above by Y-MINIMUM * 10**ORDERS-OF-MAGNITUDE, where ORDERS-OF-MAGNITUDE is real-14. (The exponent is truncated to integer before use; if that integer is not between 1 and 7 inclusive, it is arbitrarily set to 7.).

The following two options are as above, but without change to the value of those parameters not specified.

- PLOT(RID, NUMBER-OF-POINTS)
- PLOT(RID , , Y-MINIMUM, ORDERS-OF-MAGNITUDE)

- STOP(RECORD-COUNT)
  Stops the data reduction pass after reading RECORD-COUNT, integer-10, physical records, maximum. One fewer records will be processed.

- TIMEFRAME(RID, START-TIME-1, STOP-TIME-1,...,START-TIME-N, STOP-TIME-N)
  Turns on report with report identification number RID, integer-2. Accepts up to five start/stop pairs:
  START-TIME-i, STOP-TIME-i where each component of the pair is real-14. Times are input in psuedo-decimal form:
  HHMM.SSsss
  where HH is the hour in 24-hour format, MM is the minutes of the hour, SS is the seconds of the minute, and sss is the fraction of seconds. Up to five start/stop pairs may be specified for a given report. The count of five may be achieved in one or several commands, however, the lexical order must conform with the temporal order to get the results wanted. The final pair in a command may be abbreviated -- if only a start time is given, the stop time is permanently unbounded. If RID is zero, the time frame(s) bound the entire reduction, but the status of individual reports is unaffected. Histograms may not be individually delimited with the command.

- TIMEFRAME(DEBUG, START-TIME-1, STOP-TIME-1,...,START-TIME-N, STOP-TIME-N)
  Accepts up to five start/stop pairs:
  START-TIME-i, STOP-TIME-i
  where each component of the pair is real-14, to further delimit the action of debug statements. The limiting action becomes effective as soon as the program determines "what time it is". This command must be accompanied by some form of the DEBUG command or no debug output will be produced. The time frames are independent of, though not unaffected by, any other user selected timeframes.

- TITLE("NEW-TITLE")
  Accepts the string "NEW-TITLE", string-63, as the new banner to head all reports.

- TRACES(SENSE-OF-RESTRICTION, TRACE-1,...,TRACE-N)
  Excludes from, or includes in, the formatted dump only those trace subtypes listed. Exclusion is obtained when the SENSE-OF-RESTRICTION is DONTDO, alphanumeric-6; inclusion, when the SENSE-OF-RESTRICTION is DOONLY, alphanumeric-6. As many traces as required may be specified by TRACE-i, integer-3, in one or several commands; however, previous commands are discarded if the SENSE-OF-RESTRICTION changes from DONTDO.

15.6.3 Examples. The following are samples of the user optional commands:

NEW(1234)	Begins a fresh run with tape #1234 as lead tape.
DEBUG	Turns all debug statements on.
DEBUG(TIMSET)	Turns on debug statements in subroutine timeset.
HISTOG(1 ON)	Turns trailer flag on for histog 1, also turns on 1.
HISTOG(2, , 43, 1, 2)	43 buckets, lowval is 1, interval size is 2, report is on.
CPRINT(123)	123 seconds between periodic reports & plot lines.
LINES(DOONLY,"A4","T5","3225")	Format dump only lines with IDs a4|, t5| and the DRUN with ID 3225. Multiple "lines" are cumulative. Command also applies to debug during the analysis phase.
ON(1, 2 3)	Turns on reports 1, 2 and 3. (All reports are on by default).
PLOT(10, 200)	Plot 10 will stop after 200 points.
OPTION(-SUBDISPATCH)	Subtract users subdispatch time from his computed response time.
OPTION(MINUSDISKIO)	Subtract users disk I/O time from his computed response time.
OPTION(ONCHANGE)	Limits the TSS Delay - User Delay Report.
STOP(1234567)	Stop after processing 1234567 physical records.
TRACES(DONTDO, 1 9 78 43 66)	Do not format traces 1 9 78 43 and 66.
TRACES(DOONLY, 12)	Format only trace 12. "Dontdo" may be preceeded by a "doonly". Multiple "traces" are cumulative. Command also applies to debug during the analysis phase.
TIMEFR(0, 0200.)	Data reduction starts at 0200:00 with no stop time.
TIMEFRAME(10 234.58 0345.40)	Plot is bounded by 0234:58 and 0345:40.
TIMEFR(0, 0234.58, 0345.4)	Same for data reduction.
TIMEFR(DEBUG, 1142., 1150.)	If debug is on, then debug statements are bounded by 1142:00 thru 1150:00 independent of main timeframes.
ON OFF(1)	Two commands: turn on all reports except 1.
NEW(5677)	Must be last command on card.
TITLE("TSS reduction, version 7.2 - 1.000, 19 July 82")	Default title.

15.6.4 Card Input Errors. If the card input interpreter cannot understand a command keyword or data item, it will echo the entire card and mark the last correctly interpreted character position with the message "ILLEGAL COMMAND (or DATA) FOLLOWING...". Depending on the state of the interpreter, it may attempt to find the end of the incorrect command in order to pass on

to a following command on the same card, or it may reject the remainder of the card; it will report which of the two actions was taken, and will then continue with card interpretation. In no case will an action confirmed by the interpreter be reversed or "undone" by a following error.

15.7 JCL for Reduction Execution.

The data reduction procedures consist of a single FORTRAN program having a main level and multiple links.

A sample JCL deck is presented in figure 15-30. The LIMITS card should be studied to meet user needs. The run time (999) and output limit (60K) may both be altered as required by the duration of the monitoring run. As a rule of thumb, run time will be approximately one-to-one with data collection elapsed time if all three passes are made in the same activity.

TEARS requires 44K of memory in order to execute plus an additional 2K for SSA space.

If no data cards are required, the JCL card

```
$     file    05,null
```

should be used in their stead.

15.8 Multireel Processing

If the GMC collected data is on more than one tape reel, TEARS will coordinate the change with a series of messages to the computer operator. These messages first warn the operator of an impending change, then direct the change, and finally, tell the operator of the tape number error if it has occurred. You can help smooth the transition from reel to reel by including MSG2 cards in the JCL deck listing the tapes which comprise the data set. Here are the messages directed to the computer operator:

a. PLEASE GET TAPE # yyyyy. IT WILL BE CALLED FOR SOON.

   This warns of the impending change, and names the next tape in sequence (yyyyy).

b. DISMOUNT OLD TAPE; --THEN MOUNT TAPE # yyyyy ON DRIVE zzzzz.

   This directs the change, naming the new tape number (yyyyy) and the specific tape drive (zzzzz).

c. WRONG TAPE JUST MOUNTED; CAN YOU MOUNT yyyyy ON DRIVE zzzzz (Y/N).

   This message is conditional -- it is sent when TEARS finds the wrong tape has been mounted (by comparing internally generated tape

```
##a,j
$       ident   1820251/30/1878,prc-krsch
$       msg2    **********************************************
$       msg2    ** please give it a "u" on tape error.
$       msg2    ** will use tapes ja176 and jb176 twice
$       msg2    **********************************************
$       select  b29idpx0/object/tears
$       limits  999,44k,,60k
$       tape    01,x3dd,,ja176,,gmf-t90,,den16
option(-subdispatch) option(-diskio)
cprint(20)
new(ja176)
on(emulate)
option(onchange)
$       endjob
```

Figure 15-30. TEARS JCL

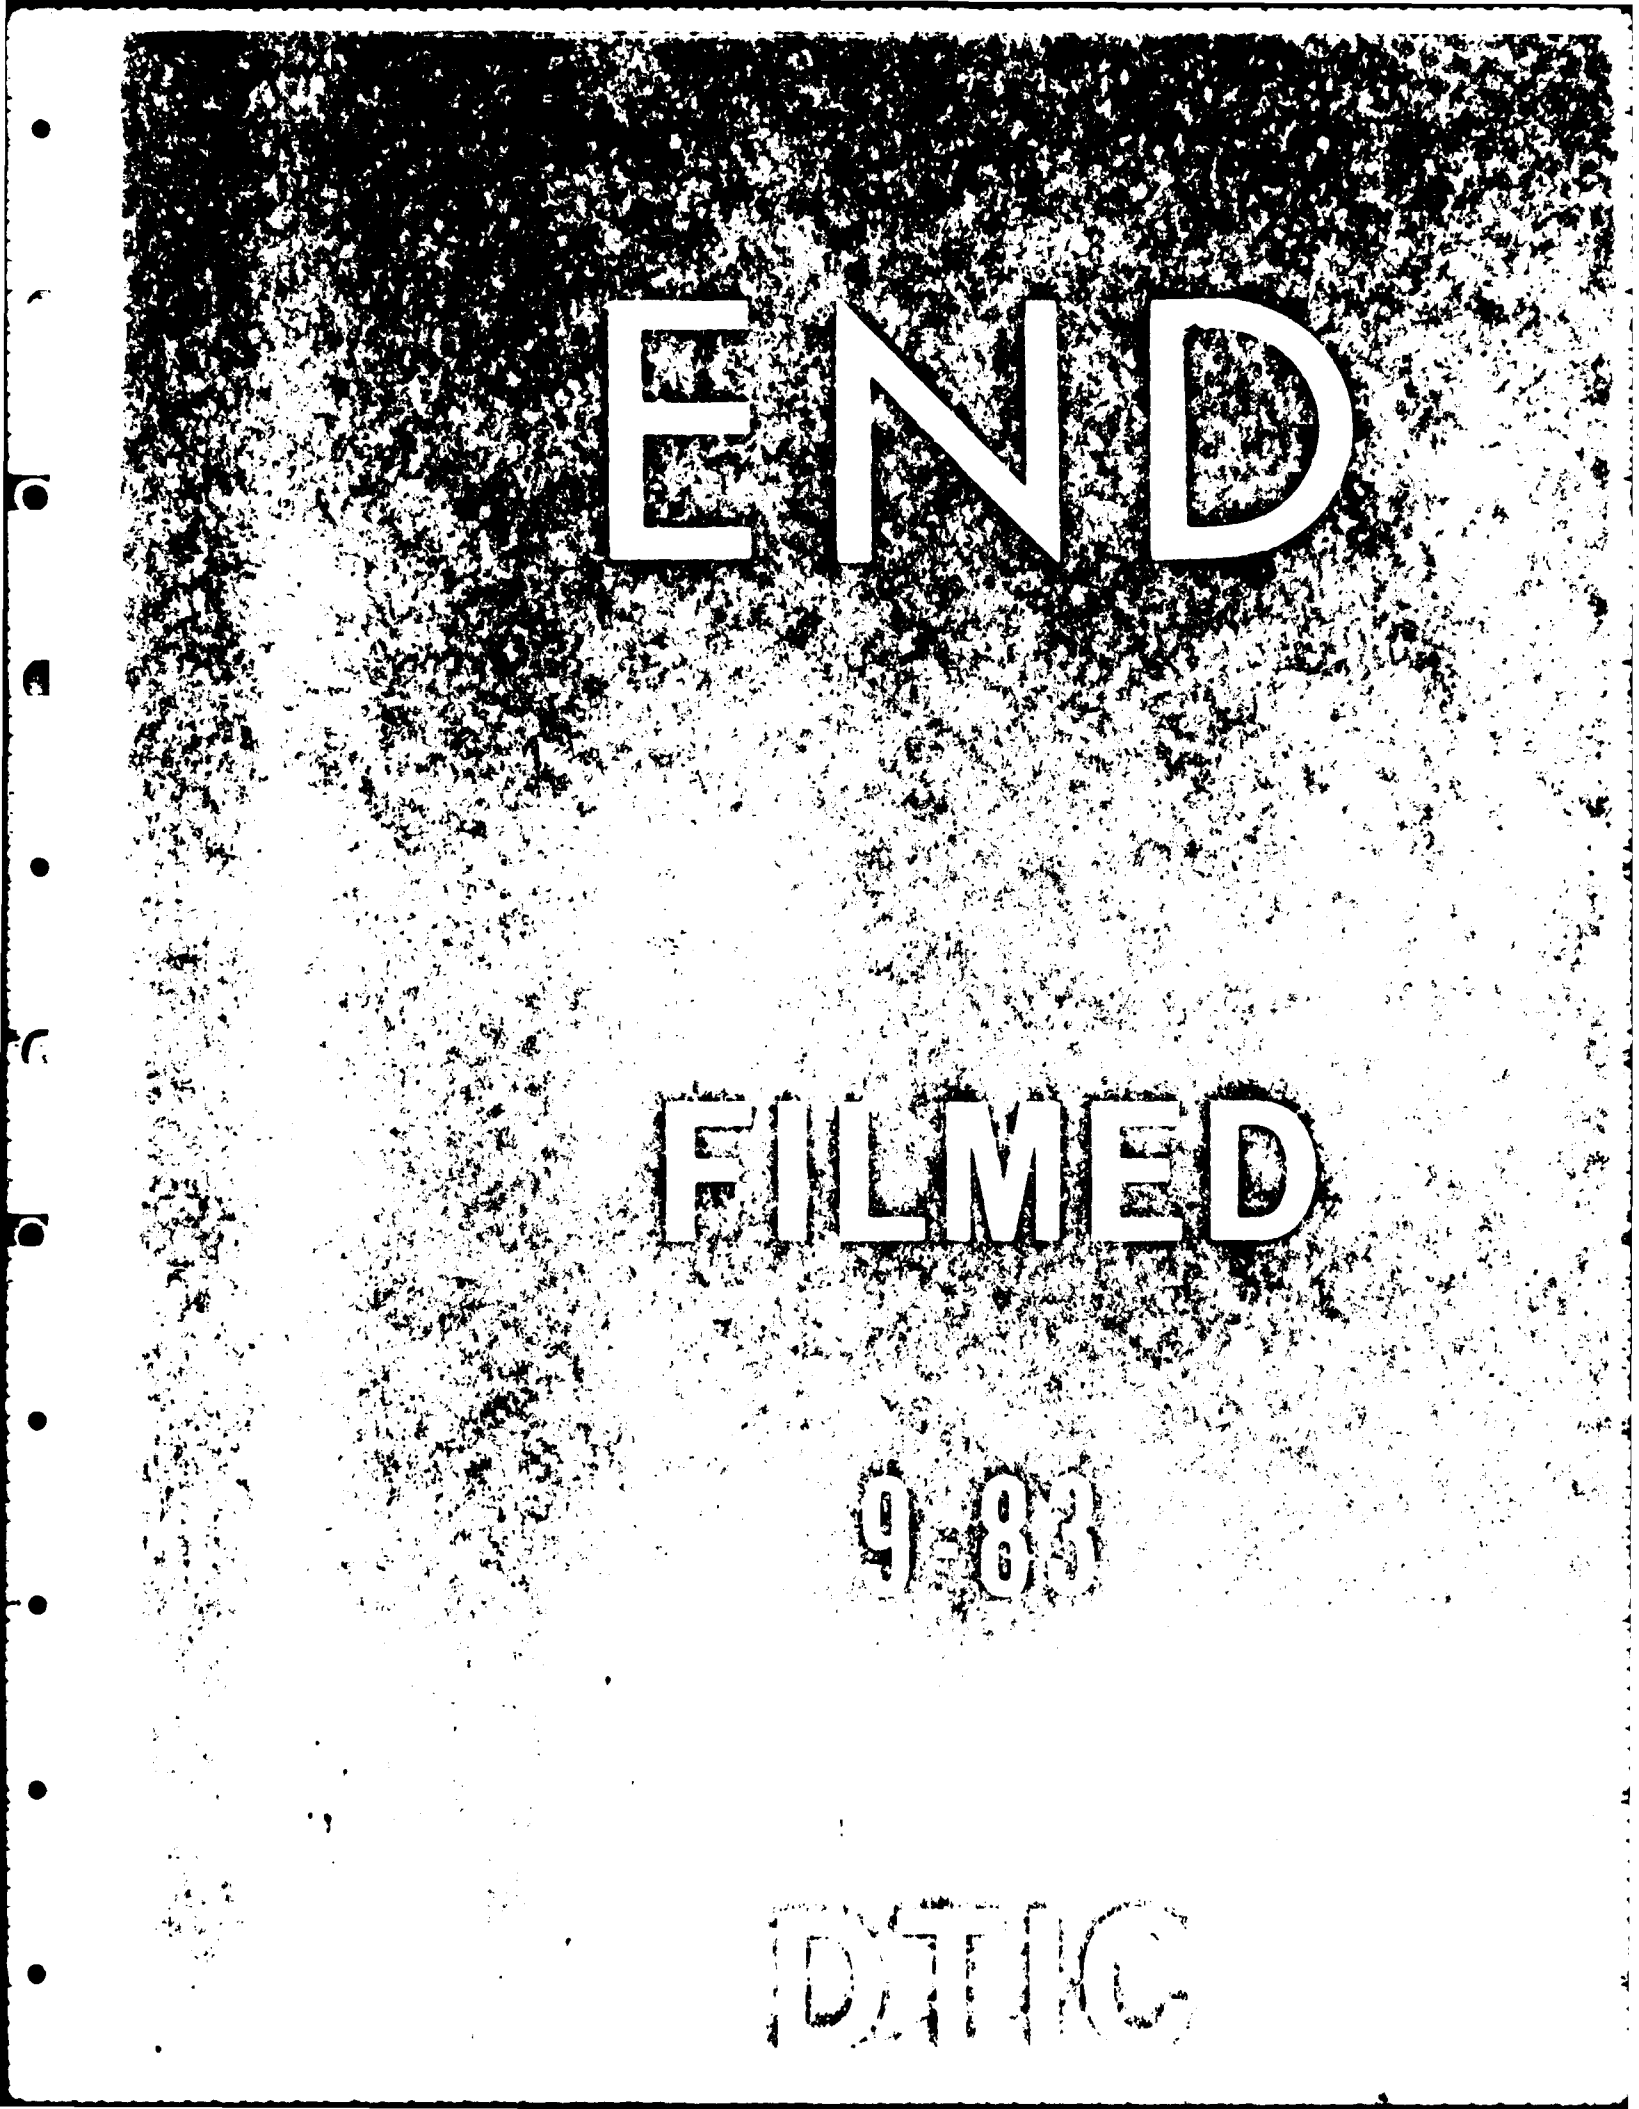
labels), or when more subtle changeover errors cause TEARS to suspect the new tape is not correct. If the operator answers affirmatively, the message sequence begins again with message (b) above. If the answer is negative, TEARS will terminate the reduction pass gracefully. Otherwise, TEARS will repeat the message.

15.9 Tape Error Aborts

During the course of processing, the operator may be required to abort the data reduction program due to an irrecoverable tape error. If such a condition occurs, the operator should abort the job with a "U" abort. This will allow the data reduction program to enter its wrap-up section and produce all reports generated prior to the tape error.

DISTRIBUTION

Addressee	
CCTC Codes	
C121 (Reference and Record Set)	2
C430	250
Defense Technical Information Center, Cameron Station, Alexandria, Virginia 22314	12
	256


END
FILMED
9-83
DTIC