

AD-A119 130

COMMAND AND CONTROL TECHNICAL CENTER WASHINGTON DC
GENERALIZED MONITORING FACILITY. CHANGE 1.(U)


F/6 17/2

UNCLASSIFIED

JUL 82
CCTC-CSM-UM-246-B2-CHG-1

NL

1-2
A
210-10





DEFENSE COMMUNICATIONS AGENCY
COMMAND AND CONTROL TECHNICAL CENTER
WASHINGTON, D. C. 20301

12

IN REPLY
REFER TO: C751

16 July 1982


TO: RECIPIENTS

SUBJECT: Change 1 to Computer System Manual CSM UM 246-82, Generalized
Monitoring Facility

1. Insert the enclosed change pages and destroy the replaced pages according to applicable security regulations.
2. A list of Effective Pages to verify the accuracy of this manual is enclosed. This list should be inserted before the title page.
3. When this change has been posted, make an entry in the Record of Changes.

FOR THE DIRECTOR:

92 Enclosures


J. DOUGLAS BOTTES
Assistant to the Director
for Administration

AD A119130

DTIC
SELECTED
SEP 10 1982
H

DISTRIBUTION STATEMENT A
Approved for public release;
Distribution Unlimited

DTIC FILE COPY

82 09 10 011

EFFECTIVE PAGES - 1 MAY 1982

This list is used to verify the accuracy of CSM UM 246-82 after Change 1 pages have been inserted. Original pages are indicated by the letter O, and Change 1 pages by the numeral 1.

<u>Page No.</u>	<u>Change No.</u>
v-xiii	1
xx-xxi	1
2-7 - 2-8	1
2-11 - 2-12	1
2-12.1 - 2-12.2	1
5-5 - 5-7	1
5-8	0
5-13 - 5-16	1
5-16.1 - 5-16.2	1
5-17	1
5-18	0
5-25	1
5-26	0
5-29	1
5-30	0
5-39 - 5-41	1
5-42	0
5-43 - 5-44	1
5-44.1 - 5-44.2	1
5-49	1
5-49.1 - 5-49.2	1
5-50	0
5-55	0
5-56 - 5-60	1
5-60.1 - 5-60.2	1
6-3	0
6-4	1
6-4.1 - 6-4.2	1
6-5	0
6-6 - 6-8	1
6-11	0
6-12 - 6-13	1
6-14	0
6-15	1
6-16 - 6-17	0
6-18 - 6-19	1
6-20	0
6-23	0
6-24	1
6-27	0

6-28
6-33
6-34 - 6-36
6-36.1 - 6-36.2
6-37
6-38
6-43 - 6-44
6-49
6-50
6-53
6-54
6-55
6-56 - 6-58
6-58.1 - 6-58.4
6-59
6-60
6-61
6-61.1 - 6-61.2
6-62
7-1
7-2
7-2.1 - 7-2.2
7-3
7-4 - 7-5
7-6 - 7-7
7-7.1 - 7-7.2
7-8
7-29
7-29.1 - 7-29.2
7-30
7-31
7-32
7-35
7-36
7-36.1 - 7-36.4
7-37
7-38
7-39
7-40 - 7-41
7-42
7-42.1 - 7-42.2
7-43
7-44
7-44.1 - 7-44.4
7-47
7-48
7-53
7-54
7-54.1 - 7-54.2

1
0
1
1
0
1
1
0
1
0
1
0
1
1
0
1
0
0
1
1
1
0
1
1
0
1
1
0
0
1
1
1
0
1
1
0
1
1
0
0
1
1

Accession For	
NTIS GRA&I	☒
DTIC TAB	☐
Unannounced	☐
Justification	
By _____	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A	

CH-1

Page No.

Change No.

8-49
8-49.1 - 8-49.2
8-50
9-23
9-24
9-27
9-28
9-29
9-29.1 - 9-29.2
9-30
9-30.1 - 9-30.4
9-35
9-36
9-36.1 - 9-36.2

1
1
0
0
1
1
0
1
1
1
1
0
1
1
1

Section		Page
5.2.7.5	Interface Requirements	5-22
5.2.7.6	Abort Codes	5-23
5.2.7.7	DATANET Monitor Software Description	5-23
5.2.7.7.1	DATANET-Host Interface	5-24
5.2.7.7.2	Monitoring of CPU	5-24
5.2.7.7.3	Host/DATANET Response Time	5-25
5.2.7.7.4	Terminal Monitoring	5-26
5.2.8	Idle Monitor	5-26
5.2.9	Transaction Processing System Monitor	5-26
5.2.9.1	TPS Trace Collection	5-27
5.2.9.2	Modifying the Transaction Processing System	5-27
5.3	Processing	5-27
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 70 - Extended	5-46
5.4.5	TM	5-46
5.4.5.1	Trace Type 50	5-46
5.4.5.2	Special Trace Type 50	5-46
5.4.5.3	Trace Type 51	5-47
5.4.5.4	Trace Type 52	5-47
5.4.6	CM	5-47
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
5.4.7.3	Special TSS Trace Type 14	5-49.1
5.4.8	GRTM	5-49.1
5.4.8.1	Trace Type 62	5-49.1
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

Section		Page
5.4.10.3	Trace Type 74	5-53
5.4.11	Special Records	5-53
5.4.11.1	Lost Data Record	5-53
5.4.11.2	Termination Record	5-54
5.4.11.3	End-of-Reel Flag	5-54
5.4.11.4	NUM Lost Data	5-54
5.4.11.5	Reconfiguration Record	5-54
5.5	GMC User Input Parameter Options	5-54
5.5.1	ON/OFF Option	5-56
5.5.2	Tape Selection Option	5-56
5.5.3	Terminal Specification Option	5-57
5.5.4	Move Option	5-57
5.5.5	CPU SNUMB Option	5-57
5.5.6	Connect Option	5-57
5.5.7	Time Option	5-58
5.5.8	Specifying High Density Tape	5-58
5.5.9	Limited Mass Store Monitor/Channel Monitor Collection	5-59
5.5.10	Request Next Data Card	5-59
5.5.11	Specifying Monitoring Requirements for the GRM	5-59
5.5.12	General Rules of the GMC Data Parameter Card	5-60
5.6	JCL for Creation of an Object File	5-60
5.6.1	Introduction to JCL	5-60
5.6.2	Creation of an Object File	5-60.1
5.7	JCL for Executing the GMC	5-62
6.	MEMORY UTILIZATION DATA REDUCTION PROGRAM	6-1
6.1	Inputs	6-1
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	6-5
6.1.5	Default Option Alteration	6-5
6.1.6	Histogram Alterations (Action Code HISTG)	6-9
6.1.7	Plot Alterations (Action Code PLOT)	6-9
6.1.8	Turn a Report On (Action Code ON)	6-12
6.1.9	Turn a Report Off (Action Code OFF)	6-12
6.1.10	Set a Time Span of Measurement (Action Code TIME)	6-12
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

Section		Page
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
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
6.3.2	System Program Usage	6-28
6.3.3	MUM Reports	6-29
6.3.3.1	Report 1 - Memory Demand Sizes of New Activities in 1K Word Blocks	6-32
6.3.3.2	Report 2 - The Memory Demand Size of All Demand Types	6-32
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 Activities Waiting Memory in Allocator Queue	6-35
6.3.3.11	Report 11 - The Time Corrected Number of Activities Waiting Memory	6-35
6.3.3.12	Report 12 - The Time Corrected Number of User Activities Waiting Memory	6-35
6.3.3.13	Report 13 - The Number of Activities Waiting Memory When a Processor Went Idle	6-36

Section		Page
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	6-36
6.3.3.18	Report 18 - The Number of Activities in Memory When a Processor Went Idle	6-36
6.3.3.19	Report 19 - The Ratio of User Activity Duration Versus Its Memory Use Time	6-36.1
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	6-37
6.3.3.23	Report 23 - The Request Size of GEMOREs	6-37
6.3.3.24	Report 24 - Not Output	6-37
6.3.3.25	Report 25 - The Memory Demand Size Versus the Memory Wait Time	6-37
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	6-38
6.3.3.29	Report 34 - The I/O Time Charged per User Activity in Seconds	6-38
6.3.3.30	Report 35 - The CP Time Charged per User Activity in Seconds	6-38
6.3.3.31	Report 36 - The Number of Times a User Activity was Swapped	6-38
6.3.3.32	Report 37 - The Total Elapsed Time a User Activity was Swapped	6-40
6.3.3.33	Report 38 - The Number of Times a System Activity was Swapped	6-40
6.3.3.34	Report 39 - The Total Elapsed Time a System Activity was Swapped	6-40
6.3.3.35	Report 40 - Number of Extra Activities That Might Fit in Memory Using Compaction	6-40
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

Section		Page
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	Peripheral Allocation Status Report	6-53
6.3.12	Plots Reports	6-53
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
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
6.4	Error Messages	6-58.3
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

Section		Page
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
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	7-42.1
7.6	Default Option Alteration	7-46
7.6.1	Monitor a Specific Device Area (Action Code AREA)	7-47
7.6.2	System Debug (Action Code DEBUG)	7-47
7.6.3	Continue Data Reduction After an Input Option Error (Action Code ERROR)	7-47
7.6.4	Specify System File Names (Action Code FILDEF)	7-47
7.6.5	End Card (Action Code END)	7-49
7.6.6	Produce the SSA Module Usage Report by Job (Action Code MODULE)	7-49
7.6.7	Record Limitation by Connects (Action Code NCONN)	7-49
7.6.8	Record Limitation by Records (Action Code NREC)	7-49
7.6.9	Turn a Report Off (Action Code OFF)	7-49
7.6.10	Turn a Report On (Action Code ON)	7-50
7.6.11	Produce Connect Summary Report by Userid/SNUMB (Action Code PROJ)	7-50
7.6.12	Reduce WW6.4 Data or Process MSMDRP on a WW6.4 System (Action Code RN)	7-50
7.6.13	Set a Timespan of Measurement (Action Code TIME)	7-50
7.6.14	Change the Time Quantum Value for the Chronological Device Utilization Report (Action Code TIMEQ)	7-52

Section		Page
7.6.15	Suppress the USERIDs (Action Code USERID)	7-52
7.6.16	Change the Time Quantum Value for the Connect Per 10 Minute Report (Action Code RATECH)	7-54
7.6.17	Turn on the Cat/File String Report (Action Code CAT)	7-54
7.6.18	Request the Connect Per 10 Minute Report for Specific User Job (Action Code RATE)	7-54
7.6.19	Limit the Processing and Output (Action Code LIMITS)	7-54
7.7	JCL	7-54
7.8	Multireel Processing	7-54
7.9	Tape Error Aborts	7-56
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

Section		Page
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
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
8.7	JCL	8-49
8.8	Multireel Processing	8-49.1
8.9	Tape Error Aborts	8-51
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

Section		Page
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
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 Report	9-29.1
9.5.14	TS1 Initial Parameter Report	9-29.1
9.5.15	Mailbox Busy Report	9-29.1
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	Terminate Input Options (Action Code END)	9-36
9.6.14	Default Options	9-36
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

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	6-31
6-13	Out-of-Range Histogram	6-33
6-14	Report 25	6-39
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	Peripheral 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
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
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	7-44.3

Figure

Page

7-25	Specific Device Area Report Card Input	7-48
7-26	Limited File Code Summary Input Card Format	7-51
7-27	Input Option TIME Card Format	7-53
7-28	MSMDRP JCL	7-55
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	8-50
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
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

2.4 System Organization

The GMF is composed of two data collectors, GMC and RMC, and associated data reduction programs. Sections 3 through 12 describe these programs. Figure 2-1 gives a system flowchart for the RMC. Figure 5-1 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 B29IDPX0. This USERID must be created with 3960 LLINKS of file space. A user restore can then be run. B29IDPX0 is subdivided into several catalogs described below:

- . GMFCOL - 595 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 - 1890 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 describe each program in detail.

- . OBJECT - 1105 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 - 18 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.

- . RMON - 345 LLINKS - This subcatalog contains all the software required to collect and reduce the data for the RMON Monitoring system. This subcatalog is further subdivided into JCL, SOURCE and OBJECT subcatalogs. The files within these subcatalogs are completely described in sections 3 and 4.

<u>FILE NAME</u>	<u>FUNCTION</u>	<u>SOURCE SIZE (LL)</u>	<u>OBJECT SIZE (LL)</u>
MUM	MEMORY UTILIZATION MONITOR DATA REDUCTION PROGRAM. REFERENCED IN CHAPTER 6.	360	207
MSM	MASS STORE MONITOR DATA REDUCTION PROGRAM. REFERENCED IN CHAPTER 7.	315	190
CM	CHANNEL MONITOR DATA REDUCTION PROGRAM. REFERENCED IN CHAPTER 8.	250	185
CAM	COMMUNICATION MONITOR DATA REDUCTION PROGRAM. REFERENCED IN CHAPTER 9.	180	110
CPU-TAPE	CPU AND TAPE MONITOR DATA REDUCTION PROGRAMS. REFERENCED IN CHAPTER 11.	389	160
GRT	DATANET 355 MONITOR DATA REDUCTION PROGRAM. REFERENCED IN CHAPTER 10.	280	155
TPETG	TRANSACTION PROCESSING DATA REDUCTION PROGRAM. REFERENCED IN CHAPTER 12.	64	70
TPEALT	AN ALTER FILE FOR ADDING TPE TRACE CODE INTO THE TPE SUBSYSTEM (NO OBJECT FILE). REFERENCED IN CHAPTER 12.	14	
TPEDUMP	A PROGRAM FOR OBTAINING A FORMATTED TRACE DUMP FROM A TPE/GMP DATA TAPE. REFERENCED IN CHAPTER 12.	38	26

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

Table 2-1. GMC Release Dependent Parameters

<u>Program</u>	<u>LINE #</u>	<u>Variable</u>	<u>Explanation</u>
GMF.TOP	90	SYS64	Used to control conditional assembly of GMC set=1 for W6.4(2H) release set=0 for W7.2(4J) release
	10220	-	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 pt 15 decimal of the dispatcher. The code being searched for is a LDAQ;STAQ;TRA0,1
	10700	-	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.
CPU.PAT	10	-	Code in this area searches for an ASA .SALT,5 Instruction in the dispatcher.
	210		In W6.4 search octal locations 1500-1700.
	230		In W7.2 search octal locations 2340-2450. In addition in W7.2 we need to find a STQ .QTOD,4 instruction
	420		between octal locations 2400-2460. For both releases we need to find 8 words of patch space. In W6.4
	720		between octal locations 3540-3740. In W7.2 between octal locations
	740		4600-5000. If not found here then
	1460		search octal locations 4150-4300 in
	1480		W6.4 or octal locations 5400-5530 in W7.2 all offsets are from the entry point of .MDISP. In addition .MFALT is searched in W7.2 for an ARL 12 instruction
	1190		between octal locations 2500-2550 (offset from the entry point). This is for gate locked timing code which is supposed to be assembled into W7.2 code.

<u>Program</u>	<u>LINE #</u>	<u>Variable</u>	<u>Explanation</u>
CAM.INIT	350	-	Beginning at 5100 octal locations from the entry point of .MDNET and continuing for 100 octal locations search for a 777777375207 instruction. This searches for # special interrupts processed code (NSIP).
	510	-	Beginning at 6700 octal locations from the entry point of .MDNET and continuing for 100 octal locations search for a 000077360003 instruction. This searches for the # of lines waiting mailbox code (R01XCT).
CAM.PAT	10	-	Code in this area searches for a LDQ M.LID,3 instruction in DNWW, followed by a ANQ =0077777,DU instruction. Octal
	140		locations 5000-6000 are searched (offsets are from entry point). Ten words of patch space must also be obtained. This patch
	370		space must be between octal locations 11100 and 11200 (offsets are from entry point). If patch space is not found then 7 words of patch space are searched for within the dispatcher. This search is
	720		performed between octal locations 3540-3740 in W6.4 and octal locations
	740		4600-5000 in W7.2 (offsets are from entry point). It should be noted that in commercial releases and WW7.2, DNWW is referred to as DNET.
MSM.PAT	10		Code in this area searches for an AOS .CRTDL instruction and an AOS .CRTBH instruction in the dispatcher. In W6.4
	390-680		and W7.2, these need to be within 300 octal locations of the label DBASE. If these instructions are not found, a search is made from octal locations 4600-5100 in W6.4 and octal locations 7164-7464 in W7.2. In addition, 8 words of patch space is needed. In W6.4 between octal
	850		3540-3740. In W7.2, between octal
	870		4600-5000. If the patch space is not found then search between octal locations
	1320		4150-4300 in W6.4, or octal 5400-5530 in W7.2. All offsets are from the entry point of .MDISP. In addition in W7.2 FMS
	1340		CACHE logic is also analyzed. See label TSFIO in routine T7 for locations required within FSIO module.

<u>Program</u>	<u>LINE #</u>	<u>Variable</u>	<u>Explanation</u>
GMP.MON	840	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.
	850	FMS2	Offset from entry point of .MFSIO pointing to the work which gives the option selection for FMS catalog cache. Used only in W7.2. Set to -15 decimal.

THIS PAGE LEFT INTENTIONALLY BLANK

Table 5-1. Required Trace Type for Each Monitor

<u>Monitor #</u>	<u>Monitor</u>	<u>Required Trace Type</u>
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)	10, 11, 21, 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*
6	GRTS Monitor (GRTM)	62*
7	Idle Monitor (IDLEM)	0, 1, 2, 3, 13, 16, 21, 22, 37, 65
8	Transaction Processing System Monitor (TPSM)	0, 1, 2, 4, 5, 6, 13, 42, 51, 65, 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 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 and require the Trace Number to be active on the \$ TRACE card.

Table 5-2. Abort Codes (Part 1 of 3)

- 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 than five characters) on data card for CPU SNUMB option.
- C3 - More than three 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 3)

- MO-M8 - 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.
- N2 - Sufficient patch space is not available in .MDISP to run the CPU Monitor. See subsection 5.2.3.
- N3 - DNWW/MDNET patch location could not be found. See subsection 5.2.6.
- N4 - Sufficient patch space is not available in DNWW/MDNET to run the Communication Analysis Monitor. See subsection 5.2.6.
- N5 - MSM patch for SSA cache analysis not found (AOS .CRTDL). See subsection 5.2.2.
- N6 - MSM patch for SSA cache analysis not found (AOS .CRTBH). See subsection 5.2.2.
- N7 - MSM patch space in .MDISP not sufficient for monitoring SSA cache. See subsection 5.2.2.
- N8 - CPU Monitor hook code for subdispatch could not be found. See subsection 5.2.3.
- NF - The Dispatcher hook code could not be found. Call CCTC/C751.
- NS - A CAM abort because it could not find NSIP (# of special interrupts) address in .MDNET.
- NR - A CAM abort because it could not find ROTXCT (number of lines found waiting mailbox) instruction.
- 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.
- 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 SIZEBUF.

Table 5-2. (Part 3 of 3)

- 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
- 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.
- TO-T8 - Required system trace is not on. The number following the T indicates the monitor having the problem.

locates that area of the dispatcher where it has been determined that the required SSA module is in the SSA Cache Buffer. The search for these instructions begins at octal location 4600 in WW6.4 or 7164 in WW7.2 and continues to octal location 5100 in WW6.4 or 7464 in WW7.2 within the dispatcher (offsets are from the entry point). 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.

The above searches do not occur if a search of ILIST finds a routine by the name of DBASE. In this case, the search starts at that address and continues for 300 octal locations. The lower half of word 7 of the dispatcher contains the address of a 10-word table, called ILIST, which points to conditionally loaded routines in the dispatcher. The first word of each routine is a BCD constant identifying that routine. During system startup, the dispatcher compresses itself to eliminate unnecessary routines (e.g. priority B if bit 18 in word 0 is not set). The search by absolute locations is done only if a search of ILIST does not find the desired routine.

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. This search begins at octal location 3540 in WW6.4 or 4600 in WW7.2 and continues until octal location 3740 in WW6.4 or 5000 in WW7.2 (offsets are from the entry point of .MDISP). 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. In this case, GMC will search the dispatcher from octal location 4150 in WW6.4 or 5400 in WW7.2 through octal location 4300 in WW6.4 or 5530 in WW7.2 after checking word 0 to insure that bit 2 has been set (offsets are from the entry point of .MDSIP).

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 than 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. 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 types (octal) 10, 11 and 21 must be enabled in the computer system boot deck on the \$ TRACE card. CPUM processes trace types 10, 11, 21, 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 type 70 is a direct transfer trace, 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. The subdispatch processing is done only when under a WW7.2 release.

The search for the ASA .SALT,5 ranges between octal locations 1500 and 1700 in WW6.4 or 2340-2450 in WW7.2 within the dispatcher. The search for the STQ .QTOD,4 instruction ranges between octal locations 2400

and 2460. 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 either abort occurs, the dispatcher code should be examined to determine if either 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 four free locations under WW6.4 or eight free locations under WW7.2 in which to create a direct transfer trace into the GMC. This search has the same ranges as that for SSA cache in MSM. If patch space is not found, an N2 abort will occur. See subsection 5.2.2 for a description of this search procedure.

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 three 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.

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.

When CM is active, it is essential that GCOS trace types (octal) 4, 7, 15, and 22 are enabled in the computer system boot deck on the \$ TRACE card. CM processes these trace types 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. In addition, all Idle Monitor traces must be active. 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. When CAM is active, it is essential that the GMC generated trace type (octal) 14 is enabled in the computer system boot deck on the \$ TRACE card. 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.

The method used by the CAM to patch DNWW/MDNET is similar to that used by the CPUM to patch the dispatcher. The CAM searches DNWW/MDNET for the LDQ M.LID,3 instruction beginning at octal location 5000 and ending at octal location 6000 (offsets from the entry point). If CAM cannot find this instruction, GMC will abort with an H3 abort. If this problem occurs, the DNWW/MDNET code should be examined to see if this instruction has been moved out of the octal range 5000-6000 due to an edit or recompile. If so, the code in CAM.PAT will need to be altered.

Upon finding this instruction, CAM then searches DNWW/MDNET patch area for 10 free locations in which to create a new system trace type 14. This search begins at octal location 11100 and continues for 100 octal locations (offsets from the entry point). If 10 free words of space are not found, then seven words of patch space are searched for within the dispatcher. This search occurs between octal locations 3540-3740 in W6.4 or 4600-5000 in W7.2 (offset from the entry point). If no space is found by either of these searches 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. Beginning at 5100 octal locations from the entry point, and continuing for 100 octal locations, CAM.INIT searches for a 777777375207 instruction. If it does not find this instruction, it will abort with an NS abort. CAM.INIT is searching for a number of special interrupts processed (NSIP). In addition, CAM.INIT will also search for a 000077360003 instruction beginning at octal location 6700 from the entry point and continuing for 100 octal locations. If it does not find this instruction, it will abort with an NR abort. In this section, CAM.INIT is searching for the ROIXCT processing (number of lines found waiting mailbox). 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

THIS PAGE LEFT INTENTIONALLY BLANK

5-16.2

CH-1

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.

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).

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. To run GRTM, GCOS trace type (octal) 62 must be enabled via the H6000 computer system boot deck on the \$ TRACE card. 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 by applying a set of alters supplied with release of the GMC software. It should be noted that in Release WW7.2 the alter cards to support the monitor are included within the standard release. If this is the case then only the following procedures need to be followed. The user must check the FMAC module to insure that variable FMON has been set to 1. The FMAC module must be recompiled and the macro library reloaded. Finally, all the GRTS modules should be recompiled. If the user is unable to find the FMON variable in the FMAC module, he should check the release bulletin to confirm whether the required alters have been included within the standard release. Upon completing this procedure the user should refer to subsection 5.2.7.4. If an old GRTS release is being edited, the user should execute the following procedure.

Files needed in order to assemble the GRTS modules are contained under catalog B29IDPX0/GMFCOL/GRT/ALTER. The object decks created from the reassembly of the GRTS modules are contained under catalog B29IDPX0/GMFCOL/GRT/OBJECT. Both subcatalogs are included on the GMF save tape. (See figure 5-2.)

B29IDPX0								
<u>GMFCOL</u>								
MAK.XXX	GMF.OBJ	PATLOOK	<u>GMF</u>	<u>MUM</u>	<u>MSH</u>	<u>CM</u>	<u>CPU</u>	<u>TM</u>
			GMF.TOP	MUM.INIT	MSH.INIT	CM.INIT	CPU.INIT	TM.INIT
			GMF.8TH	MUM.T10	MSH.PAT	CM.T04A	CPU.PAT	TM.T50
			GMF.MID	MUM.T46	MSH.REMO	CM.T22A	CPU.T70	
			GMF.MOW		MSHDOIT	CM.T07A	CPU.REMO	
							CPUDOIT	

Figure 5-2. GMC Catalog File Structure (Part 1 of 2)

- 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 H6000 to accept data)
- b. Input Accepted (input received by the H6000)
- c. Send Output (355 requesting continuation of output)

- d. Output Received (355 has received data from H6000)
- 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 (H6000 has told 355 it has data to send)
- h. Accept Direct Output, Then Input (H6000 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.

5.2.9 Transaction Processing System Monitor. The GMC Transaction Processing System Monitor (TPSM) is used to collect data on the performance of the GCOS Transaction Processing Executive (TPE)

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 M9I 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.

2

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 **,*  
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 GMCMB,DU  
TRA GMC.
```

The interface is completed by execution of the

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

is altered to look like

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

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

<u>Job Status</u>	<u>Information Collected</u>
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 .CRSN1 table

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 .MPQ04 is called; in a delayed RLSEC request, .MPQ04 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.

<u>Word</u>	<u>Bits</u>	<u>Information</u>
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.

<u>Word</u>	<u>Bits</u>	<u>Information</u>
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 than 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.

<u>Word</u>	<u>Bits</u>	<u>Information</u>
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.

<u>Word</u>	<u>Bits</u>	<u>Information</u>
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 than 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:

<u>Word</u>	<u>Bits</u>	<u>Information</u>
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.

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.

<u>Word</u>	<u>Bits</u>	<u>Information</u>
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. If an activity has a program number greater than 14 or FSTSLV, it is considered as a system program if: it is privileged (bit 18 of the .STATE is set) and if it has no J* file for SYSOUT (.SGNPA). This extension of the definition for system programs allows accumulation of processor time used primarily by copies of GEIN. Although DRL TASK jobs have no J* file, they are considered user activities because they are not privileged (the CPU monitor will accumulate processor time associated with termination of DRL TASK jobs as system CPU time since, when terminating, DRL TASK activities are privileged). 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 \$SLTA its processor time is accumulated, along with that for program number six (test and diagnostics). If a system program described above performs an initialization before it puts its SNUMB into .CRSNB, its processor time may be accumulated in the special category for miscellaneous system programs. The format for this GMC trace type record is shown below.

<u>Word</u>	<u>Bits</u>	<u>Information</u>
1	0-17	Size of record (77 or 83)
	18-26	Not used
	27-35	Trace number (octal 70)

2	0-35	Processor time for system programs not otherwise recognized (e.g. GEIN, DRL TASK termination)
3	0-35	Time for program 1 (\$CALC, core allocator)
4	0-35	Time for program 2 (\$PALC, peripheral allocator)
5	0-35	Time for program 3 (\$SYOT, SYSOUT writer)
6	0-35	Time for program 4 (\$RTIN, scheduler)
7	0-35	Time for program 5 (TS1, TS executives #1)

THIS PAGE LEFT INTENTIONALLY BLANK

5-44.2

CH-1

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

<u>Word</u>	<u>Bits</u>	<u>Information</u>
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
	0-11	Ignore
9	12-17	Status
	18-29	Input data tally
	30-35	Not used
	0-17	Slave LAL
10	18-35	Checksum
	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.
37-529	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.

<u>Word</u>	<u>Bits</u>	<u>Information</u>
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.

<u>Word</u>	<u>Bits</u>	<u>Information</u>
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 quanta
	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

5.4.8 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.

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

THIS PAGE LEFT INTENTIONALLY BLANK

<u>Word</u>	<u>Bits</u>	<u>Information</u>
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.

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.

Figure 5-3. 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) Requesting high density tape be collected.
- (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.
- (11) Specifying the monitoring requirements for the GRTM.

5.5.1 On/Off Option. This option allows the user to turn off all monitors not required for his purposes. The GMC default is to have all monitors on. 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-MF = User Developed Monitors

(See Section 13 for a discussion of user developed monitors.)

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 abort. The 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(4-9) = Specified number of reels

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.

+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

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 three 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 jobs

5.5.6 Connect Option. This option allows the user to suppress collection of tape connects or mass store connects within the CM 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.

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 Specifying High Density Tape. On tape output, the GMC will write for 1150 records, or to the end of tape mark. If a user desires to alter this procedure so that a high density write (1600 BPI 2100 records) can be performed, the following option should appear on the data card: M;. This option may appear anywhere on the data card, but must precede the time option request. The user will also need to

modify the \$TAPE card in the GMC execution JCL so as to specify 1600 BPI tape collection (see subsection 5.7).

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, TSI and \$PALC will be captured. If the user desires, he may request that data 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 Specifying Monitoring Requirements for the GRM. In order to collect GRM data, an M6 must not appear on the first data card. If the M6 is omitted from the first card, then the GRM 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)

SCH_n = Subchannel numbers associated with each HSLA entry
(n=0 to 31)

R_n = 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)

5.5.12 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 5-2. 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

THIS PAGE LEFT INTENTIONALLY BLANK

5-60.2

CH-1

Table 6-1. (Part 2 of 4)

<u>ID Number</u>	<u>Histogram Title</u>
21	Number of Activities Waiting Memory When a Processor Went Idle
22	Memory Available When a Processor Went Idle
23	Memory Demand Size Versus Memory Wait Time
24	Used in Conjunction with ID 23
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)
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

Table 6-1. (Part 3 of 4)

<u>ID Number</u>	<u>Histogram Title</u>
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
<u>ID Number/Name</u>	<u>Plot Title</u>
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 Short- fall in Peripheral Allocator Queue
28/PLOT3	Number of Activities in Core Allocator Queue vs. Number of Activities in Peripheral Allocator Queue
59/PLOT4	Average size of TSS, FTS and NCP

Table 6-1. (Part 4 of 4)

<u>ID Number/Name</u>	<u>Other Reports</u>
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

THIS PAGE LEFT INTENTIONALLY BLANK

6-4.2

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 Time-Sharing.

The remaining items that can be modified are the title and the vertical axis headers. 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. 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 requirement 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 WW6.4/2H GCOS release, an additional data card is required. This data card is not described elsewhere in this chapter. The data card should contain the letters RN:

Table 6-2. Default Values for Histograms

<u>ID #</u>	<u>Low Value</u>	<u>Interval Size</u>	<u>Number of Buckets</u>
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	4	4	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
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

<u>ID #</u>	<u>Max Size of Plot</u>	<u>Lower Plot Limit</u>	<u>Upper Plot Limit</u>
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

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)

MASTER - Define SNUMBS that are considered to be SYSTEM jobs - (all programs with a program number less than FSTSLV)

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

Card	1	A	
	2	B	
	3	C	repeat this set for each
	4	D	parameter to be changed
	2+N	E	

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.
- E = new action command (table 6-4)

Figure 6.2. HISTG Action Code Format

Several points must be stressed.

- o The set of parameter cards just described must be preceded by two data cards. The first data card contains the word PLOT and the second data card contains the ID number of the plot to be modified (see table 6-1).
- o 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 HIVAL value is selected, it should be divisible by 114.
- o 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 ID 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 ID 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 ID 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 ID 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 ID name to be affected by the time request (table 6-1). If the entire reduction (all reports) are to be time controlled, a report ID 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 ID 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

OF ACTIVITIES IN CALC QUEUE VS # IN PASC QUEUE

DELTA = 1.0

A=NCALC

B=NPALC

TIME OF DAY
0.

28.0

NUMBER WAITING
57.0

Handwriting practice lines with the letter 'Y' repeated across the page.

12:49:43.00
12:59:44.00
13:09:44.00

:

•

•

•

•

•

•

•

•

•

•

:

;

•

70

16:19:42.00

Figure 6-3. Standard Plot

Card	1	A	
	2	B	
	3	C	repeat this set for each
	4	D	parameter to be changed
	2+N	E	

A = the word PLOT
 B = plot ID number (Table 6-1)
 C = parameter control word
 D = new parameter value
 E = new action command (table 6-4)

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.

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 response to the operator action request made by GOOS in processing tape errors. If the operator performs this action, the data reduction program will abort gracefully.

Unfortunately, there are times when a tape error will cause a program abort without giving the operator a chance to respond with a "U". In these cases reports will be lost and this option will need to be used

Card 1 A
2 N M
3 B C D E ...

where

A = The word TIME

N = Report ID 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 GOS 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

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.

Immediately prior to the output of the histograms, the user will find a printout containing processing information. Included in this information is the following:

- o Printout of all input options selected by user
- o Indication of multireel tapes that are being requested and have been mounted
- o Indication of the monitors that were active during data collection
- o 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.
- o If the time frame option was used, an 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, 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

Table 6-5. File Code for MUM Reports

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)
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
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)
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)

***** THE MEMORY UTILIZATION MONITOR *****
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:

$$W = \frac{\sum_{i=1}^N \left(\begin{array}{c} \text{memory} \\ \text{available} \end{array} - \begin{array}{c} \text{demand for} \\ \text{memory} \end{array} \right) * (T_{i+1} - T_i)}{\text{TOTAL TIME}}$$

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 Peripheral 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 Allocator 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 $\frac{\text{Total Size Time Product for This System Program} * 100}{\text{Total Size Time Product for all Programs}}$

THE ELAPSED DURATION OF USER ACT IN TENTHS OF A SECOND

[illegible]

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

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 Total Memory Demand Outstanding. This report shows the sum of demand for all activities in the system including outstanding GEMOREs. 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.

For this report, an entry is made at each allocator call.

6.3.3.4 Report 4 - The Demand That Was Outstanding When a Processor Went Idle. This report is the same as Report 3, except that 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. Several of the CPU reports described in chapter 11 can be used to verify this hypothesis. IDLEM data is used to produce this report.

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.

The average value reported in this report minus the average value reported in report 3 will give a good feel for memory surplus or shortfall. A positive result will indicate a surplus while a negative result will indicate a shortfall. The MUM heading report also gives a surplus/shortfall indicator. 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.

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 3. 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 activities 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.

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 activities 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.

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 activities 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 Activities 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.

6.3.3.15 Report 15 - The Number of User Activities in Memory. The activities 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. See report 10 for additional user options in defining system jobs and user jobs.

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 activities for 5 minutes, and 8 activities for 1 minute. The average number of activities in memory is $(6+20+8)/3=11$. If we correct for time however, we get $((6)*(2)+(8)+(20)*(5))/8=100/8=12.5$ activities 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 processor 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 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.

THIS PAGE LEFT INTENTIONALLY BLANK

6- 36.2

CH-1

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 additional 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.

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.

6.3.3.24 Report 24 - Not Output. Report 24 is used within the data reduction as a buffer for the two levels of information necessary for Report 25 and is not an available report for display.

6.3.3.25 Report 25 - The Memory Demand Size Versus the Memory Wait Time. This report is the only dual-axis (averaged) histogram and displays the relationship of memory demand size and the wait time for its allocation. This report will show if the allocator or workload is biased in its services. The vertical axis represents a single-axis histogram and contains the memory demand sizes. The information to the left of the sizes is a count of the entires in each interval. The horizontal axis displays the averaged wait time for all the entires of a particular size interval. This is in contrast to the single-axis histogram which shows merely the occurrence distribution. The scale of this axis is determined from the data; the lowest and highest scale values represent the shortest and longest averaged wait times that have occurred. The time interval of each X is given as the DELTA on the report and the actual averages obtained for each entry are displayed under AVERAGE.

The histogram is derived by accumulating the sum of the wait times for each interval size and dividing by the total number of entries in that interval. This supplies the average memory wait time for that interval of demand size. The statistics shown below the histogram pertain to

the vertical axis followed by the statistics of the average wait times of the horizontal axis. The minimum and maximum times shown are those for all wait times and are not the averages.

For this report, an entry is made whenever an allocation of memory is made (refer to figure 6-14).

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.

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.

In both this report and the following report, two special SNUMBs may appear. These SNUMBs are GWAKE and WAITL. These SNUMBs will appear for any activity that was in GWAKE or waiting original core allocation during the entire monitoring session. 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), LU (Last 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.

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-XX----

with the left and right asterisks representing the upper and lower addresses of the job whose SNUMB is JOBOL, with an activity number of XX. 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:

FOLLOWING INFO MAY BE INCOMPLETE DUE TO LOST, DATA, NO EOF, OR ACTIVITY WAS ACTIVE WHEN MONITOR ENDED.

UNIT	TIME	STATUS	PROGRAM	DESCRIPTION	REMARKS
1	12:39	55	55	55	55
2	12:39	55	55	55	55
3	12:39	55	55	55	55
4	12:39	55	55	55	55
5	12:39	55	55	55	55
6	12:39	55	55	55	55
7	12:39	55	55	55	55
8	12:39	55	55	55	55
9	12:39	55	55	55	55
10	12:39	55	55	55	55
11	12:39	55	55	55	55
12	12:39	55	55	55	55
13	12:39	55	55	55	55
14	12:39	55	55	55	55
15	12:39	55	55	55	55
16	12:39	55	55	55	55
17	12:39	55	55	55	55
18	12:39	55	55	55	55
19	12:39	55	55	55	55
20	12:39	55	55	55	55
21	12:39	55	55	55	55
22	12:39	55	55	55	55
23	12:39	55	55	55	55
24	12:39	55	55	55	55
25	12:39	55	55	55	55
26	12:39	55	55	55	55
27	12:39	55	55	55	55
28	12:39	55	55	55	55
29	12:39	55	55	55	55
30	12:39	55	55	55	55
31	12:39	55	55	55	55
32	12:39	55	55	55	55
33	12:39	55	55	55	55
34	12:39	55	55	55	55
35	12:39	55	55	55	55
36	12:39	55	55	55	55
37	12:39	55	55	55	55
38	12:39	55	55	55	55
39	12:39	55	55	55	55
40	12:39	55	55	55	55
41	12:39	55	55	55	55
42	12:39	55	55	55	55
43	12:39	55	55	55	55
44	12:39	55	55	55	55
45	12:39	55	55	55	55
46	12:39	55	55	55	55
47	12:39	55	55	55	55
48	12:39	55	55	55	55
49	12:39	55	55	55	55
50	12:39	55	55	55	55
51	12:39	55	55	55	55
52	12:39	55	55	55	55
53	12:39	55	55	55	55
54	12:39	55	55	55	55
55	12:39	55	55	55	55
56	12:39	55	55	55	55
57	12:39	55	55	55	55
58	12:39	55	55	55	55
59	12:39	55	55	55	55
60	12:39	55	55	55	55
61	12:39	55	55	55	55
62	12:39	55	55	55	55
63	12:39	55	55	55	55
64	12:39	55	55	55	55
65	12:39	55	55	55	55
66	12:39	55	55	55	55
67	12:39	55	55	55	55
68	12:39	55	55	55	55
69	12:39	55	55	55	55
70	12:39	55	55		

Figure 6-16. SNUMB-IDENT Report

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 ≤ 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:

- o Line number (always 50 if out of core report is run without memory map)
- o SNUMB and activity number of job waiting memory
- o Memory demand and urgency
- o Whether the job could fit in memory if available memory were compacted (FWC)
- o Whether it could fit into an existing hole of memory (FWOC)
- o Whether the job is a new request (NEW), a swapped job (SWAP), or a job GEWAKE (GWKE)
- o 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

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 the amount of time the activity was in memory. 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 Peripheral Allocation Status Report. This report will track an activity as it proceeds through different phases of Peripheral 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 it completed that phase of allocation, the total time spent 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 Peripheral Allocation process. Following is a list of the more common phases of peripheral 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 simply 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

Only activities found to be in a PALC 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. Three 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

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, WAO24, CSE, 250, MCCANN, SECRET
2044T- 1		81				2.8		FCCCSE	1010, WA803, CSE, 264, UARRETT, UNCL
2057T- 1		64				3.4	47	FCCOSD	1010, NS0033, CSD, 283, PADILLA, UNCL
2083Q- 1		56				1.0		FCCCSE	1020, WAO24, CSE, 250, MCCANN, SECRET
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

SNUMB	MEMORY	STATUS	PERIPHERAL ALLOCATION STATUS REPORT			DEVICE TYPE	NUMBER
			TIME IN	TIME OUT	TOTAL TIME(SEC)		
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. Peripheral 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:

$$\frac{(\text{Upper Plot Limit} - \text{Lower Plot Limit})}{114}$$

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 surpluses 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.

PLOT- 2

DEFLTA= 4.0 A= CALC B=PALC

TIME OF DAY
().

MEMORY (K WORDS)

340.0

112.0

228.0

$$Y \cdots Y \cdots Y \cdots Y \cdots Y \cdots Y \cdots Y \cdots Y \cdots Y$$

12:49:21.50
12:59:22.30
12:59:22.30

— 4 —

14:19:18.40

A B

A B

A B

B

A

A

“

MAX :

0.1 :

0.3

CURVE	MINIMUM
CALC	0.
+ PALC	0.

MAXIMUM	0.1950E 03
	0.3000E 03

20

16:09:15.30

Figure 6-23. Standard Plot

UNCLASSIFIED

CCTC-CSM-UM-246-82-CHG-1

NL

2nd 2

END
DATE
FILMED
10.82
DTIC

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.

6.3.15 Distribution of Urgency Over Time Report. This reports (figure 6-27) is always produced and cannot be turned off. An entry in this report is made under the same time interval constraints 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 activities (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.

Core and peripheral allocation are extremely complex mechanisms that cater to high priority jobs. In the case of core allocation, a great amount of overhead in the form of CPU and I/O time is generated to respond to an urgent job's demand for memory. Additionally, peripherals that could be used by a routine job sit idle waiting to be allocated to a priority job. This is an acceptable practice for the occasional urgent job for which the system was designed, but when the majority of jobs are running at or above an urgency level that is considered critical, complications arise.

Core allocation defines an urgent or priority job as one with an urgency greater than or equal to decimal 32. This is considered the threshold core allocation value.

Requests for core are prioritized in a list by urgency. Each time core allocation is attempted, this list is scanned from high to low urgency in an effort to allocate for the most urgent jobs first. If a job's urgency has been raised higher than the highest urgency from the last update, it is processed immediately. An example of this is a transcript print job enabled at a 60 urgency by TLCF.

Once the most urgent job has been determined an attempt is made to find the smallest unused hole in core into which that job will fit (best fit algorithm). If a large enough core hole cannot be found, many complex swap and compaction decisions must be made. Swap/compaction is performed automatically if the urgency of the job is greater than 5. If the job's urgency equals 5 and no other jobs have been denied or bypassed, swap/compaction will be attempted. In the latter case, if there are bypassed jobs, the smallest of those will be allocated first. Swap/compaction will not be attempted for jobs with urgency less than 5.

The number of jobs that will be moved during memory compaction (regardless of the urgency of the moved job) varies depending on overall system memory size. Default values are 2 for systems of 512K or more and 10 for systems up to 512K. A site option patch is available to either prevent memory compaction altogether or to specify a value other than the defaults listed.

On the first pass of the swap/compaction process, jobs with urgencies of 0 or 1 will be eligible to swap. If the core request is in the first 64K or quadrant 0 or location 0 of a nonzero quadrant, jobs with urgency less than 63 will be eligible to swap. Jobs requesting these segments of core are TSS, \$POLTS, \$TOLTS, \$FSYS, ect.

If the core request is for a nonurgent job and core cannot be allocated using the above guidelines, the job is bypassed and the next lower urgency job is considered. If a priority job is being processed, however, further attempts will be made. In this case, a second pass is made through the available/allocated core table where jobs with an urgency less than six will be marked to swap. If this fails to free up sufficient memory, jobs with up to an urgency of six less than the requesting job will be considered for swapping. In all cases, memory compaction will be attempted unless a swap is possible, in which case a job will be swapped before compaction takes place.

Once a core hole has been found, the job is loaded into that memory area. Before exiting, the allocator checks to determine if a job has been refused core or bypassed. If so, the permanent urgency of that job, if under the priority threshold, is incremented by one. The highest urgency job still requiring core is then found. If the urgency of this job is 56 or greater, the allocation process is repeated for this job. This loop will continue as long as there is a sufficient high urgency job waiting memory after assignment of core to another job.

While high urgencies cause additional processing during peripheral allocation, they also result in unused peripherals, wasting valuable resources in both cases. Priority jobs are given precedence over routine jobs so that in some cases routine jobs will be ignored until an urgent job is allocated.

The peripheral allocator considers three separate threshold urgencies depending on current processing and job requirements:

- o THRS4-(40) - Considered the overall job urgency threshold. Special consideration is given to all jobs over this threshold as will be described later.
- o THRS3-(30) - Is the print/punch activity core urgency. THRS3 is the minimum default urgency given to jobs requiring a dedicated card reader, line printer, or card punch. If a job's own urgency is greater than THRS3, the higher urgency will be used. THRS3 is patchable via a site option.
- o THRS2-(20) - Is the core priority threshold. When PALC passes a job to core allocation, it normally uses a default urgency of 5. If the job's urgency is greater than THRS2, the higher urgency will be used.

The peripheral allocator maintains a job stack that is maintained in urgency order similar to the core request table. Upon entering PALC's main pass the job stack is scanned in decreasing urgency order for candidates for peripheral allocation. When an eligible job is found, its urgency is compared to THRS4 (40) to determine if it is a priority job. If urgent, an attempt is made to allocate the job immediately, otherwise additional checks are made.

During normal allocation of a new job, if more than 10 jobs are waiting core or, if more than 5 are waiting core and this job's core requirements exceed the largest hole in memory, it will be bypassed. However, if the job's urgency is greater than THRS2 (2), it will be considered irregardless. In addition, sieve limits are not checked for an urgent job with urgency more than THRS4 (40).

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.

6.4 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.

THIS PAGE LEFT INTENTIONALLY BLANK

6-58.4

CH-1

MEMORY STATISTICS TABLE 1 DUANG6

DATE	START	STOP	JOBS	CPU/IO RATIO	AVG ACT SIZE(R)	USER MEM AVAILABLE	SYSTEM MEMORY	EXCESS MEMORY CALC	EXCESS MEMORY PASC	RATIO OF DURATION VS MEMORY TIME	Z SLAVE MEM USED	TIME USER ACT SWAPPED
811217	1001	1130	1.48	0.336	23.72	210	294	87	86	1.150	75.757	0.010 SEC
AVG # USER												
ACT WAITING MEM												
1.000												
AVG # SYSTEM												
ACT WAITING MEM												
0.												
AVG # USER												
ACT IN MEM												
3.765												
AVG # SYSTEM												
ACT IN MEM												
8.771												
DURATION OF												
USER ACTIVITY												
70.050 SEC												
WAIT TIME FOR												
ORIGINAL ALLOCATION												
25.500												
ACT/HOUR												
(THROUGHPUT)												
177.920												
USLN ACT/HOUR												
(THROUGHPUT)												
121.770												
SWAPS PER												
HOUR												
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)	CEMORE	-99999	FTS	1	35	82
11.78417(114703)	MET	0	FTS	1	36	83
11.78578(114709)	CEMORE	-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

URGENCY	DISTRIBUTION OF URGENCY REPORT OVERTIME										
	0	5	10	15	20	25	30	35	40	45	50
14:06	0.	41.9	1.5	1.2	1.3	0.7	0.9	0.	0.1	0.	0.
14:16	0.	39.9	0.	0.	0.2	0.	5.8	0.	0.	0.	0.
14:26	0.	44.3	0.	0.	0.3	0.	2.2	0.	0.	0.	0.
14:36	0.	40.3	0.	0.	0.1	0.1	5.2	0.1	1.4	0.	0.
14:46	0.	46.5	0.8	1.1	2.3	0.1	1.9	0.	0.1	0.5	4.2
14:56	0.	40.6	0.	0.	0.1	0.	6.1	0.	0.0	0.	7.2
15:06	0.	45.3	3.1	0.	0.2	0.	3.4	0.	0.	0.5	1.1
15:16	0.	46.0	0.2	1.1	2.9	0.	5.7	0.	0.	0.	0.
15:26	0.	45.0	0.	0.	3.5	0.	7.4	0.	0.	0.	0.
15:36	0.	41.9	0.	0.	0.0	0.	7.3	0.	0.	0.4	6.3
15:46	0.	48.5	0.	0.	0.2	0.	5.5	0.	0.0	1.0	3.1
15:56	0.	48.5	0.	0.	0.4	0.	5.9	0.	0.	1.3	0.

SUMMARY OVER THE TOTAL DURATION

0.	40.2	0.	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

6-61.2

CH-1

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 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.

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:

- o Discovery of improper configurations (software or hardware), e.g., not alternating usage of dual physical channels configured on a subsystem.
- o 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.
- o Discovery of open file allocation on a device which can cause long and frequent arm movement.
- o 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) (NAME=SYSFILES)
- 10 Individual Module Activity Report (File 21) (NAME=SYSFILES)
- 11 SSA Module Usage Report by Job (File 21)
- 12 File Code Summary Report (File 23) (NAME=FILECODE)
- 13 CAT/File String Report (File 23)
- 14 Connect Summary Report by Userid/SNUMB (File 23)
- 15 Activity Summary Report (File 24) (NAME=ACTIVITY)
- 16 Device Area File Code Reference Report (File 22) (NAME=AREA)
- 17 Device File Use Summary Report (File 21) (NAME=FILEUSE)
- 18 Chronological Device Utilization Report (File 26) (NAME=CHRONO)
- 19 FMS Cache Report (File 21)
- 20 Connects Per 10 Minute Report (File 20) (NAME=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)

Table 7-1 (Part 2 of 2)

- 25 Special FTS File Access Time Report (File 42)
- 26 TSS Swap File Usage Over Time Report (File 42)
- 27 Device Seek Movement Summary (File 29)

THIS PAGE LEFT INTENTIONALLY BLANK

7-2.2

CH-1

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 (CHDRP). 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 CHDRP 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 CHDRP 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).

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 B29IDPX0/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.

$(\text{number of devices actually configured}) * 2 + 8$

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/I/O 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 produces a series of 24 reports listed in table 7-1 over which the user has limited 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 * * *
 * * * V E R S I O N 0 1 - 8 2 * * *

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
HJM	120	2.35
MSH	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

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:

$$\frac{(\text{CPU TIME Used by Monitor})}{(\text{TOTAL Elapsed Time}) \times (\text{Number of Processors})}$$

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 HCP. 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

 * * * MSM SYSTEM SUMMARY REPORT * * *

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

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



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.

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 SHUMB 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
DJ8X170202	0	0	0.
DJ8X170203	3	587	1.106
OPNSUTIL	1	3	0.006
	1	2088	3.932
DMSTA	0	0	0.
NEWRM	11	7352	13.846
8174T	1	1891	3.561
TSS	1	2224	4.108
XPALC	1		

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/
\$PASC- 1	5	1	0	3195	906/	1030	3.53/
\$SYOT- 0	17	6	0	0	512/	512	0. /
\$RTIN- 0	40	1	0	5141	3/	3	1713.67/
TS1 - 1	105	7	0	51245	3281/	3281	15.62/
\$LOGN- 0	50	1	0	1485	13/	13	114.23/
NCP - 1	87	9	5	76381	1114/	1114	68.56/
TLNT - 1	8	1	0	0	1/	1	0. /
FTS - 1	183	20	18	83149	3151/	3151	26.39/
TLCF - 1	90	1	0	2	3/	3	0.67/
7088T- 1	5	2	0	2946	143/	151	20.60/
FSPIN- 2	8	1	0	0	2/	2	0. /
MUM - 1	0	0	0	284	0/	47	0. /
SCHED- 0	0	0	0	10453	135/	149	77.43/
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 MSMGRP (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).

TIME	DEVID	PROB	DEVID	PROB	DEVID	PROB	DEVID	PROB	DEVID	PROB	DEVID	PROB	DEVID	PROB
11:35:38:00	4	0.36	2	0.24	8	0.12	11	0.08	15	0.08	16	0.08	0.08	0.08
11:36:38:00	4	0.44	19	0.24	2	0.09	16	0.06	8	0.05	11	0.05	0.05	0.05
11:37:38:00	4	0.32	20	0.17	16	0.09	2	0.07	11	0.06	12	0.05	0.05	0.05
11:38:38:00	21	0.49	4	0.07	1	0.07	15	0.07	14	0.04	9	0.04	0.04	0.04
11:39:38:00	21	0.55	4	0.07	14	0.05	9	0.04	16	0.04	15	0.03	0.03	0.03
11:40:38:00	21	0.60	4	0.10	1	0.05	2	0.05	14	0.03	8	0.02	0.02	0.02
11:41:38:00	21	0.73	4	0.09	2	0.06	1	0.03	11	0.01	8	0.01	0.01	0.01
11:42:38:00	21	0.68	4	0.09	2	0.05	1	0.03	11	0.02	14	0.02	0.02	0.02
11:43:38:00	21	0.72	4	0.11	2	0.05	8	0.02	1	0.02	14	0.01	0.01	0.01
11:44:38:00	4	0.21	21	0.21	1	0.13	7	0.11	8	0.07	2	0.06	0.06	0.06
11:45:38:00	4	0.21	15	0.18	21	0.15	7	0.07	1	0.07	2	0.05	0.05	0.05
11:46:38:00	21	0.30	4	0.21	8	0.18	1	0.10	14	0.04	16	0.03	0.03	0.03
11:47:38:00	21	0.26	4	0.23	1	0.15	8	0.08	10	0.04	2	0.03	0.03	0.03
11:48:38:00	4	0.21	17	0.18	21	0.12	1	0.09	8	0.06	9	0.06	0.06	0.06

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. | $TBS + 650 = HCM \text{ Required}$ |

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:

- o 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.
- o 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.
- o 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.
- o 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:

- o Only permanent files are allocated on RMVBL devices; i.e., no temporary files are allocated on RMVBL devices.
- o 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.
- o 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:

- o Initially use the system-supplied (default) value for n (one buffer for each four NIAST devices).
- o 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.
- o 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%

HIAST BUFFER EFFICIENCY REPORT			
NUMBER OF BUFFERS ALLOCATED	BUFFER ATTEMPTS	BUFFER BUSIES	AST TABLE IN MEMORY
12	18	0	8

AST TABLE IN MEMORY BUT WAS BUSY

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

INDIV. NUMBER	CUMUL. NUMBER	CUMUL. PRC	INDIV. PRC	DEVICE ID	PERCENT OF OCCURRENCE	REPORT
12	12	0.000	0.000	0-	0	
18947	18959	9.835	9.828	1-	1	
8206	27165	14.091	4.257	2-	2	
20269	47434	24.606	10.514	3-	3	
34033	81467	42.260	17.654	4-	4	
238	81705	42.384	0.124	5-	5	
25580	107285	55.653	13.269	6-	6	
.	.	.	.	.	.	
.	.	.	.	.	.	
.	.	.	.	.	.	

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											REPORT
					00	05	10	15	20	25	30	35	40	45	50	
0	0	0.	0.	0-	1	1	1	1	1	1	1	1	1	1	1	2
5935	5935	2.066	2.066	1-	1	1	1	1	1	1	1	1	1	1	1	
7879	13814	4.809	2.743	2-	1	1	1	1	1	1	1	1	1	1	1	
10424	24238	8.438	3.629	3-	1	1	1	1	1	1	1	1	1	1	1	
13957	38195	13.297	4.859	4-	1	1	1	1	1	1	1	1	1	1	1	
13675	51870	18.058	4.761	5-	1	1	1	1	1	1	1	1	1	1	1	
14260	66130	23.023	4.965	6-	1	1	1	1	1	1	1	1	1	1	1	
14658	80788	28.126	5.103	7-	1	1	1	1	1	1	1	1	1	1	1	
13735	94523	32.908	4.782	8-	1	1	1	1	1	1	1	1	1	1	1	
13192	107715	37.500	4.593	9-	1	1	1	1	1	1	1	1	1	1	1	
12161	119876	41.734	4.234	10-	1	1	1	1	1	1	1	1	1	1	1	
11550	131426	45.755	4.021	11-	1	1	1	1	1	1	1	1	1	1	1	
10956	142362	49.569	3.814	12-	1	1	1	1	1	1	1	1	1	1	1	
10396	152778	53.189	3.619	13-	1	1	1	1	1	1	1	1	1	1	1	
9580	162358	56.524	3.335	14-	1	1	1	1	1	1	1	1	1	1	1	
9380	171738	59.789	3.266	15-	1	1	1	1	1	1	1	1	1	1	1	
8894	180632	62.886	3.094	16-	1	1	1	1	1	1	1	1	1	1	1	
8319	189951	65.782	2.896	17-	1	1	1	1	1	1	1	1	1	1	1	
7945	196896	68.548	2.766	18-	1	1	1	1	1	1	1	1	1	1	1	
7182	204078	71.048	2.500	19-	1	1	1	1	1	1	1	1	1	1	1	
6454	210532	73.295	2.247	20-	1	1	1	1	1	1	1	1	1	1	1	
.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	
.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	.	
567	278696	97.026	0.197	49-	1	1	1	1	1	1	1	1	1	1	1	

287238 ENTRIES TOTAL AVERAGE = 16.89080 VARIANCE = 35200.476 STANDARD DEVIATION = 187.618

Figure 7-21. Elapsed Time Between Seek Report

DATA TRANSFER SIZE

[illegible]

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 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. In multiple copies of Timesharing are present, I/O activity generated by all copies will be recorded under the single entry of TSI. 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). In this 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

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 is 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, regardless 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

DATA TRANSFER SIZES FOR TSS SWAP FILES

[illegible]

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 ACCOMPLISH THE READ/WRITE OF THE FILES.
 ** THIS INCLUDES NO NETWORK TRANSFER TIME ON READS.

SNUPB	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)

7-44.1

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

CH-1

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

7-44.3

CH-1

Figure 7-24.3. Device Seek Movement Summary Report

THIS PAGE LEFT INTENTIONALLY BLANK

7-44.4

CH-1

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.

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.

<u>Device Type</u>	<u>Numbers Cylinders</u>	<u>Number Sectors/ Cylinder</u>	<u>Number Sectors/ Block (LLINK)</u>
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

Card 1 A
Card 2 B
Card 3 C D E F . . .

Where

A = The word TIME
B = Number of different times appearing on Card 3
C,D,E = 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 TSI. 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 SUMMARY. 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 SUMMARY on a card immediately following the LIMITS card.

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.

THIS PAGE LEFT INTENTIONALLY BLANK

7-54.2

CH-1

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.

8.7 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.

THIS PAGE LEFT INTENTIONALLY BLANK

8-49.2

CH-1

COL	1	8	16
\$		IDENT	1820251/30/3044
\$		SELECT	829IDPX0/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

[illegible]

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, Subsystem, Session Length, Response Time, # Inputs, # Outputs. If a terminal was logged onto a subsystem when CAM started, then there is no way for CAM to determine the subsystem being used by the terminal. In this case, the subsystem will be reported as UNKNWN. If a user logs onto a subsystem and then JDAC's to a different subsystem, CAM is unable to determine this switch. The entire user session will be reported as being on the subsystem originally logged onto. 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:

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

TERMINAL ACCEPTED DISCONNECT ALL LINES INPUT ACCEPTED ACCEPT FINAL OUTPUT ACCEPT DIRECT OUTPUT REJECT REQUEST-TEMP DISCONNECT ACK BREAK ACKNOWLEDGED ACCEPT NEW TERMINAL SEND INITIAL OUTPUT CONNECT TO SLAVE BREAK ***** END REACHED *****	OP-CODES COUNT FOR NMCC2 ON 122181 AT 16:03: 34.749 164 DISCONNECT THIS LINE 2 ACCEPT CALLS 20412 ACCEPT OUTPUT 12 OUTPUT NOT AVAILABLE 27614 ACCEPT DIRECT OUTPUT-INPUT 49 TERMINAL REJECTED 26 6000 INITIALIZED 219 TERMINAL CHARACTERISTICS CHANGE 556 DISCONNECT LINE 2 SEND OUTPUT 160 ACCEPT DIRECT INPUT 219 TERMINAL CHARACTERISTICS CHANGE ACK 58	117 4 242 1 20501 392 4 58 26 27841 20461 58
---	---	---

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 in general (i.e., averaged over all DAC subsystems), 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.
USER	- average number of users 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 prevalent 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 INDEX ARRAY." This means that an internal array has been filled. This is usually the terminal ID array. The parameter SMAX must be increased to enlarge the required arrays. The current size of SMAX is 100. This can be exceeded if there are a large number of users on the system when the CAM is started. If a terminal is logged on when the CAM is started, the terminal is logged on to subsystem UNKNWN. When this terminal disconnects and reconnects, it is now logged on to a valid subsystem and a different entry is made as this is now a complete terminal session. All complete terminal sessions are collected in one entry, but any unusual session required a separate entry. 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 (see figure 9-16.1).

9.5.13 Abort Report. This report indicates each time the DN355 aborts and is reinitialized. It also presents each time line IDs 01,02 disconnect. These disconnects indicate a WIN problem since WIN has lost its network connection (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 total number of special interrupts that have occurred during the time frame and the average number of unanswered interrupts, requests waiting mailboxes and lines waiting mailboxes. A line is produced for each datanet (see figure 9-16.4).

THIS PAGE INTENTIONALLY LEFT BLANK

9-29.2

CH-1

RESPONSE TIME REPORT FOR NMCC ON 030382

TOD	SYSTEM		TSS		USER		RESP	I/R		I	#	0	USER	RESP	I/R	FTS	I	...	OPREJT	OPREJP	RJM	RRSVP
	I/R	#	I/R	#	I/R	#		I/R	#													
12:38:01	0	659	7	13	1	8	0	1	1				1	1	0							
12:43:01	0	835	10	15	2	2	0	2	1				2	1	0							
12:48:01	0	849	6	18	2	43	1	4	1				4	1	6							
12:53:01	2	604	6	18	2	32	2	4	1				4	1								
12:58:02	2	469	7	18	9	14	1	4	1				4	1								
13:03:02	2	247	4	18	39	5	2	4	1				4	1								
13:13:03	1	976	6	17	7	33	1	4	1				4	1								

Figure 9-16. Response Time Report

H6000 - DN355 REJECT REPORT FOR NMCC ON 030382

TERMINAL ID	RJM	RSVP	ECHO	OPCRJT	OPCRJP	COMMANDS
D7	0	1	0	2	0	0
O1	0	0	0	27	0	0
EG	0	3	0	1	0	0
DM	0	19	0	3	0	0
DE	0	3	0	0	0	0
B4	0	0	0	138	0	0
B4	0	5	0	6	0	0
D5	0	3	0	0	0	0
O2	0	9	0	1	0	0

Figure 9-16.1. H6000-DN355 Reject Report

ABORT REPORT FOR NMCC ON 030382

TERMINAL ID 01 ON SUBSYSTEM UNKNOWN DISCONNECTED AT 13:05:19
TERMINAL ID 02 ON SUBSYSTEM UNKNOWN DISCONNECTED AT 13:05:19
TERMINAL ID 02 ON SUBSYSTEM NCP DISCONNECTED AT 13:05:24
TERMINAL ID 01 ON SUBSYSTEM NCP DISCONNECTED AT 13:05:27
***** ABORT 355-0 AT 13:28:35 AT RDCNT 144
TERMINAL ID 01 ON SUBSYSTEM NCP DISCONNECTED AT 13:28:35
TERMINAL ID 02 ON SUBSYSTEM NCP DISCONNECTED AT 13:28:35
***** ABORT 355-0 AT 13:28:35 AT RDCNT 144
355-0 RE-INITIALIZED AT 13:32:14 AT RDCNT 149
TERMINAL ID 02 ON SUBSYSTEM NCP DISCONNECTED AT 13:32:32
TERMINAL ID 01 ON SUBSYSTEM NCP DISCONNECTED AT 13:32:32
TERMINAL ID 02 ON SUBSYSTEM NCP DISCONNECTED AT 13:44:32
TERMINAL ID 01 ON SUBSYSTEM NCP DISCONNECTED AT 13:44:32

Figure 9-16.2. Abort Report

TS1 INITIAL PARAMETER REPORT

SIZE PARAMETERS

INITIAL TS1 MAX SIZE	180K	SIZE GROWTH	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	1200LINKS	MIN SWAP FILE GROWTH SIZE		300LINKS

ALLOCATED DEVICES

FILE	DEVICE/FILE NAME
#D	DQ1
#P	
#Q	
.D	.D
SS	
#S	
#T	

Figure 8-1. TS1 Initial Parameter Report

MAILBOX BUSY REPORT FOR NMCC2 ON 062582					
TIME OF DAY	DN	SPECIAL INTERRUPTS	UNANSWERED INTERRUPTS	REQUESTS WAIT MBX	INES FOUND WAITING MBX
11:10:37	0	22	0	0	0
11:10:37	1	13	0	0	0
11:11:37	0	5525	0	0	0
11:11:37	1	239	0	0	0

Figure 9-16.4. Mailbox Busy Report

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 GMP 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 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.

9.6.14 Default Options. The default options for the variable input are as follows:

<u>ACTION CODE</u>	<u>Option</u>	<u>Default Value</u>
TIME	Timeframes	None, total tape processed.

DELTA	Delta Time-frames	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

9-36.2

CH-1