

ID-A249 194



Digital Hardware
Architecture Implementation
Annual Report


Best Available Copy



COLSA CORPORATION

6726 Odysseus Drive

San Jose, California 95130



92-10143





**Digital Hardware
Architecture Implementation
Annual Report**

DTIC
ELECTE
APR 28 1992
S D D

**Sponsored by:
U.S. ARMY STRATEGIC DEFENSE COMMAND**

15 April 1992

**Prepared for:
United States Army Strategic Defense Command
GBR Project Office
P.O. Box 1500
Huntsville, AL 35807-3801**

**Prepared by:
COLSA CORPORATION
6726 Odyssey Drive
Huntsville, AL 35806**

**This document has been approved
for public release and sale; its
distribution is unlimited.**

**This report has been prepared by COLSA Corporation as part of Contract
Number DASG60-91-C-0006 for the U.S. Army Strategic Defense Command.
The views, opinions, and/or finding contained in this report are those of the
author(s) and should not be construed as an official Department of the Army
position, policy, or decision unless designated by other documents.**

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

a. REPORT SECURITY CLASSIFICATION Unclassified			1b. RESTRICTIVE MARKINGS			
a. SECURITY CLASSIFICATION AUTHORITY			3. DISTRIBUTION / AVAILABILITY OF REPORT			
b. DECLASSIFICATION / DOWNGRADING SCHEDULE						
4. PERFORMING ORGANIZATION REPORT NUMBER(S) CI-2071-92-0012			5. MONITORING ORGANIZATION REPORT NUMBER(S)			
a. NAME OF PERFORMING ORGANIZATION COLSA Corporation		6b. OFFICE SYMBOL (If applicable)	7a. NAME OF MONITORING ORGANIZATION			
c. ADDRESS (City, State, and ZIP Code) 6726 Odyssey Drive Huntsville, Alabama 35806			7b. ADDRESS (City, State, and ZIP Code)			
a. NAME OF FUNDING / SPONSORING ORGANIZATION US Armt Strategic Def. Cmd		8b. OFFICE SYMBOL (If applicable)	9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER			
c. ADDRESS (City, State, and ZIP Code) GBR Project Office P.O. Box 1500 Huntsville, AL. 35807-3801			10. SOURCE OF FUNDING NUMBERS			
			PROGRAM ELEMENT NO.	PROJECT NO.	TASK NO.	WORK UNIT ACCESSION NO.
1. TITLE (Include Security Classification) Digital Hardware Architecture Implementation Annual Report - Unclassified						
2. PERSONAL AUTHOR(S) Lance Cooper, George Paramithas, Max Patin, Stuart Venters						
3a. TYPE OF REPORT		13b. TIME COVERED FROM 04-91 TO 04-92	14. DATE OF REPORT (Year, Month, Day) 92-4-15		15. PAGE COUNT	
6. SUPPLEMENTARY NOTATION						
7. COSATI CODES			18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)			
FIELD	GROUP	SUB-GROUP	Ground Based Radar			
			Signal Processing			
			Parallel Computing			
			Benchmarks			
			Throughput Analysis			
			Data Processing			
9. ABSTRACT (Continue on reverse if necessary and identify by block number) The work documented in this report was performed for U.S. Army Strategic Defense Command Ground Based Radar Project (GBR). COLSA performed extensive analysis of the digital hardware and computer architectures of the data and signal processing equipment. Included in the analysis is a detailed evaluation of critical throughput for highly parallel computer systems.						
10. DISTRIBUTION / AVAILABILITY OF ABSTRACT ☒ UNCLASSIFIED/UNLIMITED ☐ SAME AS RPT. ☐ DTIC USERS			21. ABSTRACT SECURITY CLASSIFICATION UNCLASSIFIED			
2a. NAME OF RESPONSIBLE INDIVIDUAL Lance W. Cooper			22b. TELEPHONE (Include Area Code) (205) 922-1512		22c. OFFICE SYMBOL	

FOREWORD

The Ground Based Radar (GBR) digital hardware architecture implementation annual report is presented to the GBR Project Office, U.S. Army Strategic Defense Command (USASDC) under Contract Number DASG60-91-C-0006.

Technical questions regarding this GBR hardware support effort should be directed to the GBR Project Office Point of Contact identified below:

Commander
U.S. Army Strategic Defense Command
ATTN: SFAE-SD-GBR (Jack Remich)
P.O. Box 1500,
Huntsville, AL 35807-3801
Telephone: Autovon: 788-1867
 Commercial: (205)955-4370

Accession For	
NTIS CRA&I	☒
DTIC TAB	☐
Unannounced	☐
Justification	
By	
Distribution /	
Availability Codes	
Dist	Avail and/or Special
A-1	

Statement A per telecon
Dug Sloan Army Strategic Defense Command
ATTN: SFAE-SD-GBR-E
Huntsville, AL 35807-3801
NWW 4/23/92

TABLE OF CONTENTS

1.0	EXECUTIVE OVERVIEW.....	1
2.0	INTRODUCTION	2
3.0	TECHNICAL EVALUATIONS	3
3.1	Communications Analysis.....	3
3.2	Processor Support	5
3.2.1	FPS Computing.....	6
3.2.2	CDC	7
3.2.3	MASPAR Computing	8
3.2.4	CSPI RTS-860 Multiprocessor.....	10
3.2.5	GBR Benchmark Results	11
3.3	Data Processor Support	12
3.3.1	VAX-9000	12
3.3.2	CDC 4680MP	14
3.4	BSG Replacement/Control Processor Evaluations.....	15
3.5	Microprocessor Analysis	20
3.5.1	CPU Performance.....	20
3.5.2	Memory Systems	21
3.5.3	Caches.....	21
3.5.4	Bus Snooping	21
3.5.5	Flat and Segmented Address Space	21
3.5.6	MIPS R4000.....	22
3.5.7	IBM RS/6000	23
3.5.8	Hewlett-Packard	23
3.5.9	INTEL i860	24
3.6	Computer Architectures	25
3.6.1	Pipeline Architecture	25
3.6.2	Super Pipeline Architecture	25
3.6.3	Superscalar Architecture	25
3.6.4	Parallelism at the Network Level	25
3.6.5	IBM Power Visualization System.....	26
3.6.6	ALLIANT FX-2800	26
3.6.7	CDC 4680MP.....	27
3.6.8	CONVEX C3200 & C3400.....	27
3.6.9	CRAY.....	27

3.7	TMTR Support.....	28
3.7.1	HWIL	30
3.7.2	APTEC IOC-200	30
3.7.2.1	Data Exchange Bus	31
3.7.2.2	High-Speed Memory	31
3.7.2.3	I/O Processing	32
4.0	SUMMARY OF TRIPS AND MEETINGS	33
5.0	CONCLUSIONS AND RECOMMENDATIONS	34

LIST OF ILLUSTRATIONS

<u>Figure</u>	<u>Title</u>	<u>Page</u>
1	Outline of Final Report	2
2	Unipolar, Bi-Polar Configuration	3
3	USAKA Communications	4
4	GBR Test Site Communications	4
5	ARC Communication Equipment	5
6	Signal Processor Flow Diagram	5
7	FPS 500	6
8	CDC Hardware	7
9	CDC Software	8
10	Advanced Parallel Processor	9
11	MASPAR Computer	9
12	RTS-860 Processor	11
13	GBR FFTs	13
14	VAX 9000	14
15	Control Processor Block Diagram	16
16	Control Processor Configuration	17
17	Current HWCI Control Flow	18
18	Proposed HWCI Control Flow	19
19	RCS Configuration	19
20	Massively Distributed System	26
21	Testbed as a Risk Reduction Tool	29
22	GBR HWIL	30
23	APTEC IOC-200	31
24	GBR HWIL Configuration	32

LIST OF TABLES

<u>Table</u>	<u>Title</u>	<u>Page</u>
1	Computer Performance List	12
2	CDC 4000 Series	15
3	GBR HWCIs	17
4	Raytheon Designed Cards	18
5	Processor Key Characteristics	20
6	Microprocessor Architecture Characteristics	22
7	IBM Power Station	23
8	HP - PA 1.1 Models	24
9	Computer Evaluations	27
10	Computer Evaluations	28

1.0 EXECUTIVE OVERVIEW

During the past several years COLSA has supported the GBR Project Office in the evaluation of the digital hardware and software designs. Previous designs were proposed for the development of the GBR-X system to be built in USAKA. During the past year, the program has changed direction to include several radars to be built and tested. This expansion has grown into the concept of the GBR Family of Radars. The scope of the GBR systems definition has broadened beyond the National Missile Defense (NMD) role with the addition of the Theater Missile Defense (TMD) program. As the program continues, the need for design analysis of digital hardware and software becomes a much more important factor in the development of critical equipment for US. defense systems.

This annual report is the results of the analysis COLSA has provided in support of the design considerations for the TMD and NMD radar system concepts. The COLSA contract DASG60-C-91-0006 titled "Digital Hardware Architecture Implementation" was initiated April 15, 1991 to provide expert analysis and modeling of the critical hardware and software designs within the digital components of the GBR architecture.

COLSA has provided support in technical areas which will be described in detail within this annual report. The report is segmented into five sections: 1.0 Executive Overview, 2.0 Introduction, 3.0 Technical Evaluations, 4.0 Trips & Meetings, and 5.0 Conclusions & Recommendations. The Technical Evaluations section is the principal focus of this report, describing in detail a number of subjects important to sound GBR designs.

*COLSA Provides the GBR Project Office
With a Wide Range of Expertise and Experience
Related to Digital Hardware and Software Evaluations*

There are several areas in which COLSA has provided analysis and recommendations for design considerations. In each of the areas, designs were evaluated for technical merit and how they relate to the technical requirements proposed for the specific item. They were then evaluated as to the soundness of design, with parameters such as throughput, timing, and processing capability being of prime importance. These designs then were evaluated as to the feasibility of fabrication and testability in meeting the design requirements criteria.

The technical areas evaluated include GBR communications; data processor (DP); signal processor (SP); beam steering generator (BSG), receiver, exciter, test target generator (REXTTG); antenna equipment (AE); proposed control processor specifications; latest results of microprocessor designs as described by industry; and finally, HWIL configurations which will assist the GBR program in validation & verification of GBR hardware & software.

It is anticipated that during the next year, a prime contractor will be chosen and COLSA will continue to support GBR Project Office needs as they arise. It is COLSA'S objective not only to accumulate and evaluate prime contractor design information, but to also introduce new ideas, and hardware configurations, in order to enhance GBR system performance.

2.0 INTRODUCTION

All work under contract DASG60-91-C-0006 was performed in response to specific task assignments by the GBR program office to COLSA personnel. The scope of work includes the analysis of GBR digital hardware equipment, to ensure mission performance requirements are met. The report will describe in detail COLSA's effort and methodology, in trying to achieve such goals during the contract period that began the 15th of April and expanded through the end of January 1992.

This report summarizes the contents of the periodic reports submitted to the GBR Project Office during the contracting period. An outline of the report contents is depicted in Figure 1.




Figure 1. Outline of Final Report

3.0 TECHNICAL EVALUATIONS

The technical analysis in this report covers the areas of GBR communications, signal processor and data processor evaluations, the possible substitution of the Beam Steering Generator with an off-the-shelf computer, and the possible incorporation of a Control Processor into the GBR system. The GBR hardware-in-the-loop tests, which are to be performed at the GBR testbed within the ARC, will also be addressed.

3.1 COMMUNICATION ANALYSIS

COLSA investigated the cost and time required to install T-1 communication lines between the GBR test site and the ARC. The following paragraphs describe in brief a T-1 carrier configuration which is suitable for use as the terrestrial portion of the GBR communications.

A T-1 carrier is comprised of twenty-four, eight bit PCM (pulse code modulation) words and a framing bit, which results in 193 bits ($24 \times 8 = 192 + 1 = 193$) of information per frame. The sampling rate of each PCM word is 8 kHz. Each of the 24 channels provides data at a rate of 64 kbits/sec ($8 \times 8 \text{ k/sec} = 64 \text{ kb/sec}$), for a total of 1.53 Mb/sec ($24 \times 64 \text{ k} = 1.536 \text{ Mb/sec}$) for all channels plus 1 frame bit (the frame bit has a data rate of 8 kb/sec). The complete output data rate results to 1.544 Mb/sec ($1.53 \text{ Mb} + 8 \text{ k} = 1.544 \text{ Mb/sec}$).

The wave form used with a T-1 communications line is of a bipolar format shown in Figure 2.




Figure 2. Unipolar, Bi-Polar Configuration

The bipolar format can be described as the "1" state alternating between a positive and a negative voltage. The advantages to the use of a bipolar format become evident when T-1 is applied to a long haul circuit such as GBR requires.

During the investigation it became apparent that installing T1 communication lines can be an unnecessarily expensive proposition if these lines were to remain idle for long periods of time. As a result, two alternatives were investigated: a) Use existing communication lines and work in unison with other agencies b) Install and use T1 lines, only on heavy periods of GBR data transmissions.

Figures 3 and 4 describe the communication possibilities that presently exist between the ARC and USAKA, and the ARC and WSMR, respectively.




Figure 3. USAKA Communications




Figure 4. GBR Test Site Communications

One testbed option was that the Sim. Center based CDC 990 be used as the GBR data processor. Therefore, COLSA investigated what form of communication equipment would be needed to make the existing T3 link between the Sim Center and the ARC operable and conducive to GBR needs.

Figure 5 depicts the communication equipment that must be purchased, and the hardware that is presently within the ARC.




Figure 5. ARC Communication Equipment

According to an ARC communication hardware inventory, the purchase of two Routers and two 703 Link Adapters will be required, at a cost of \$89,000, and if the Sim. Center based CDC990 is connected to the Ethernet, a CDCNet will also be needed and will cost an added \$108,000. Other equipment such as the Lit3Bs, KG95s, and Light Gates are at present available at the ARC, as a result no added purchases will be necessary.

3.2 PROCESSOR SUPPORT

Figure 6 depicts the GBR signal processor's basic functions and configuration in a flow diagram.




Figure 6. Signal Processor Flow Diagram

The GBR signal processor should take large quantities of radar wave forms from the receiver and pick out a small quantity of interesting detection events for the data processor. The signal processor performs this by first compressing the pulses using multipass FFT processing and then detecting events through threshold calculation and thresholding. Several investigations were performed in order to pinpoint the best suited computer, with the results listed in the following paragraphs.

3.2.1 FPS COMPUTING

The floating point systems presentation of their proposed GBR signal processor revealed a number of interesting points. The FPS 500 series is a family of high performance computer systems that tie SPARCscalar, vector and parallel matrix processors into a single Integrated Heterogeneous Super computer. The foundation of the FPS 500 series is the Scalable Interconnect Architecture (SIA) as shown in Figure 7, which allows modular upgrades to the system to accommodate user needs. The systems processors adhere to the Scalable Processor Architecture (SPARC) standard created by Sun Microsystems, Inc.




Figure 7. FPS 500

The 500 series, packs 67 MIPS into a single RISC processor. With FPS system interconnect, up to seven additional SPARC processors can be added to produce 533 MIPS. With FPS software, all this power can be applied in parallel to a single task or to many tasks in a symmetric multiprocessing environment. It should also be pointed out that SPARC is the most widely used RISC processor architecture today so, a wide array of applications, software, network and chip vendors are available and that keeps prices competitive.

+FPX manages all system resources so that multiple processors can be used in several different ways. With symmetric multiprocessing, multiple job streams may be run. By compiling a task to use two or more processors in parallel, the operating system efficiently mixes parallel and symmetric multiprocessing workloads together, to maximize system throughput. The Matrix processor used delivers 13,440 Megaflops of peak performance to algorithms operating on multidimensional data arrays. The Matrix processor ties these processors into a cohesive parallel processing unit to deliver massive processing power to applications with

matrix content. FPS provides a complete software development environment with compilers, debuggers, the FPSMath library and loaders to aid code development for either individual or parallel sets of i860s.

Memory in the 500 Series is configurable from as low as 64 megabytes and a four gigabyte virtual address space that removes artificial memory constraints from large tasks. The 500 Series can grow to one gigabyte of real memory which with interleaving and one gigabyte/sec bandwidth of the system interconnect ensure timely memory service to all processors. In addition, FPS has implemented the High-Performance Parallel Interface (HIPPI). It has four 100 megabyte/sec simplex channels, and serves as a platform for the FPS DMASS disk array system. DMASS systems may contain up to 256 gigabytes of disk storage and transfer rates of up to 128 megabytes/sec are available.

The 500 series can provide a useful advantage to GBR, with its modularity. It allows the user to add single modules instead of upgrading the entire system, such as additional scalar, vector and matrix units, quadruple its speed and expand memory. Its modularity also extends to the Matrix Processor. Available with 4 to 168 Intel i860 RISC processors, it can be configurable to many applications and workloads which can be considered helpful to supporting TMD/DEM-VAL, TMD/FD and GBR-T versions. The FPS processor seems to be well suited to the current GBR requirements. Being a Scalable Interconnect Processor, it allows modular upgrades to the system, thus readily accommodating changing GBR needs.

3.2.2 CDC

CDC, throughout the year, provided reports presenting their ideas on what type of signal and data processor should be used by the GBR program. The highlights of their descriptions are summarized in Figures 8 and 9.




Figure 8. CDC Hardware

It is important to point out that CDC has taken into consideration the use of previously purchased GBR support hardware and even though future modifications on hardware will be required, since the program objectives have changed, the cost of such modifications may be less than that of a total replacements.




Figure 9. CDC Software

CDC's configurations of using the APPs in clusters, depicted in Figure 10, seem to be of sound engineering judgment, and the processors as such, will be adequate in handling GBR requirements for all three systems.

The only drawback is that CDC's suggested signal processor for GBR-T will have a larger footprint than that of MasPar or FPS. It should also be understood that the reconfiguration of this system is presently under development, which makes it somewhat riskier than the already proven systems of MasPar and FPS.

3.2.3 MASPAR COMPUTING

MasPar computing provided information and written materials on their MP-1 computer. The information included a paper on Radar Tracking of the MasPar MP-1. In conclusion, the paper demonstrates that SIMD computers were effective in implementing tracking and that a statistically load-balanced implementation was sufficient to achieve Real Time performance.

The results of the write-up show that over 100,000 tracks can be processed in about a third of the sweep time of an Over the Horizon Radar (OTHR) system. As a result, it is

not unusual to say that today's SIMD computers have the capability of tracking more than 200,000 returns in real time.




Figure 10. Advanced Parallel Processor

The MasPar MP-1 computer is based on a massively parallel processing scheme which is shown in Figure 11.




Figure 11. MASPAR Computer

MasPar assigns one processing element to each data element (or to each set of data elements), and carries out the instructions concurrently on all processing elements. This approach is very powerful for solving problems where each element of a problem obeys the same basic principles. It's high performance results from the replication of simple data processing elements (from 1,024 to 16,384 processors in a scalable array), each with it's own dedicated data memory. A single control unit fetches, decodes, and broadcasts instructions to the array of processing elements. All processing elements execute the instructions simultaneously using their own local data.

The processor has a distributed memory architecture, suitable for signal processing algorithms which makes it a competitive alternative to the more conventional vector processors. Scalability is still an inadequately-defined parallel processing concept, one that has and will result in changing perceptions of performance as it becomes understood and applied.

The MasPar computer, as described earlier, is a highly parallel machine that performs one instruction at a time, over multiple processors simultaneously. The operations are not performed as fast as sequential computers, however the number of processors can allow for higher throughput capabilities than larger single CPU systems. In the largest MasPar, up to 16K processors can be present in a configuration.

Per COLSAs request, MasPar provided tables of FFT performance. All results are for complex, 32-bit floating point data. The "embarrassingly parallel" routines calculate a separate FFT on each processor. One can solve from 1 to 16,384 FFTs simultaneously using these routines. The performance is directly proportional to the number of FFTs being calculated, although the time is the same no matter how many FFTs are computed. Since each FFT must fit into the memory of a single processor, these routines are limited to short FFTs (up to 1024 points on systems with 16 KB of memory per processor).

The "minimum latency" routines solve a single FFT on all the processors of the computer. FFT size is only limited by the total memory in the computer. These routines calculate a single FFT as fast as is possible by using all the parallelism in the MP-1.

The "maximum throughput" FFTs solve from 1 to thousands of FFTs simultaneously. They are a generalization of the previous two types for routines. If the number of FFTs to be solved is greater than one but less than the total number of processors, these routines provide the greatest performance.

3.2.4 CSPI RTS-860 MULTIPROCESSOR

One processor with interesting debugger capabilities is the CPSI RTS-860 Multiprocessor and is depicted in Figure 12. The processor is an i860-based solution for performing compute intensive applications requiring high speed floating-point operations.

It comprises a chassis for 6U sized boards, wired for both VME bus and several VSB clusters, a 68030 based CPU for performing control and I/O, and from 1 to 16 i860 Super Card vector processors to perform the floating-point calculations. The pSOS+ real-time kernel and Unison operating system provide multitasking and multiprocessing capabilities. The RTS-860 can be configured to provide from 80 MFLOPS to 1.28 GFLOPS of peak processing power.

With CSPI's transparent multiprocessing, the user can do something not possible until now: debug an entire multiprocessing and multitasking application from a single window on a workstation. In effect CSPI has automated the process of building the application so that the component programs can be shuffled easily to accommodate the number and types of processors available in the system. As a job is divided into individual tasks, those tasks can be placed on any processor with no change in code. And tasks may be moved from one processor to another to suit application needs.




Figure 12. RTS-860 Processor

Interactive debugging is provided through a point-and-click user interface to CSPI's remote real time system and source-level debugger, called REMEDY. REMEDY runs on the SPARC processor in the Sun workstations, and provides a full cross-development capability with network support. Implemented with a window and mouse-based interface, it is fully symbolic and handles multiple task debugging, dynamic task display, and task breakpoint. Essentially, REMEDY provides the same level of debug facilities that would be available for a single processor.

3.2.5 GBR BENCHMARK RESULTS

As shown on Table 1, memory and I/O speeds were some of the areas evaluated in order to determine true signal processor capabilities. Most of the present day processors including the ones of the table, are very capable machines and would more than adequately meet the reduced GBR requirements. The only definitive way to determine the best suited processor, is to run GBR parallel benchmarks on the processors and evaluate the response times for each one.

Raytheon provided a listing of a signal processing benchmark program. These benchmarks were run on various types of processors presently available at ARC, in order to determine which machine would best fit GBR requirements. SP_SIM program was written in FORTRAN to be used on GBR benchmarks. After generating a simulated radar signal, the

program performs a series of typical signal processing functions on the data, some specified number of times for each of a number of different FFT sizes as shown in Figure 13.

Table 1. Computer Performance List

Name	Type	Memory	I/O	Peak Performance
MasPar	SIMD	1.0GB	1.3GB/S	1.3GFLOPS @ 26000MIPS
FPS	SPARC	4GB	200 MB/S	13.7GFLOPS @ 533MIPS
CDC	APP	4 GB	200MB/s	

The series of functions performed in the processing loop, includes weighting sequence application, Fourier transform, amplitude computation and rescaling, calculations of average CFAR (constant false-alarm rate) threshold, and peak detection. Also, outside the loop, the time is output for benchmarking purposes. Subroutine SIGGEN performs the simulated generation. It calculates the In-phase and Quadrature-phase components (I & Q) of a time -domain signal as if sampled after chirped local oscillator mixing. If desired, SIGGEN can add noise to the samples.

Note that only one simulated signal is generated, i.e. SP_SIM calls SIGGEN just once. Then SP_SIM's processing loop begins each iteration by applying a weighting sequence to that original signal.

3.3 DATA PROCESSOR SUPPORT

Candidate GBR data processors for the GBR program were another part of COLSA's technical analysis. The information generated from this analysis will be described in the following paragraphs.

3.3.1 VAX-9000

The VAX-9000 computer is the company's flagship. It can be expanded from one to four closely coupled processors and each CPU can be optimized with vector processing that gives up to 500 MFLOPS of number crunching performance. Pipeline is the primary technique used to maximize work done per cycle, based on a six-stage approach. This provides for more internal parallelism during the instruction processing stages, by enabling work to be carried out on six instructions simultaneously.

The incorporation of 64-bit data paths, twice as wide as on previous VAX systems, means that double precision arithmetic takes place at nearly the same speed as single precision. The machine balances operations between the vector and scalar job components by

operating the scalar and vector processors at the same time, optimizing overall processor performance.

At the heart of every system acting as the traffic manager, is the System Control Unit (SCU) a 2 GByte per second crossbar switch. It provides for the fastest processor, memory, and I/O interconnect in the VAX family. The very large main memory size of the VAX 9000 reduces paging and balances performance by enabling the caching of small databases and frequently referenced files. Up to 512 MBytes of memory are available, using 1 Mbit DRAMS on 64 MByte arrays.




Figure 13. GBR FFTs

Architectural support is provided for up to 2 GBytes using 4 Mbit DRAMS. Unlike other VAX systems, the VAX 9000 employs one to four XMI buses solely for I/O to and from storage, networks, and VAX cluster systems. The XMI bus does not carry any interprocessor communications or processor to memory traffic, leaving the entire bandwidth of the XMI available for the I/O traffic. As a result, the VAX 9000 delivers sustained I/O throughput in a balanced manner. Each XMI is capable of moving data at 80 MBytes per second for a system total of up to 320 MBytes per second.

In conclusion, Digital claims optimal super computer performance for numerically intensive, vectorizable applications, having the best software development environment and ADA compiler on the market today. Figure 14 depicts a block diagram configuration of the VAX 9000.




Figure 14. VAX 9000

3.3.2 CDC 4680MP

On November 14 1991, CDC presented their 4000 series computers to both COLSA and GBR program office personnel. The CDC4680MP is a probable substitute to the CDC 990.

The 4000 series starts with the basic workstation type 4320, to the multiple processor types 4370MP(1-4 processors), 4375MP(1-8 processors), and the GBR proposed 4680MP(1-4 processors) a high performance RISC processor. It is an air cooled machine, with dimensions of 6ftx4ftx2ft, requiring 120VAC power for operation. It is an expandable system with a total of 6 VME buses and 24 VME Controllers. It has an Intelligent Peripheral Interface(IPI), an Ethernet link(10 Mbits/sec), a Fiber Distributed Data Interface(FDDI, 100 Mbits/sec), HPPI, and SCRAMNet reflective memory.

Table 2 describes the key characteristics of the 4 different CDC systems. It is important to point out that by mid 1992 the company will release a new version of the 4680MP with 32 bit R6000's, a clock frequency of 80 MHz, an instruction cache of 512 KB, and a data cache of 2 MB.

If the CD4680MP is used, CDC commits to common user interface, application availability and ease of connectivity, since the computer configures to open system standards. In terms of connectivity, the computer uses what is called the SCRAMNet which is a fiber optic communication link, at 250 Mbits/sec.

Table 2. CDC 4000 Series

KEY CHARACTERISTICS				
Model	4336	4339	4360	4680MP
Processor	R3000	R3000	R3000A	R6000
Clock Freq.	24MHz	24MHz	33MHz	60MHz
NO. of CPU	1-4	1-8	1	1-4
Memory MB	8-128	8-256	32-256	32-384
Bus MB/Sec	100	100	133	240
Instr. Cache	32-128KB	32-128KB	64KB	64KB
Data Cache	"	"	"	"
I/O EXP.	5VME	10VME	7VME	24VME
SPECmarks	17.6/cpu	17.6/cpu	32.1/cpu	56/cpu

The system operates on TC/IX-Real Time Unix that uses the LynxOS Kernel built for real-time, and provides real-time primitives. A full Verdix Ada Development System (VADS) is used, with ANSI/MIL-STD-1815A and symbolic debugger.

CDC claims that existing GBR software development can be finished on the 990 computers and then ported to the 4680MP. The vectorization portion will be the only part of the software that will have to be modified, and that is approximately 10% of the total software.

3.4 BSG REPLACEMENT/CONTROL PROCESSOR EVALUATIONS

COLSA was instructed to investigate the replacement of certain GBR subsystems such as the Beam Steering Generator, with off-the-shelf computers. It has been COLSA's opinion all along, that a control processor replace a number of GBR subsystems such as the BSG, the TTG, etc. Figure 15 shows such a concept in a block diagram configuration. It has not yet been determined if this possibility exists, but the merits of such design changes would be more than just convenient. It would eliminate the necessity of stocking a number of different subsystems as spares. It would also provide for the hardening of one unit rather than many, which is especially important for TMD-GBR. The design of hardware and software/firmware would be reduced to a minimum, and designs would already be proven, if off-the-shelf processors can be used.

Among other advantages in such a design, the principal one would be that all software could be co-located in one control processor, thus making it easily accessible to testing as

well as facilitate future design changes. Additional hardware will be required and can be fabricated by the prime contractor (control hardware portions as shown in Figure 16, but they would be of minimal amount and complexity, compared to present GBR designs. The hardware can be located within the RF section of the system and connected to the digital section through fiber optics, which will not only make the transmissions/receptions of commands less susceptible to noise, but will also make the connect/disconnect time between the two sections faster.

COLSA is also pursuing the possibility of using one computer to perform both functions of the GBR data and control processors. That would mean the use of only one super computer within the GBR system.




Figure 15. Control Processor Block Diagram

On that basis, computer manufacturers were contacted and technical information describing maximum capabilities of their products was ordered. Cray, Alliant, Convex, CDC, and Encore were some of the computer companies in consideration.

It should be pointed out that there is a possibility that an Alliant Campus/800 computer will be installed within the ARC, which can prove helpful to the GBR program, if it is found that a Campus/800 or Alliant 2800 can be used as either the GBR control processor, data processor or both, thus alleviating the need of purchasing a second computer for testbed use alone. In such a best case scenario, with hardware-in-the-loop testing, the program office will be capable of evaluating the majority of the GBR system software as well as all major digital hardware subsystems.

The control processor configuration as shown in Figure 16 separates the digital and RF hardware. The present Raytheon Radar System Control (RSC) design should be modified not only for the reasons stated above, but also to a) reduce the quantity of card designs and spares, b) de-embed the control software to make it more testable, and c) reduce system block diagram complexity to ease simulation.

The current RSC design includes a combination of digital hardware and software which takes a stream of radar requests from the data processor and causes the RF and analog hardware to perform the radar requests. The Raytheon RSC is embedded inside the HWCIs of the radar. Specifically, it consists of the parts of HWCIs as shown in Table 3.




Figure 16. Control Processor Configuration

Table 3. GBR HWCIs

TCE	All
BSG	All
AE	Only DMC chips
SP/HSR	Only control system
REX/TTG	Only control system
TCU	Only control system

The RSC is implemented in the Raytheon designed cards, most of which were designed specifically for GBR. The list of the cards necessary to implement this function is shown in Table 4. It is evident that there is major design effort involved in the Raytheon RSC design. It may therefore be applicable to reduce or replace such efforts with commercial off the shelf computer equipment.

Table 4. Raytheon Designed Cards

IPI Interface	G404079-1	RBFN/TBFN	G500342-1
ISI Interface	G404078-1,2,3	Clock Driver	G500107-1
CBI Interface	G404108-1	Radar Time Register	G404074-1
RBI Interface	G404109-1	Radar Register	G404073
A/D Interface	G404101-1	Master Clear/System Clock	G404075
Line Driver Controller	G404161-1	Status Register	G404165
BSG Instruction Handler A	G500340-1	Receiver/Antenna Interface	G404162
BSG Instruction Handler B	-	Bus Monitor	G500341-1
High speed Controller 1	G500337-1	Beam Processor	G500334-1
High speed Controller 2	G500338-1	Time & Control	G500343-1
Array Matrix Driver	G500335-1	Bus Terminator	G438946-1
TDU Driver	G500336-1	68000 Processor Board	-

The current RSC takes a stream of commands from the Data Processor and broadcasts it to each HWCI via CDC ISI like interfaces. Each HWCI receives this stream and processes it with a HWCI dependent set of 68000 series microprocessors. The processors in turn, output the formatted actions to special control hardware. At the correct time, the control hardware output the actions to the Radar analog hardware.

The HWCI control software is embedded inside the HWCI's. It's inputs and outputs are available only within the HWCI's which makes the testability and modification of the software rather difficult. Figure 17 depicts the current Radar HWCI Control Flow.




Figure 17. Current HWCI Control Flow

The suggested RCS configuration would use a single (initially commercial) computer to run all of the control software, combined with a consolidated set of special control hardware. Figure 18 describes the proposed Radar HWCI control flow.




Figure 18. Proposed HWCI Control Flow

COLSA's RSC configuration depicted in Figure 19 would place the control software in a single isolatable and testable location. This case would eliminate the need of designing custom microprocessor cards and buses.




Figure 19. RCS Configuration

The hardware designs, will be required to perform the following functions: a) distribute radar time and system clock throughout the radar, b) distribute specific radar commands from the control computer throughout the radar, c) meet the specific hardware control interface requirements for the RF and analog hardware, and d) provide paths for testing individual parts of the RF, analog and control hardware.

The Raytheon design uses approximately 50 Motorola 68000 series processors to perform the required tasks. As a result, the commercial computer in question would have to possess at least as much horsepower, in an easily usable form. This means that the horsepower is to be provided in less than 50 processors, running with a shared data memory. The processor

count limit guaranties that the horsepower will be at least as concentrated as it was, while the shared memory guaranties a simple method of communication among cooperating processors.

3.5 MICROPROCESSOR ANALYSIS

The Raytheon RSC as described earlier, uses a number of Motorola processors (20 MHz 68030s). Since the 68030 was designed, there have been many advances in microprocessor technology. The result is that there are available processors today with 4 to 8 times the processing speed of the 68030. The processors selected for comparison in this report and to be addressed later are: The R4000, the IBM RS/6000, the HP-PA, the Intel i860, and the SPARC.

Multiprocessor technology has advanced so rapidly that it has allowed CPU performance to double almost every year. This speed of change has brought to light new concepts and much marketing hype, including unrealistic manufacturer claims that can confuse basic but none-the-less important issues, when trying to determine which system is best suited for a certain application. This section of the report, will attempt to suggest which processors, and what type of architecture would best be suitable for the GBR RCS modification. Table 5 depicts a number of different processors with their individual characteristics.

Table 5. Processor Key Characteristics

	MIPS R3000	MIPS R6000	SPARC	HP-PA	INTEL i860	IBM RS/6000
Clock Speed	25MHz	66	40	50-66	40	20-40
MIPS	27	-	28	57-76	64	29-56
SPECmarks	16	57	25	55-72	30	32-72
MFLOPS	3	-	4	50-66	10	9.2-25.2

3.5.1 CPU PERFORMANCE

The computing time for any program depends on: a) the MHz rate of the chip, b) the number of instructions, c) the number of clock cycles per instruction. A formula which is often used to predict the time to complete a program is shown below where CPI=clock cycles per instruction.

$$\text{Compute Time} = (\text{Instructions/program}) \cdot (\text{CPI/clock Rate (MHz)})$$

The main battleground in RISC design today is in increasing MHz and decreasing cycles per instruction. For example, the 68030 is 20 MHz and 40 CPIs, while some latter date processors are greater than 40 MHz and less than 2 CPIs.

3.5.2 MEMORY SYSTEMS

DRAM (Dynamic Random Access Memory) is used as CPU memories. The densities of DRAM have been quadrupling every three years but the access time has decreased by only a factor of three. If the speed of the CPU is dramatically increased while keeping the memory speeds constant, the CPU will spend precious clock cycles waiting for memory to respond. That's where caches come in to play.

3.5.3 CACHES

A cache is a high speed memory device which resides between the microprocessor and the CPU memory. Its function is to hold often used data, so that most of the microprocessors requests for data can be quickly satisfied. A cache is a higher speed SRAM (static RAM), a more expensive subset of the memory that it represents. First level caches can respond to the CPU in a time which introduces no wait states. The CPU memory and the caches must maintain accurate copies of each other's data, with few exceptions. Cache data is read in as a chunk of data called a cache line. Reading data in lines takes advantage of the fact that the next data wanted, is likely to be near to the data just accessed, a concept known as spatial locality.

The most flexible approach to caching data is to allow any line of CPU memory to replace any line of the cache. The cache controller can then determine the last line used and replace it, an approach called the fully associative cache. This approach is very complex to implement. The IBM RS/6000, for example, implements a 4-way set associative cache. Each line of CPU memory can inhabit any of 4 lines in cache. The simplest approach to cache mapping is the one way set associative cache, also called direct mapped cache. A direct mapped cache typically masks off the bottom bits of the cache line and uses them as an index to the cache. A 64 KB four way associative cache generally has the same miss rate as a 128 KB direct mapped cache.

3.5.4 BUS SNOOPING

Bus snooping is a technique which allows each bus interface referencing memory to monitor all bus traffic. On a read request, the bus interface with the latest copy of the data responds by placing it on the bus. If the data is not contained in any of the caches, CPU memory has it. If one or more caches contain the data but have never written into that memory location, any cache or CPU memory can respond. The CPU can then either write the result back to CPU memory immediately (write-through), or wait until the cache line is either replaced by new data or requested on the bus (write back).

3.5.5 FLAT AND SEGMENTED ADDRESS SPACE

Virtual memory systems allow applications to run which are larger than the available physical CPU memory. A running program generates a memory reference which is called the effective address. The maximum memory that a program can address is its virtual address space. The maximum memory that a program can directly address in a single instruction is the effective address space. The maximum amount of memory that can be configured on a system is its physical address space. The maximum virtual and effective address spaces are architectural parameters. The maximum physical address space is an implementation parameter. The maximum virtual address space is different from the width of the instruction and data paths. The IBM RS/6000, has a 128 bit instruction path. The Intel i860 has a 64 bit data pathway. These

parameters define the number of bytes of data and instruction which can be read in a single clock cycle.

Large virtual address spaces are achieved through either a flat or segmented architecture. A flat virtual address space is one in which the maximum effective address size is equal to or greater than the maximum virtual address space. A flat virtual address space allows a program to address any location in the virtual address space directly. A segmented address space is a work around to achieve a larger virtual address space than the effective address. In a segmented virtual address space, the program can only directly address the virtual memory in a single segment, or a small number of segments. To address virtual memory in other segments, the program must swap segments.

In the IBM RS/6000, the 32 bit effective address consists of 28 bits of segment offset and 4 bits of segment ID. Each segment is 256 MB long and up to 16 segments may be mapped simultaneously. Table 6 illustrates the address limitations and segment sizes for the major microprocessor architectures for 1991.

Table 6. Microprocessor Architecture Characteristics

	MIPS R4000	HP-PA 1.1	INTEL i860	IBM RS/6000
FLAT / SEGMENTED	Flat	Segment	Flat	Segment
MAX. VIRTUAL ADDRESS SPACE	64 bits	48 bits	32 bits	52 bits
EFFECTIVE ADDRESS SIZE	64 bits	32bits	32 bits	32 bits
PHYSICAL ADDRESS SIZE	36	32	32	32
SEGMENT SIZE	N/A	32 bits	N/A	28 bits
TOTAL SEGMENTS	N/A	16 bits	N/A	20 bits

3.5.6 MIPS R4000

The MIPS R4000 microprocessor contains on-board instruction and data caches, with an initial configuration of 8 KB per cache. The external clock rate of the chip is 50 MHz. The internal super pipeline runs at double the external clock rate 100 MHz. The instruction and data paths are 64 bit wide. Two instructions are brought in every external clock cycle allowing the internal 100 MHz super pipeline to run at full speed. A full double precision floating point load or store can be performed in one internal clock cycle.

The MIPS R4000 implements a complex instruction called cache operation which manages first and second level caches, as well as requests from other processors. The R4000 can be used to implement a wide variety of caching schemes, and a joint or split instruction and data caches.

3.5.7 IBM RS/6000

The RS/6000 consists of a nine chip set running at 20 and 25 MHz clock rates. In late 1990 IBM announced the POWERStation 550 based on a 41.6 MHz version of the RS/6000. IBM is currently working on integrating the nine chip RS/6000 into a single multichip module capable of running at 100 MHz. This implementation is expected to be available in systems in early 1993. It is a 52 bit virtual address machine, with an effective address length of 32 bits. A process can map only sixteen (4 bit) segments of 28 bits (256 MB) per segments at a time. The other 20 bits of address are used to select segmentation register, of which only 16 can be mapped into the effective address at once. The segment registers allow the program to move from 28 bit window to 28 bit window. Changing segment registers requires a system call. A large data set for simulation, for example cannot be larger than 4 GB in size, residing in 16 segments. Programmers working with large datasets must first know which segment the data they want to access resides, and then map the data to an effective address that may change depending on how the segments are selected.

The IBM RS/6000 cannot grow by expanding its virtual addressing in a linear manner. Changing the segmentation approach would require a change in the instruction set architecture and therefore a radical change in both operating system and applications. IBM has effectively locked itself into a 32 bit programming model for the foreseeable future. IBM is now modifying the RS/6000 chip to support symmetric multiprocessing with rumors of an asymmetric multiprocessor in the near term. Table 7 describes the capabilities of the IBM POWERStation workstations.

Table 7. IBM Power Station

	320	320H	520	530	550
MHz	20	25	20	25	41.6
MIPS	29.5	37.1	29.5	37.1	56.0
MFLOPS	9.2	11.7	9.2	15.2	25.2
SPECmark	32.6	41.2	32.6	43.4	72.2

3.5.8 HEWLETT-PACKARD

The HP-PA system has broken no new ground in microprocessor implementation. HP has used few advanced techniques, but has concentrated primarily on increasing the system performance through higher MHz rates and simple optimization techniques. The focus of the HP-PA system was to produce the highest performing workstations at the low end of the market and they have succeeded. Beyond the MHz increases on the HP-PA 1.1 chip set, new compilers are the key to the HP performance.

In simulation codes the HP compilers perform extensive optimization related to loop unrolling, instruction scheduling, cache management and branch prediction. The HP-PA 1.1 has two new instructions, FMPYADD and FMPYSUB, which allow the

microprocessor to perform a floating point add/subtract and multiply in one instruction. The weaknesses of HP is that it is a proprietary chip set. The narrow base of the HP workstation means that it will be difficult to get applications on the architecture. It also does not support multiprocessing. There are no synchronization primitives and although the caches are physically tagged, indications are that the main goal of the implementation is low cost high performance in the uniprocessor side.

The HP-PA 1.1 architecture supports a 48-bit virtual address space. Effective addresses are 32 bits long with 32 bits of maximum physical memory. There are 16 bits of segments supported. The systems currently available come in three models as shown in Table 8, the 720 a 50 MHz desktop version, the 730 a 66 MHz desktop system and the 750 a 66 MHz desk side implementation.

Table 8. HP - PA 1.1 Models

	720	730	750
MHz	50	66	66
MIPS	57	76	76
MFLOPS	17	22	22
SPECmark	55.5	72.2	72.2

3.5.9 INTEL i860

The i860 XR microprocessor was the industry's first 64 bit, microprocessor and was the first to integrate integer, floating point and graphics capability on a single chip. The i860 family shares the same high performance architecture which includes a 64 bit data bus and a 32 bit address bus. A 32 bit RISC integer core executes one instruction per clock cycle using a four stage pipeline. The floating point unit of the i860 architecture combines both scalar and vector pipelined processing to maximize efficient execution of floating point instructions. Special dual operation instructions provide parallel execution in the floating point adder and multiplier units so two floating point results per clock cycle are possible. Fast program execution is sustained by on chip data and instruction caches.

Chip reviewers have commented that the i860 has a small "sweet spot" or a small range of code that can actually take advantage of the chip's performance characteristics. Another problem may be the fact that in a general purpose environment, no one has yet been able to produce a compiler and system that can hit the i860's "sweet spot". The following comment was a response from Preston Briggs, of Rice University, May 16 1991: "The i860 can smoke a Sparc, but it takes a smart (mostly non existent) compiler and having the right applications."

3.6 COMPUTER ARCHITECTURES

The following description will help understand the type of computer architectures that are available in today's market. They can also assist in determining which is the most preferable architecture for the GBR program.

3.6.1 PIPELINE ARCHITECTURE

Pipeline (MIPS R3000) is a technique which allows a microprocessor to execute one instruction per cycle even though it may take longer than one cycle for any individual instruction to complete. The pipeline breaks an instruction into component parts called stages. Each stage requires one clock cycle or less to complete. In theory, pipe lining offers a peak rate of one instruction per clock cycle. In practice there are a number of hazards such as data, control and structural hazards, which prevent it from reaching this goal.

3.6.2 SUPER PIPELINE ARCHITECTURE

Super pipe lining (MIPS R4000) allows a designer to increase the MHz rate of the CPU. Typical internal clock rates are of 100 MHz and beyond. It is a simpler optimization technique than that of superscalar. It doesn't require a great deal of control logic and instructions need not be fed to the CPU in any particular order. The simplicity of this architecture allows the chip to be run at a higher MHz rate and achieve a better ratio of actual CPI to theoretical peak CPI. The disadvantage of super pipe lining is that the pipeline becomes more complex and is somewhat more likely to stall due to the same hazards described earlier. However the degradation is generally overwhelmed by the ability to push the chip to higher clock rates due to the simplicity of the chip.

3.6.3 SUPERSCALAR ARCHITECTURE

Superscalar is a technique designed to decrease the number of cycles required on average to complete an instruction. Superscalar CPUs commonly run in the 25 MHz to 40 MHz range. The IBM RS/6000 is a good example of a superscalar microprocessor. It has four functional units which can operate in parallel. To support these functional units (branch, condition code, floating point and integer units) the RS/6000 has a 128 bit wide data path from the instruction cache into the CPU. Four instructions worth of data can be read in during one clock cycle and up to four instructions can be executed per clock cycle.

The IBM RS/6000 can achieve a theoretical peak CPI of 0.25. To do this however, the CPU must be fed a precisely balanced mix of branch, condition code, floating, and integer instructions. In the real world this is impossible to obtain. The CPI of the IBM RS/6000 approaches one in normal operation. The disadvantage of superscalar design is its complexity. Only recently was IBM able to deliver its production compiler for the RS/6000. This architecture must be closely matched with the appropriate optimizing compilers in order to be effective.

3.6.4 PARALLELISM AT THE NETWORK LEVEL

A massively distributed system as shown in Figure 20 consists of two or more multiprocessing workstations connected in such a way that they appear to be a single system. The system has a physical memory equal to the sum of all physical memories on each of the systems. Massively distributed systems allow the user to create a system of indefinite size by linearly adding systems and cost. Each system is connected by a special purpose high speed optical network. Bandwidth requirements across the network dedicated to memory accesses are reduced through the use of a directory based cache coherency scheme. Most multiprocessing systems use a snoopy bus technique to maintain cache coherency. A directory based system increases the amount of memory and logic allocated to cache coherency in order to reduce the traffic on the interconnection bus.

3.6.5 IBM POWER VISUALIZATION SYSTEM

The IBM POWER Visualization Server seems to be very well suited for the purpose of being used as the GBR control processor. It is a parallel machine with up to 32 Intel i860 RISC processors. It has a very wide (256 bit) data path that achieves internal data transfers at speeds exceeding 1 GByte/sec. Internal global shared memory can reach 1 GByte about four times as much as typical high-end graphics workstations. In addition, each processor in the server has its own 16 MByte local memory. For a 32 processor system (16 MB/processor), it totals an additional 512 MBytes of high speed memory. This local memory acts like a very large cache that keeps data flowing smoothly between memory and processors.




Figure 20. Massively Distributed System

High Performance Parallel Interface (HPPI) channels on the Server sustain transfer rates of up to 100 MBytes/sec. HPPI channels link the Server with a parallel disk array and a video controller, as well as provide capability for attachment to an external mainframe and super computer. For the GBR control processor configuration, the Model No. 002 with 16 i860 processors should be capable of providing adequate throughput.

3.6.6 ALLIANT FX-2800

The FX-2800 meets most of our set requirements. Its I/O bandwidth can meet GBR requirements. It can be expanded to hold 29 processors with an estimated SPECmark of 648, and it does have HPPI and fiber optic capabilities. The only drawback to this machine is that the processors fetch instructions over that same bus and that makes the performance of the computer non-linear in scale, as additional processors are installed.

3.6.7 CDC 4680MP

The CDC computer has marginal throughput and no room for expandability. It produces an estimated SPECmark of 205, but comes with only four processors maximum. It uses the R6000 processor with clock frequency of 80 MHz and a MIP rating of 68. Its bus bandwidth is 240 MBytes/sec and has a HPPI interface and fiber optic capabilities.

3.6.8 CONVEX C3200/C3400

The Convex computers seem to be the least suitable for our application. They are both 25 MHz processor machines with only 8 processors maximum and a throughput which is quite low.

3.6.9 CRAY

Cray was unable to suggest a machine that would come close to the control processor requirements stated earlier. The only system that could conceivably be considered, is that of Floating Point Systems, which Cray now owns. This computer is a Spark Scalar processor but with 8 processors maximum it has marginal throughput for our application. Tables 9 and 10 show the information received to date and can be used as a quick reference chart.

Table 9. Computer Evaluations

CO. NAME	Cray "FPS" early 1992	Alliant FX-2800	Convex C3400	CDC 4680MP
Comments	-	288SPECmarks	-	-
Computer Memory	4GBytes	4GBytes	2Gbytes	179GBYTES
Architecture	64bit scalar/vector	Symmetric multiprocessor Shared memory	RISC C3 Series	RISC
Clock Frequency	-	40MHz	50 MHz	80MHZ
Type Of Processors	-	i860 64 bit RISC	C3400 64 bit	R6000 32bit
No. Of Processors	8	28	8	1-4
Processor Speed	70Mips	40Mips	-	68MIPS
Processor Memory	0.5Gbytes	2MB/Module	-	32-384MBYTES
Dual Port Mem.	-	-	-	-
Proc. Intercomm.	Yes	Yes	-	-
Bus Bandwidth	260MBytes/sec	24 processors max if HPPI is used (100MBytes/sec)	200MBytes/sec	240MBytes/sec
Fiber Optics	Yes	Yes	-	SCRAMNet
HPPI Interface	Yes	Yes	-	Yes
Contacts	Larry Holsman 922-9300	Mark Benjamin 703-847-5300	Curtis Calwell 539-9430	Pete Ogden (617)466-6154

Table 10. Computer Evaluations

CO. NAME	ENCOR 91 SERIES	SILICON GRAPHICS 4D/ 480	IBM Visualization
COMMENTS	-	166SPECmarks	-
COMPUTER MEMORY	4 GBytes	256MBytes	-
ARCHITECTURE	32 bit RISC	Fully symetric processor	Parallel machine RISC
CLOCK FREQUENCY	25MHz	40MHz	40MHz
TYPE OF PROCESSORS	Motorola 88100 25Mhz RISC 32Bit.	MIPS R3000	i860
NO. OF PROCESSORS	4/board up to 16CPUS	8 processors 2/board	32
PROCESSOR SPEED	80+Mips/board of 4CPUS	300Mips	-
PROCESSOR MEMORY	16MBytes shared min.	256MBytes	16MBytes/proc
DUAL PORT MEM.	Single port access	-	-
PROC. INTERCOMM.	-	Yes	Yes
BUS BANDWIDTH	100MBytes/sec	100MBytes/sec	100Mbytes/sec
FIBER OPTICS	-	Yes	No
HPPI INTERFACE	No	Yes	Yes
CONTACTS	Mel Geiger 837-8250	David Crafton 830-5400	Eric Coldbie 830-6025

3.7 TMTR SUPPORT

During the elapsed contract period, COLSA personnel assessed key technologies to be utilized by the GBR testbed located at the ARC. Analysis was performed, to determine the utility of the testbed as a risk reduction tool during the design, development, and testing phases of all GBR systems.

Most of the information listed in Figure 21 was presented to Col. Ryan and the GBR program office, with a detailed description of ARC capabilities and it's future contributions to the program. Hardware-in-the-loop testing and design validation & verification for hardware and software alike, will be a major area of concentration which will require the purchase of additional hardware such as the GBR signal processor, and data processor.

The areas described below are in need of additional investigation, and are areas in which testbed based GBR hardware could assist program office personnel in verifying the validity of designs of interest.

- Independent validation & verification testing of signal processor design standards and constraints including signal processor software.
- Evaluation of signal processor functional & performance requirements.
- Validation & verification testing of SPCI interfaces.
- Development or verification of the Test Requirements Specification (TRS) document.
- Running of the CI HW/SW test plan in order to verify that the chosen signal processor will work as expected.

RISK REDUCTION CATEGORY	Testbed Activity	Type of Risk		
		Perf	Cost	Sched
HARDWARE	Technology Insertion & Transfer	X	X	X
	HWIL Testing	X		X
	Design Validation	X		X
SOFTWARE	Software Development	X		
	Software Validation	X	X	X
	Mission Planning	X	X	
ANALYSIS	System Simulation & Modeling		X	
	Data Analysis	X		X
	Potential ECCM	X	X	X
OPERATIONS	Operational Training	X		X
	Observe Mission Conduct	X	X	X
	BM/C3	X	X	X

Figure 21. Testbed as a Risk Reduction Tool

Other areas through which independent testbed evaluation will enhance prime contractor design direction for the signal processor are listed in detail below: Latency, noise threshold, peak detection, extended target detection, range interpolation, amplitude interpolation, CFAR threshold, RCS threshold, measurement accuracy, A/D sampled data handling, track FFT, track detection processor, pre-detection scaling and noise average computation.

As described earlier in the report, COLSA did provide information on candidate signal processor architectures under consideration by Raytheon, for the T and TMD systems. These systems vary in design such as SIMD, MIMD, and pipelined configurations. To be able to fully evaluate the designs under consideration, benchmarking and analysis of real application processing is required. It is planned that a GBR testbed located at the ARC, will provide a

convenient facility for evaluating candidate designs throughout the competitive phases of the new GBR program acquisition.

3.7.1 HWIL

Hardware-In-The-Loop (HWIL) testing can support critical hardware and software evaluations, for a number of GBR subsystems. Several components will be needed for a HWIL configuration including a DP, SP, I/O computer, disk drives, and software simulations for the radar modeling. If the TRW developed RADSIM is not useful, simulations of the potential targets, antenna, transmitter, receiver, and other analog components will have to be generated and used, along with the existing software of the two processors mentioned above. Figure 22 depicts a preliminary configuration of the GBR hardware and software needed for the HWIL testing.




Figure 22. GBR HWIL

3.7.2 APTEC IOC-200

COLSA investigated on a number of hardware systems needed within the testbed to perform HWIL tests. APTEC can be considered the GBR I/O computer due to its high speed capabilities. This computer is being presently used at the ARC by personnel who possess extensive experience and skill in not only operating the system, but are able to modify software and hardware alike, in order to fit a specific application such as GBR. It is our plan if possible to utilize the APTEC by testing it first, using the RADSIM simulation developed by TRW. Various other computers are also available within the ARC and can be used by the program, to develop additional code when needed. COLSA plans to utilize its assets in developing additional hardware, so as to make the GBR hardware compatible with peripherals, including other systems such as BM/C³.

A problem with running real-time simulations of the GBR equipment will be the high bandwidth needed for driving the signal processor. Using the estimated GBR-X design parameters of 30 Mhz data on three parallel channels, it has been established that the only existing I/O computer system presently available at the ARC which could be used to provide the simulated digital data necessary to perform HWIL testing on the GBR signal and data processors, is the APTEC IOC-200.

The APTEC is a high I/O throughput computer, and could easily be upgraded to provide support on GBR applications. It allows high speed peripherals and processors to be more fully utilized by providing a 200 MByte/sec internal bus, programmable I/O processors (IOPs) that communicate with external peripherals at up to 12 MBytes/sec, and high-speed memory with access rates of 50 MBytes/sec per memory board. A typical IOC-200 is shown in Figure 23. It should be pointed out that the APTEC presently based at the ARC will have to be upgraded in order to handle GBR requirements, at an approximate cost of \$160K (VME controllers \$10k ea, I/O processors \$19.2K ea, VSP2 \$28k ea, and a disk storage unit \$75k ea).




Figure 23. APTEC IOC-200

3.7.2.1 Data Interchange Bus

High bandwidth is built into the I/O computer through the Data Interface Bus (DIB), the backbone of the IOC-200. The DIB is a high performance, parallel, synchronous bus with a maximum sustainable bandwidth of 200 MBytes/sec. Unique, I/O directed bus commands help achieve high throughput. The DIB employs multi-word addressing to transfer as many as four 32-bit data words with a single address. This allows simultaneous reads and writes at full bus bandwidth, and maximum utilization of the IOC-200's separate read and write buses.

3.7.2.2 High-Speed Memory

The IOC-200 possesses a 32-bit random access High-speed Memory available in 2 MByte increments. It provides the fast, readily available memory needed

for I/O intensive tasks. Attached peripherals, via respective IOPs, can also access IOC Memory at a random, aggregate rate of 50 MBytes/sec per memory board. As a result overall IOC-200 performance can activate 200 MBytes/sec with a minimum of 4 memory boards. High-speed memory acts as a shared database and as a FILES-11 structured device for the VAX and all devices attached to the IOC-200. It provides an interim staging area for high speed transfers between attached peripherals or between the VAX and peripherals. In addition, it can be used to store programs to be executed by IOPs. To reserve a memory block, an IOP can test and set a flag simultaneously. This permits fast memory handling and virtually eliminates contention for shared high-speed memory among multiple processors.

3.7.2.3 I/O Processors

The I/O processors are single board 16-bit processors that provide the intelligence necessary to orchestrate I/O operations and offer computing capabilities independent of the VAX host. Each IOP can be programmed to manipulate data on-the-fly as it enters the IOC domain through a dedicated private bus on each I/O processor board. Two different IOPs are available for the IOC-200 and function identically, except the private bus in order to simplify system programming. IOPs support standard UNIBUS peripherals available from many manufacturers, or higher performance peripherals via Aptec's 12 MByte/sec OPENbus I/O Processor.

A preliminary HWIL configuration is described in Figure 24. Even though the final GBR designs have not been finalized, it is anticipated that the HWIL configuration will essentially remain as shown.




Figure 24. GBR HWIL Configuration

The simulations will be developed in various computers available at the ARC, stored in disk storage units, and upon commencement of a specific mission, the digital information will be routed through the APTEC as Sum, Alpha, and Beta inputs to the signal processor. The signal processor will in turn perform its tasked calculations and pass the results to the GBR data processor. These results will be evaluated, further calculations will be performed and new information will be passed back to the APTEC (thus closing the feedback loop), in order to generate updated digital data such as additional target information, antenna adjustments, or possibly change mission scenarios.

4.0 TRIPS AND MEETINGS

COLSA personnel attended a short Georgia Tech presentation on their recently developed Parallel Function Processor (PFP), intended to be used as a general purpose emulator. The machine is a prototype processor that consists of 2 crossbar switches, 64 nodes, and contains 4 processors: INTEL, GT-FPP, GT-FPX, and GT-1860. COLSA plans to attend a Georgia tech demonstration of the PFP the 18th of July, an opportunity to collect additional technical data. Further technical information will be presented as it becomes available.

As requested by the GBR program office, COLSA personnel also attended a Theater Air Command and Control Simulation Facility (TACCSF) briefing that was presented for the TMD program. The presentation concentrated on the capabilities the TACCSF testbed possesses in the areas of testing, MIL simulators, all execution-level air defense functions, and integrated weapons and C2 elements. The testbeds assets are: 57 tactical consoles, 29 host computers, 77 display processors and 27 array processors, 1.4+ million lines of executable code distributed architecture, IAD system models (CRC, MCE, E-3, TSQ-73, HAWK, PATRIOT, F-15, SIS), and data links (ATDL-1, PADIL, TADIL-B, TADIL-J). TACCSF also possesses radar models of radar types such as the pulse, CW, swept CW, pulsed doppler, doppler, monopulse, etc. It also has in place weapons such as the PATRIOT, HAWK, and F-15. BM/C3 type C2 levels and systems are in place with autonomous weapon systems bridge, battalion, composite BN, FACP, CRP, CRC, AWACS and SOC. It is an operational facility, the largest Air Defense Simulation Facility in the world, and it's MIL capabilities could provide insight and answers unobtainable with analytical simulations. Additional and more detailed information can be obtained from the handouts distributed at the briefing.

During the elapsed performance period, COLSA gave three quarterly review presentation to GBR program office personnel. Addressed subjects included cost summaries, the contract organization, and the technical performance of the GBR-T/TMD and TMTR systems. The technical portions of the briefings covered the areas of testbed communications, signal and data processor evaluations, analysis of Raytheon hardware, and COLSA's plans for GBR testbed support.

5.0 CONCLUSIONS AND RECOMMENDATIONS

During the past contract period much time was spend in researching the market and evaluating digital hardware that could prove to be of use to the GBR program. Since the final design has not yet been established, all of our research was based on existing designs.

It is necessary to point out that the replacement of proprietary type hardware, as designed in the past, is desirable. It has always been COLSA's opinion that it is preferable for any system such as GBR to use proven digital equipment and computers, for the sake of meeting deadlines, and simplifying designs. The market research performed on digital hardware, is a vehicle which can be used by the prime contractor to finalize and simplify the GBR designs. Many companies in today's market as described earlier, have capable computers which can be used by GBR and consideration should be given to utilizing their capabilities rather than spending money on re-designing already proven systems.