CI-2071-93-0008

DIGITAL HARDWARE ARCHITECTURE IMPLEMENTATION

ANNUAL REPORT

Sponsored by:
PROGRAM EXECUTIVE OFFICE
GLOBAL PROTECTION AGAINST LIMITED STRIKES

15 February 1993

Prepared for:
PROGRAM EXECUTIVE OFFICE
GLOBAL PROTECTION AGAINST LIMITED STRIKES
GBR Project Office
P.O. Box 1500
Huntsville, AL 35807-3801

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

This report has been prepared by COLSA Corporation as part of Contract Number DASG60-91-C-0006 for the Program Executive Office, Global Protection Against Limited Strikes. 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.

FOREWORD

The Ground Based Radar (GBR) digital hardware architecture implementation annual report is presented to the GBR Project Office, Program Executive Office, Global Protection Against Limited Strikes, 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 Space and Strategic Defense Command
ATTN: SFAE-GPL-GBR (Jack Remich)
P.O. Box 1500
Huntsville, AL 35807-3801
Telephone: Autovon: 788-1867
Commercial: (205) 955-4370

TABLE OF CONTENTS

Section Page

1. EXECUTIVE OVERVIEW 1

2. INTRODUCTION 2

3. TECHNICAL EVALUATIONS 3
3.1 GBR Communications 3
3.1.1 Fiber Optics 5
3.1.1.1 Transmitter Receiver Modules 6
3.1.1.2 Metropolitan Fiber Systems 6
3.1.1.3 Asynchronous Transfer Mode 6
3.1.2 Fiber Distributed Data Interface (FDDI) 6
3.1.2.1 FDDI Standards 7
3.1.2.2 Token Ring Operation 9
3.1.2.3 FDDI Media 12
3.1.2.4 FDDI Concentrators 13
3.1.2.5 FDDI Bridges and Routers 14
3.1.3 Ethernet Communications 14
3.1.4 High Performance Parallel Interface (HIPPI) 15
3.1.4.1 HIPPI Layers 15
3.1.4.2 HIPPI Protocol Data Unit (PDU) 16
3.1.4.3 HIPPI Communications in GBR 17
3.1.4.4 HIPPI Testing Devices 20
3.1.4.5 HIPPI Interfaces 21
3.1.4.6 Interface Summary 25
3.1.4.7 Optimization Methods for HIPPI Communications 26
3.1.4.8 HIPPI Conclusions 26
3.1.5 Technology Update On System Buses 26
3.2 Computer Analysis 27
3.2.1 Signal Processor (SP) Analysis 27
3.2.1.1 SP Candidate Comparisons 27
3.2.1.2 MasPar MP-2 28
3.2.1.3 GBR Proposed Design 28
3.2.1.4 MP-2 Footprint 28
3.2.1.5 MP-2 Specifications 28
3.2.1.6 MP-2 Architecture 29
3.2.1.7 MP-2 Communications 30
3.2.1.8 MP-2 Memory System 31
3.2.1.9 MasPar MP-2 Benchmarks 31

TABLE OF CONTENTS — (CONCLUDED)

Section Page

3.2.1.10 MasPar Hardware/Software Relationship 31
3.2.1.11 MasPar Software Statements 32
3.2.2 Data Processor Analysis 35
3.2.2.1 DEC VAX 7000 Footprint 35
3.2.2.2 DEC VAX 7000 Specifications 36
3.2.3 GBR Workstation Candidates 36
3.2.3.1 DECstation 5000/240 37
3.2.3.2 DEC VAX 4000 Workstations 37
3.2.3.3 HP-APOLLO 9000-720CRX 37
3.2.3.4 IBM Powerstation 320 38
3.2.3.5 Silicon Graphics IRIS Indigo 39
3.2.3.6 3-D Graphics Software Applications 41
3.2.4 Microprocessor Evaluations 43
3.2.4.1 DEC Alpha Chip 43
3.2.4.2 INTEL 486DX2 44
3.2.5 Storage Device Forecast 45
3.2.6 High Speed Data Recorders 46
3.2.6.1 Datatape/HIPPI Iinterface 47
3.2.6.2 Datatape LP400 Critical Path 49

4. TRIPS 51

5. CONCLUSIONS AND RECOMMENDATIONS 52

LIST OF FIGURES

Figure		Page
1	Outline of Final Report	2
2	COLSA/ARC Communication Network	4
3	USAKA Communications	5
4	Program Test Site Communications	5
5	GBR SP/DP FDDI Communications	7
6	FDDI Layers	8
7	IEEE 802.5 Token Passing Method	10
8	Token Ring Format	10
9	Ring Operation	11
10	FDDI Self-healing Configuration	12
11	Fiber Optic Cable	12
12	Light Transmission Through a Fiber Optic Cable	12
13	FDDI Concentrator Connections	13
14	COLSA's Ethernet Layout	14
15	Four HIPPI Protocol Layers Specified For GBR Use	15
16	HIPPI PDU (Protocol Data Unit)	17
17	Block Diagram of HIPPI Interfaces	18
18	Model 1200 Fiber Optic Extender	19
19	Possible HIPPI/Fiber Optic System	19
20	CASA Testbed for High Speed WANs	20
21	Block Diagram of Raytheon's REXTTG-HIPPI Interface	22
22	The HIPPI MasPar Interface	22
23	Block Diagram of the HIPPI-MasPar Interface	23
24	Block Diagram of the HIPPI-Datatape Interface	25
25	MasPar MP-2 Significantly Different from MP-1	28
26	The MP-2's Footprint	29
27	MasPar MP-2 Architecture	29
28	Three Modes of MasPar Communications	30
29	MasPar MP-2 Preliminary Benchmark Results	
	(a) Benchmarks for the 4,096 Processor Configuration	31
	(b) Benchmarks for the 16,384 Configuration	32
30	MasPar's Spread Function	33
31	Multistage Interconnection Network	34
32	Physical Characteristics of the VAX 7000	35
33	DEC VAX 7000 Architecture	36
34	DECstation 5000 Specifications	37
35	Silicon Graphics Indigo Dedicated Graphics Board	39
36	Silicon Graphics Indigo CPU Board	40

LIST OF FIGURES — (CONCLUDED)

Figure		Page
37	DEC's Alpha Chip	44
38	Intel's 486DX2 Chip and a Sample Motherboard	44
39	Probable HSR Positions and Roles in the GBR Loop	46
40	Four Leading HSR Systems	46
41	Datatape's (IOSC) HIPPI-VRB Adapter	48
42	GBR Recording Block Diagram	49

LIST OF TABLES

Table		Page
1	ANSI/ISO FDDI Standards	9
2	Computer Manufacturers That Endorse HIPPI	15
3	Maturity of HIPPI Interfaces	25
4	A Comparison of Signal Processor Candidates	27
5	POWERserver Series 300 Specifications	38
6	Microprocessor Comparison	43
7	Hard Drive Forecast	45
8	Floppy Drive Forecast	45
9	Magnetic Versus Optical Storage Devices	45
10	A Comparison of Four Leading HSR Systems	47
11	A Comparison of Leading HSRs by Mission	47
12	Price and Footprint for the Datatape LP-400 Complete System	48

LIST OF ACRONYMS AND ABBREVIATIONS

2D	Two-Dimensional
3D	Three-Dimensional
A	Amperes
ACU	Array Control Unit
ANSI	American National Standards Institute
API	Application Programming Interface
ARC	Advanced Research Center
ATM	Asynchronous Transfer Mode
BIX	Bus Interface Transfer RAMS
BSG	Beam Steering Generator
CDC	Control Data Corporation
CHiL	Channel Interface Logic
CI	Configuration Item
CPU	Central Processing Unit
DAS	Dual Attachment Stations
DEC	Digital Equipment Corporation
DEM/VAL	Demonstration/Validation
DoD	Department of Defense
DRAM	Dynamic Random-Access Memory
DSCS	Defense Systems Communication Systems
ECC	Experimental Control Center
ED	Ending Delimiter
EISA	Extended Instruction Set Architecture
EMC	Electromagnetic Compatibility
EMI	Electromagnetic Interference
EMP	Electromagnetic Pulse
FC	Frame Control
FDDI	Fiber Distributed Data Interface
FE	Front End
FIFO	First-in-First-Out
FS	Frame Status
Gb	Gigabyte
GBR	Ground-Based Radar
GBR-T	Ground-Based Radar-Terminal
HIPPI	High Performance Parallel Interface
HIPPI-FP	HIPPI-Framing Protocol
HIPPI-LE	HIPPI-Link Encapsulation
HIPPI-MES	HIPPI-Mechanical, Electrical, Signaling
HIPPI-PH	HIPPI-Physical

LIST OF ACRONYMS AND ABBREVIATIONS (CONTINUED)

HIPPI-SW	HIPPI-Switching
HSR	High Speed Recorder
HVA	HIPPI-VRB Adapter
HWIL	Hardware in the Loop
I/O	Input/Output
IC	Integrated Circuit
IOSC	Input/Output Systems Corporation
ISO	International Standards Organization
IV&V	Independent Verification and Validation
LAN	Local Area Network
LLRC	Length Longitudinal Check Word
M	Meters
MAC	Media Access Control
MAC	Media Access Controller
Mbps	Mega Bytes per Second
MFLOPS	Million Floating Operations Points Per Second
MIB	Management Information Base
MIC	Media Interface Connectors
MIPS	Million Instructions Per Second
MMU	Memory Management Unit
MPL	Massively Parallel Language
MTTF	Mean Time To Failure
NMD	National Missile Defense
OSI	Open System Interconnection
OTHR	Over the Horizon Radar
PDR	Preliminary Design Review
PDU	Protocol Data Unit
PE	Processor Element
PEXlib	PEX Library
PHIGS	Programmer's Hierarchial Interactive Graphics
PHY	Physical Layer Protocol
PMD	Physical Layer Medium Dependent
RAM	Random Access Memory
RECV	Receive
REXTTG	Receiver/Exciter Test Target Generator
s	Seconds
SAS	Single Attachment Station
SD	Starting Delimiter

LIST OF ACRONYMS AND ABBREVIATIONS (CONCLUDED)

SG	Silicon Graphics
SIA	Scalable Interconnect Architecture
SIMD	Single Instruction Multiple Data
SMT	Station Management
SP	Signal Processor
SPARC	Scalable Processor Architecture
ST	Straight Tip
THT	Token Holding Timer
TMD	Theater Missile Defense
TRT	Token Rotation Timer
TXH	Valid Transmission Timer
UOE	User Operational Evaluation
USAKA	United States Army Kwajalein Atoll
V	Volts
VHD	Very High Density
VMS	Virtual Memory System
VRB	Variable Rate Buffer
WAN	Wide Area Network
WSMR	White Sands Missile Range

1. EXECUTIVE OVERVIEW

During the past year's Ground-Based Radar (GBR) development cycle, COLSA supported the Project Office in many technical areas, which will be described in detail within the course of this annual report.

The report is segmented into five sections: Section 1 - Executive Overview, Section 2 - Introduction, Section 3 - Technical Evaluations, Section 4 - Trips & Meetings, and Section 5 - Conclusions and Recommendations. Section 3 - Technical Evaluations, is the principal focus of this report, describing in detail a number of subjects important to a sound GBR design.

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

The report will describe in detail the probable hardware candidates that were assessed during the contract proposal period, along with the actual GBR configurations, which became evident upon contract award. The subjects to be covered will include descriptions of hardware communications; signal and data processors; workstations; and the latest results of up-to-date microprocessor designs, as advertised by the industry.

2. INTRODUCTION

All work under contract DASG60-91-C-0006 was performed within the technical scope of work, which includes the analysis of GBR digital hardware equipment, in order to ensure mission performance requirements were met. This report will describe in detail COLSA's effort and methodology to achieve this goal during the contract period of April 15, 1992 to April 15, 1993.

An outline of the report contents is depicted in Figure 1.

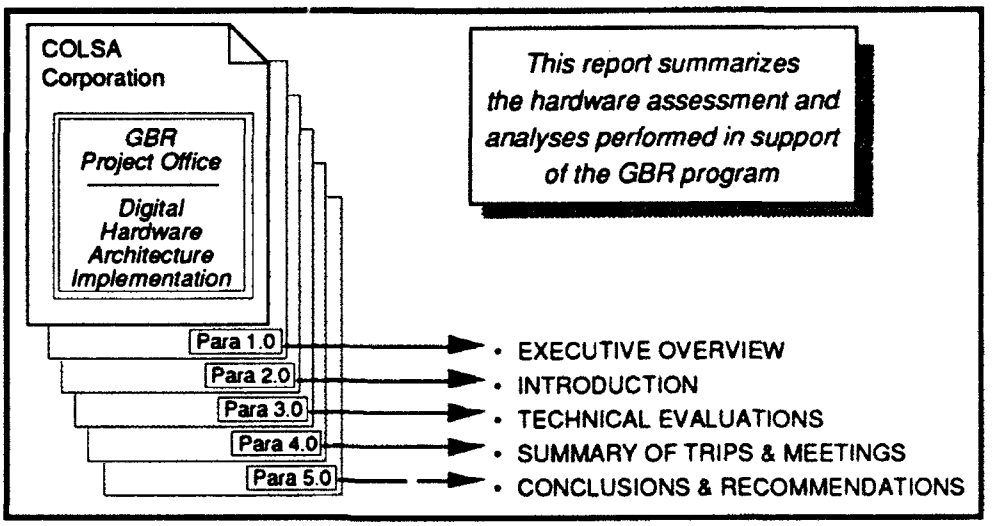

Figure 1. Outline of Final Report

It is anticipated that during the immediate future, COLSA will continue to support the program office with specific evaluations of the digital hardware as they become available throughout the design phase. In order to enhance the GBR system performance as much as possible, it is COLSA's objective to not only accumulate and evaluate the prime contractor's hardware performance specification, but to also introduce new ideas and alternate configuration improvements.

3. TECHNICAL EVALUATIONS

In order to keep the GBR program office informed of the latest digital hardware technologies, COLSA spent a considerable amount of time analyzing as many candidate manufacturers as possible. As a result, the Raytheon digital hardware choices were not only somewhat expected, but were thoroughly presented in a number of COLSA's quarterly presentations to the program office.

The GBR digital hardware configuration as described by the Raytheon designs include the following subsystems: A Receiver/Exciter Test Target Generator (REXTTG), a Beam Steering Generator (BSG), two Data Processors, a signal processor, a number of workstations, a High Speed Recording system and a number of communication links. The REXTTG and the BSG will be of a Raytheon proprietary design but even so, a number of off-the-shelf components and designs will be implemented in these proprietary designs (communication buses, Motorola microprocessors 68020s, etc.) in order to ensure correct system performance. The two VAX 70000 computers (one is used as back-up) made by Digital, will be used as the GBR Data Processor, while a number of MasPar-made computers (four for Theater Missile Defense (TMD) and five for Ground-Based Radar - Terminal (GBR-T) will be used as the radar's signal processing system. The prescribed communication links along with their required interface hardware include Fiber Distributed Data Interface (FDDI), High Performance Parallel Interface (HIPPI), Ethernet and possibly T1 and 9600 Baud Local Area Network (LANs) and Wide Area Networks (WANs), in order to connect the GBR testbed to the future test sites.

All of the areas mentioned above, will be described in detail in this report with particular emphasis on system and subsystem descriptions, design capabilities, throughput, latency, and interface requirements.

3.1 GBR Communications

It will be required that COLSA Corporation install a communications network, to be used in conjunction with the GBR hardware located at the Advanced Research Center (ARC). T1 lines, 9600 Baud lines, Satellite links, Ethernet, and FDDI communications will most probably all be part of the installation efforts in the near future. The ARC will house a number of computers and workstations connected to each other by the use of FDDI rings. The GBR testbed-based computers shall be accessible to employees who are located in COLSA'S Corporate building. These employees will be responsible for performing software development and software Independent Verification and Validation (IV&V) evaluations from their office workstations. Figure 2 depicts such a network.

As shown in Figure 2, the network shall have users connected to it, which are located in both the ARC and COLSA's corporate building. It should be pointed out that the users located in the corporate building will not have access to SECRET information or data whatsoever. It will, therefore, be necessary to install encryption devices on all of the connected workstations and routers between

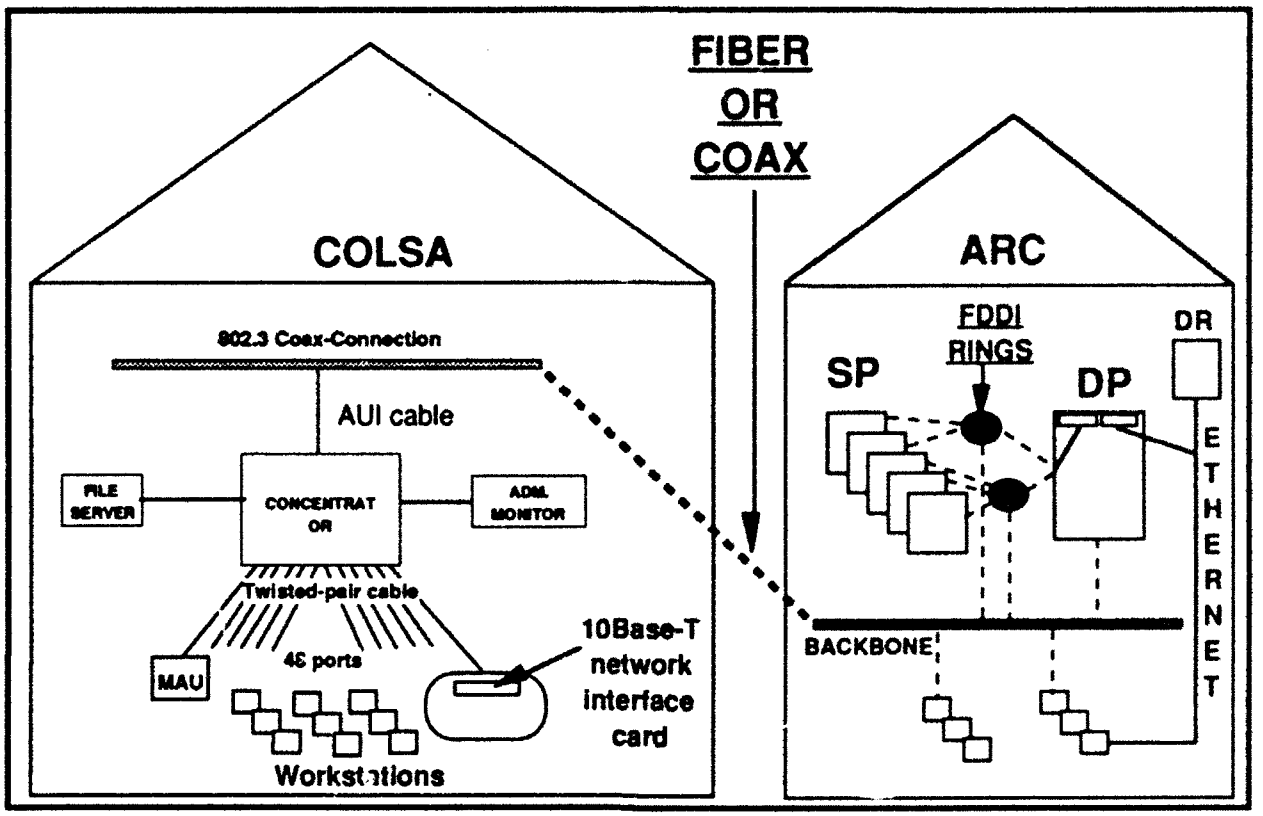

Figure 2. COLSA/ARC Communication Network

COLSA and the ARC. Multilevel security controls will be embedded in the network. It is COLSA's opinion that the network be designed to meet the A1 evaluation criteria of the Department of Defense (DoD) Trusted Computer System Evaluation Criteria. Both single level and multilevel users may share the resources of the network with assurance that the system security constraints will be maintained. A centralized network management node will allow the network and security administrators to control network operation and security. Operational status of the network will be maintained by the network management, while security audits and alarms will be processed at the network management node. Self-tests should be performed at start up and at periodic intervals, to continually verify the system's security level. The network will have the capability to ensure that data is not delivered to devices unless the devices or workstations are authorized for the data, and that data transmitted from a particular device or workstation is within the device's authorized range of security levels. The network should also provide for user authentication, (i.e., the individual workstations that will be located at COLSA's main building) for audits and security alarms, protection against denial of services and assurance of data integrity, along with the capability to remove or add users and devices to the system while in operation.

There is also a strong possibility that additional communication lines will be installed at the ARC in order to connect the testbed with the various test sites around the country. These configurations are depicted in Figures 3 and 4.

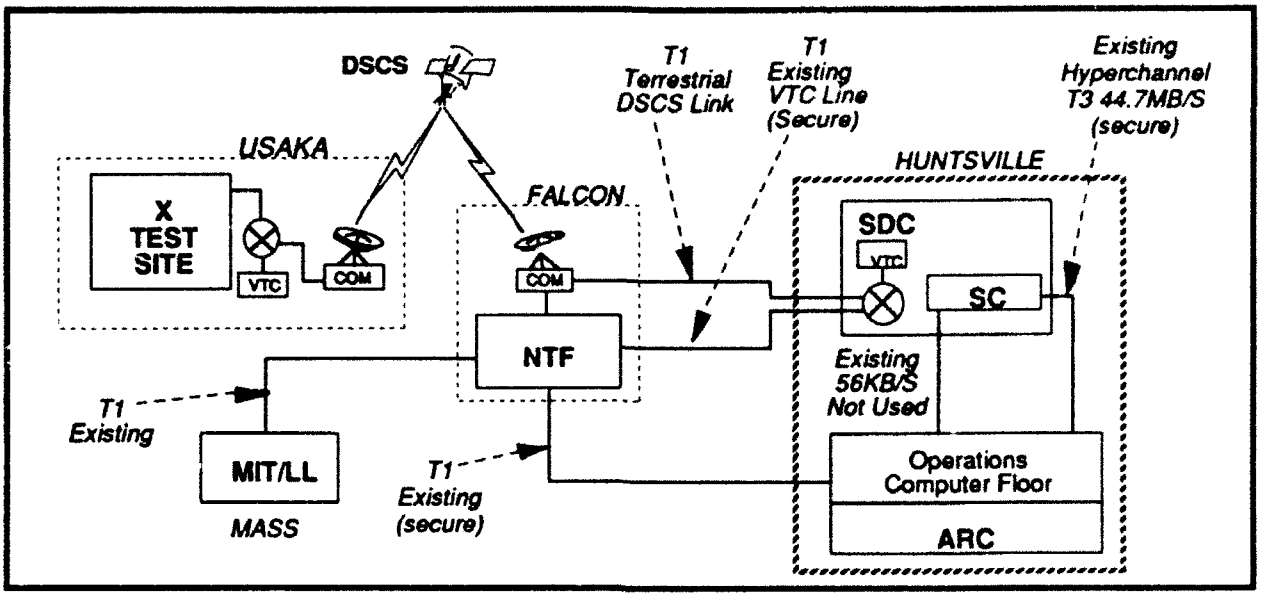

Figure 3. USAKA Communications

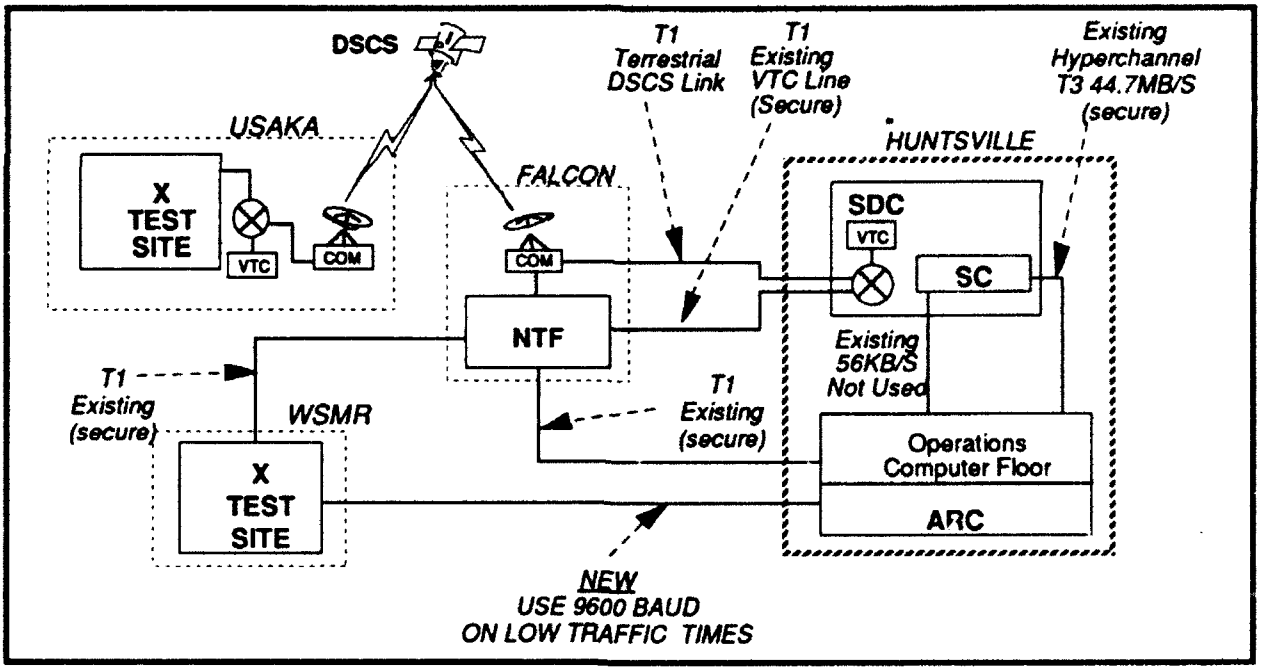

Figure 4. Program Test Site Communications

3.1.1 Fiber Optics

It has been COLSA's opinion that fiber optic links be used by the GBR program when possible, to reduce external noise interference, reduce cost, and increase ease of hardware installation/disconnect. As a result, much time has been spent in evaluating the latest innovations in fiber optic technology. A number of candidate manufacturers and their hardware capabilities will be described in the following pages.

3.1.1.1 Transmitter Receiver Modules

Finisar Corporation has developed fiber optic transmitter and receiver modules with low-cost multimode optical components that can send data at 1.5 Gb/s over fiber as long as 500m. These devices use a proprietary method for modulating laser diodes to transmit the desired data over inexpensive fiber cable. The laser diodes that the modules use as optical source are the same as those found in compact disk players and have a Mean Time To Failure (MTTF) of 200,000 hours. The modules connect to standard 50/125 and 62/125mm multimode fiber. The units are robust enough to tolerate a link having multiple connectors and a mix of fiber types. Data transfer rates range from 100 Mb/s to 1.5 Gb/s over distances as much as 500m. The bit error rate is <10-12 with a link-loss budget of 10 db. The controller Integrated Circuit (IC) monitors the status of the transmitter and receiver modules. It provides signal compensation such as receiver gain control and transmitter bias voltages. The link controller also provides a system-test interface that lets you monitor the data link's status and performance without using special optical test equipment. Both modules directly convert between electrical and optical signals, so you can use either encoded or NRZ data formats. Both modules use 5V and 12V supplies and draw 300mA at 5V and 40mA at 12V. They measure 0.685 x 0.870 x 2.725 in., and have pin connectors that can be socketed or soldered to a pc board. The transmitter module costs $310, the receiver module costs $300, and the controller IC costs $60.

3.1.1.2 Metropolitan Fiber Systems

A new service provided by Metropolitan Fiber Systems will allow computers using FDDI communications to be interconnected on urban "fiber-optic highways." Data rates will be in the 100 Mb/s range. Focus on FDDI will increase the degree to which it is industry proven, and will likely lead to improvements on the technology; both of these factors will be helpful if GBR is to employ FDDI technology.

3.1.1.3 Asynchronous Transfer Mode

FDDI communications may be seriously challenged by a standard called Asynchronous Transfer Mode (ATM). The new standard, sponsored by the Naval Research Laboratories, boasts a transfer rate of 155 Mbps as compared to the 100 Mbps of FDDI. ATM adapter cards are already available for Sun and Digital Equipment Corporation (DEC) systems at a cost of approximately $5,000.

3.1.2 Fiber Distributed Data Interface (FDDI)

Computers have revolutionized not only military systems but modern life in general. There is an enormous amount of information exchanged between the computers that control military equipment and personnel with access to high speed workstations. Reliable data communications are, therefore, very much in demand. Fiber optic networks are continually being developed and updated so as they can operate equally well in the harshest of environments, as well as benign office conditions. Fiber optics have become the medium of choice in many applications which demand more than restricted

coaxial or twisted cable capabilities. These include hazardous or electrically noisy environments, high security areas, networks with greater bandwidth requirements, etc. Fiber optic cost is decreasing as production increases making fiber optic communications much more competitive with the coaxial type links.

As depicted in Figure 5, the GBR program will attempt to take advantage of these FDDI merits, by including in its network two FDDI rings, a possible fiber optic backbone and an undetermined as of yet, length of optical cable. Additional FDDI equipment which will be described in the report will also be required to ensure that information will flow smoothly to all connected LAN users. At 100Mb/s, FDDI is the fastest standard available today and is quickly becoming the network of choice in large network installations.

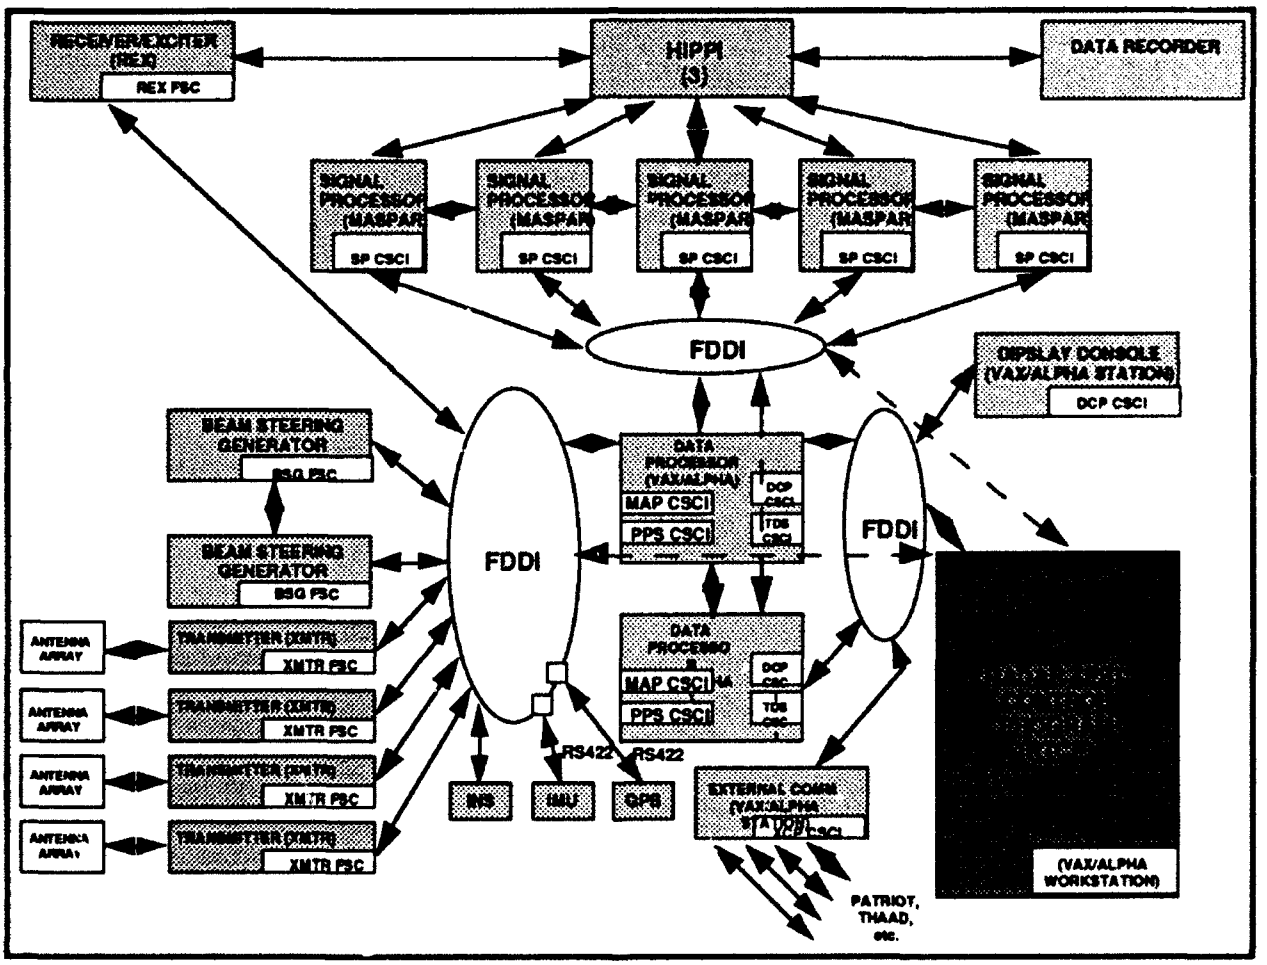

Figure 5. GBR SP/DP FDDI Communications

3.1.2.1 FDDI Standards

Network systems such as FDDI, are usually designed based upon a *layer* system. These layers perform individual services within the network. The combining of layers is known as *Network Architecture*.

System Architecture dictates how to sort communication functions into logical groups (layers), in order for them to perform peer-to-peer functions. Each layer has its own *Protocol*. A protocol can best be defined as a set of rules that a layer is obligated to follow in order to regulate activity within the layer. Protocols also direct communications and data transfer between layers, in a directly above or below hierarchy. The *Open System Interconnection (OSI) Reference Model*, which is shown in Figure 6, follows a seven layer open network architecture. *Open Network Communications* are networks with the ability to transmit information to multivendor products, a very important principle for designers and users alike. FDDI networks utilize three layers: the *Physical Layer*, the *Data Link Layer* and lastly, the *Network Layer*. *The Physical Layer* specifies the physical interface between nodes. Part of the components of the physical layer are connectors, type of wire to be used, etc. *The Data Link Layer* manages the transmission of data by introducing control information into the messages, to be distributed between adjacent nodes. The Network Layer assists in the communication of open system network nodes. It provides for the node addressing and routes frames between even dissimilar networks.

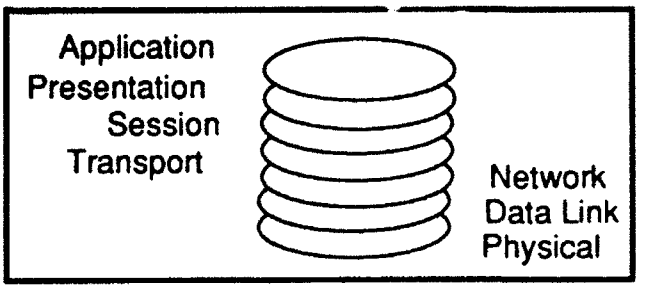

Figure 6. FDDI Layers

FDDI is a set of LAN standards. The standard was developed under the procedures of the American National Standards Institute (ANSI). *FDDI offers an industry* standard solution for organizations that need flexible, robust, high performance networks. It is a 100 Mb/s timed-token-passing LAN, constructed of two independent counter rotating rings that are connected together in case of ring failure. FDDI can connect up to 500 stations for a total fiber distance of 200 km. The FDDI standards provide the required rules for manageability, high speed, low error rates, and high fault tolerance. *Multimode* fiber optic cable was the first transmission medium defined in the FDDI standard, followed by the *single mode* fiber cable. Support for copper media is under development by the ANSI Committee.

The following ANSI standards (also shown in Table 1), define the components of FDDI: (1) *Physical Layer Medium Dependent* (PMD) standards define the medium and the connectors for multimode and single mode fiber optic cable. (2) *Physical Layer Protocol* (PHY) standard defines the rules for encoding of data. (3) *Media Access Control* (MAC) standard defines the protocols for operation of token ring and construction of frames and tokens. (4) *Station Management* (SMT) standards define the protocols for managing the PMD, PHY, and MAC entities. A large portion of the FDDI standard dictates the operation of the SMT software. Functions such as connection management, configuration management, and fault isolation/recovery are performed with the use of SMT software. The software also defines the operations to be performed on managed objects, (get, change, remove, add, etc.) and how certain events are to be handled. As of today, the latest version of SMT is 6.2 and is always evolving. SMT also defines the *Management Information Base* (MIB). A MIB can contain information such as the number of ports in a station, most recent station configuration, number of frame errors, and so on.

TABLE 1. ANSI/ISO FDDI STANDARDS

ANSI/ISO STANDARD	FDDI STANDARD
PMD:X3.166-1990/ ISO 9314-3:1990 SMF-PMD:X3.184-1991 LCF-PMD:TBD TP-PMD:TBD	PHYSICAL LAYER MEDIUM DEPENDENT: Four standards correspond to the lower portion of the OSI Physical Layer and define the transmit/receive power levels, transmit and receive interface requirements, error rates, and cable and connector specifications. These standards are: • Physical Layer Medium Dependent (PMD). • Single Mode Fiber Physical Layer Medium Dependent (SMF-PMD). • Low Cost Fiber Physical Layer Medium Dependent (LCF-PMD). • Twisted Pair Physical Layer Medium Dependent (TP-PMD) under development.
X3.184-1988/ ISO 9314-1:1989	PHYSICAL LAYER PROTOCOL: This medium independent standard corresponds to the upper portion of the Physical Layer. PHY defines symbols, linestates, encoding/decoding techniques, clocking requirements, and data framing requireennts.
X3.139-1987/ ISO 9314-2:1989	MEDIA ACCESS CONTROL: This standard corresponds to the lower portion of the OSI Data Link Layer. MAC defines data link addressing, frame formatting, media access, error detection, and token handling.
X3T9.5/84-49	STATION MANAGEMENT: This standard defines the system management services for the FDDI protocols. SMT includes facilities for the connection management, node configuration, recovery from error conditions, and the encoding of SMT frames.

The FDDI standard also defines several types of networking devices that include concentrators, dual attachment stations, and single attachment stations. *Concentrators* gather several lines in one central location. They attach directly to the FDDI dual ring, that is, to the primary and secondary. They also provide highly fault tolerant connections to the ring. *Dual Attachment Stations* (DAS) devices connect directly to the dual ring or to a concentrator, where as *Single Attachment Stations* (SAS) devices connect to the ring only through a concentrator. Extended LANs are typically created by connecting LANs together with bridges and routers.

3.1.2.2 Token Ring Operation

The IEEE 802.5 standard as shown in Figure 7, is based upon the token ring technology. Ring initialization and token generations are done by a master system. A token grants its holder the right to transmit for a period of time defined by the token holding rules. The token simply circulates around the ring when no transmissions are being performed between stations.

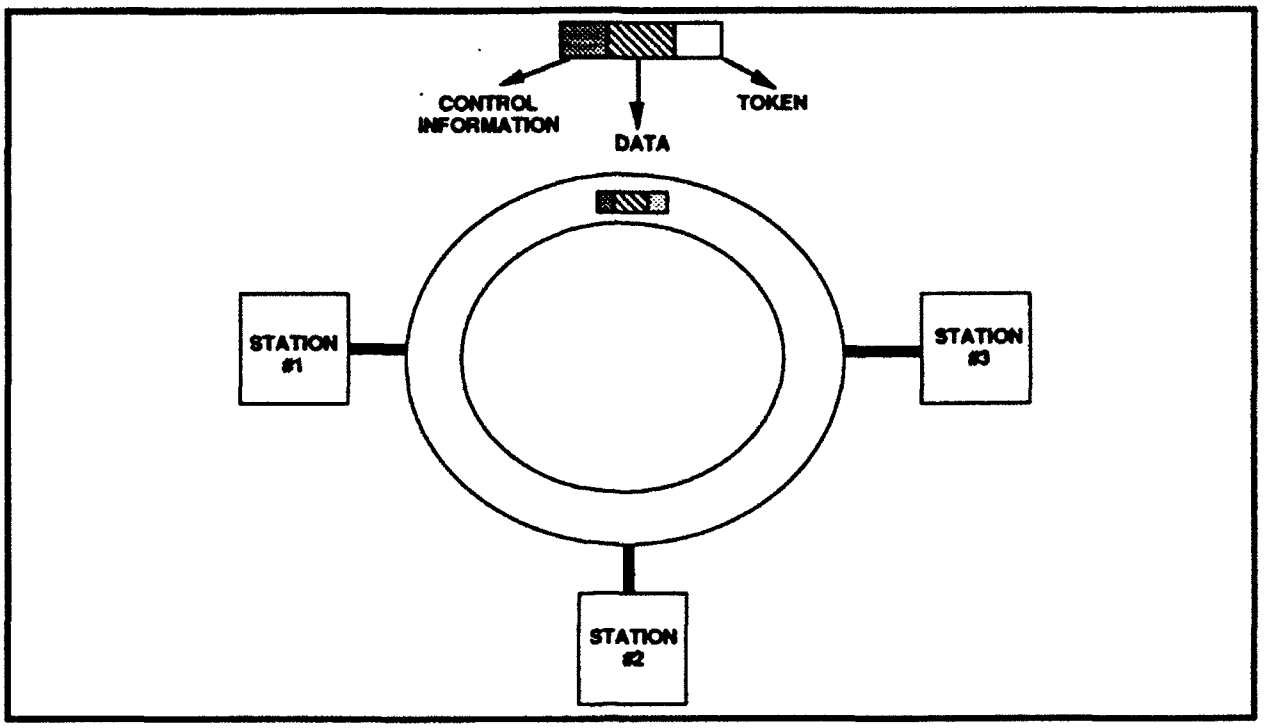

Figure 7. IEEE 802.5 Token Passing Method

The format which is used by the 802.5 token ring is depicted in Figure 8. The frames always begin with the *Starting Delimiter (SD)* and end with the *Ending Delimiter (ED)*. When both are used at the same time, this is interpreted as an abort of transmission condition. When T bit is set to zero (within the *Access Control field*), the token is captured and the station can transmit its particular frames to the destination station. The M bit is used by the active network monitor, the selected station that makes sure that the ring is functioning properly. The *Frame Control (FC)* uses its first two bits to indicate the type of data present at the information field. The last six bits dictate information pertinent to the ring management operations. The destination and the source bytes respectively describe the destination address for which the information is destined, while the source address describes the address from which the information originated. The *Frame Status (FS)* byte contains the A and C flags, which are used for additional addressing and copy the frame type information.

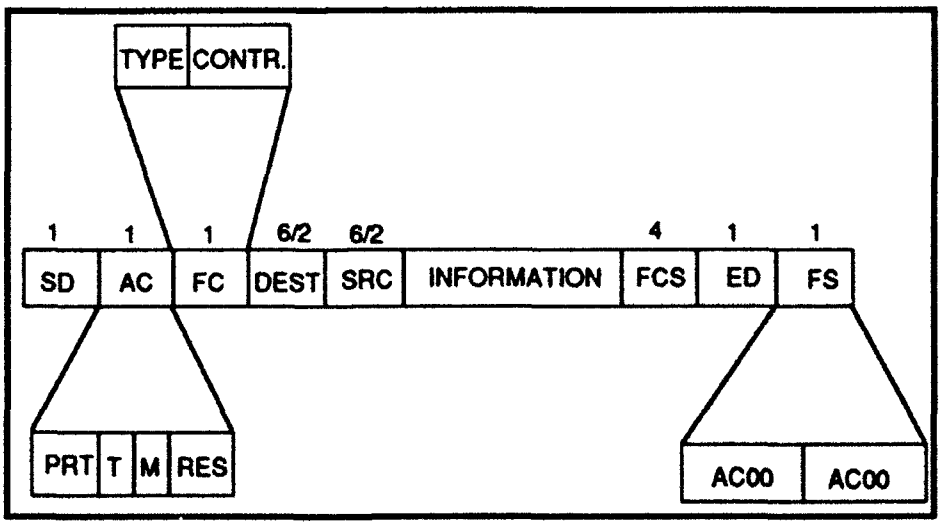

Figure 8. Token Ring Format

Similarly with the IEEE 802.5 token ring operation, the FDDI ring operation depicted in Figure 9, can be described in four easy steps: (1) Station A has a frame of data F1 to transmit so it captures the token, (2) station A begins transmitting frame F1 destined for station D, (3) at completion of transmission, station A issues token while station D copies the particular frame F1, (4) the next station which has a massage, will then capture the token and the transmission cycle is repeated.

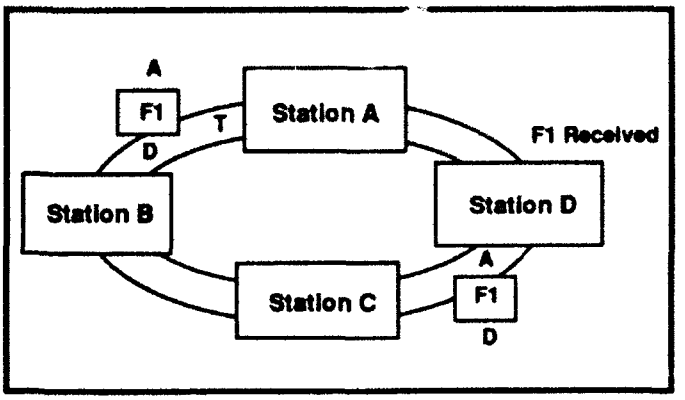

Figure 9. Ring Operation

The FDDI ring operation consists of making the *connection establishment, ring initialization, steady-state operation* and finally, *ring maintenance*. In order for these functions to be regulated, timers are used. Each connected station uses three timers to regulate its operation. The *Token Rotation Timer* (TRT) is used to time the duration of operations within the particular station. This timer controls ring scheduling during normal operation and fault recovery when the ring is idle. The *Token Holding Timer* (THT) controls the duration of time the station can initiate asynchronous frames. A station holding the token can transmit asynchronous information only if the THT has not expired. The *Valid Transmission Timer* (TVX) times the period between valid transmissions on the ring. TVX also has the capability to detect excessive ring noise, token loss, and other such type faults. Each time the station receives a valid frame or token, the valid transmission timer resets. When TVX expires, the station starts the ring initialization sequence so as to restore the ring to its normal operation.

At *ring power up*, the connected stations must establish connections with their neighboring stations. A link is established after a station (a) exchanges information on port type and connection rules, (b) negotiates the length of the link confidence test, which checks the quality of links between its neighboring stations, (c) runs the link confidence test for the negotiated time, and (d) exchanges results.

The steady-state operation of an FDDI *timed-token ring* simply follows the same procedure as that of IEEE 802.5 token ring, with one important exception. The token is available to each station for an agreed upon time limit. This time limit is negotiated by the stations each time an additional station is connected to the ring. In FDDI, a station transmits by first capturing the token, then transmitting as many frames as are allowed by the token rules, and finally, reissuing the token. Only one token can be present on the ring at one time, but multiple frames which form multiple systems or stations can be on the ring. The timed-token passing technique is an efficient use of bandwidth.

To minimize delay, the station reads and repeats the frame fields as it receives them. The first part of the frame contains frame class information followed by the destination address, source address, and finally the data. When a station realizes it is being addressed to, it strips the remainder of the frame from the ring. But the first part of the frame has already been repeated, thus creating the so

called frame fragment. These fragments are removed by the operations of a repeater filter of each station's PHY or by a transmitter.

Ring Maintenance is performed by all of the connected stations. They monitor ring integrity, ring inactivity longer than the valid transmission time, or a physical/logical break.

Self-healing is one of the most important aspects of this topology. One ring is used as the primary, while the other is the system's backup. When a failure occurs, as shown in Figure 10, the active devices on either side of the failure use the backup ring to *wrap around* on themselves, thus isolating the failed component.

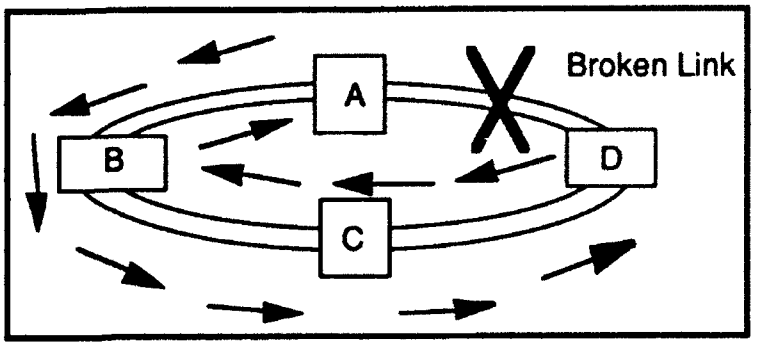

Figure 10. FDDI Self-healing Configuration

3.1.2.3 FDDI Media

The wavelengths most suitable for optical communications are called windows. These windows are 850nm, 1300nm and 1550nm. Multimode fibers operate at 850nm and 1300nm. Multimode fibers have a core diameter of 62.5 micrometers and a cladding diameter of 25 micrometers and is referred to as 62.5/125. Only the three sizes of 62.5/125, 50/125, and 100/140 microns are used by today's industry. Figure 11 depicts a typical Fiber Optic Cable.

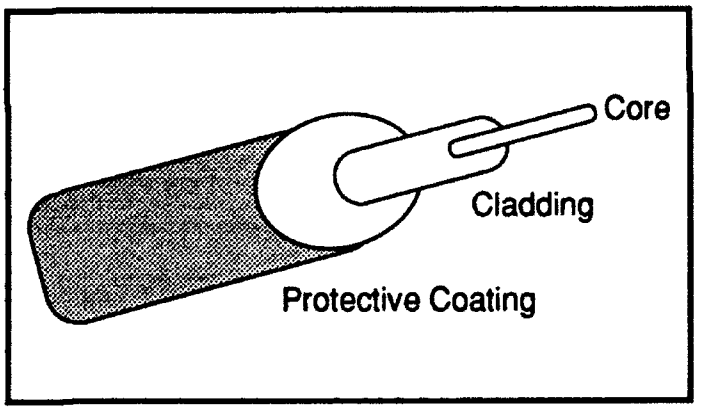

Figure 11. Fiber Optic Cable

FDDI specifications recommend 62.5 micron, graded index multimode fiber. Fiber optic cables can be of multimode or single-mode design. Multimode fibers, as shown in Figure 12, allow multiple rays (frequencies) to be transmitted at the same time, while single mode fibers allow only one light frequency to pass through them. The latter can be used for the longer distance applications of up to 60 km. Graded-index fiber, through its construction, allows for optical pulse correction.

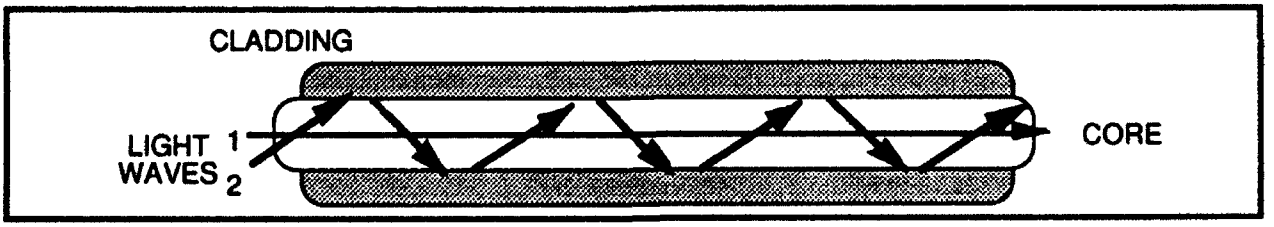

Figure 12. Light Transmission Through a Fiber Optic Cable

The FDDI specifications also define the type connectors fiber optic cables are terminated with. There are two types of connectors. The Media Interface Connectors (MICs) consist of two terminated fibers within a protective shroud. These are keyed connectors that prevent the mixing of the primary and secondary rings and the input and output ports. The Straight Tip (ST) bayonet connectors are of a BNC configuration. STs are more difficult to install and do not protect the glass ends with a shroud as do MICs. Additional hardware design concepts such as attenuation, power loss, and link loss are all important factors and should be taken into account when designing an FDDI network. *Attenuation* describes the amount of optical power lost during the transmission. It is expressed in DBs and can be calculated by knowing the unit attenuation and link length (the longer the cable the higher the loss). The *power budget* of a system is determined by the minimum transmitter power and minimum receiver sensitivity. It is good design practice to remember that power is lost when light passes through connectors and splices. The *link-loss budget* is the total amount of loss that can be introduced into a working optical system.

3.1.2.4 FDDI Concentrators

Concentrators require a single connection to the main trunk counter rotating ring, with station connections through its single attachment ports. These single attachment ports are named the A and B ports. The station ports are called M for master, and the connected stations are referred to as S for slaves. Figure 13 depicts the various port connections of a typical concentrator.

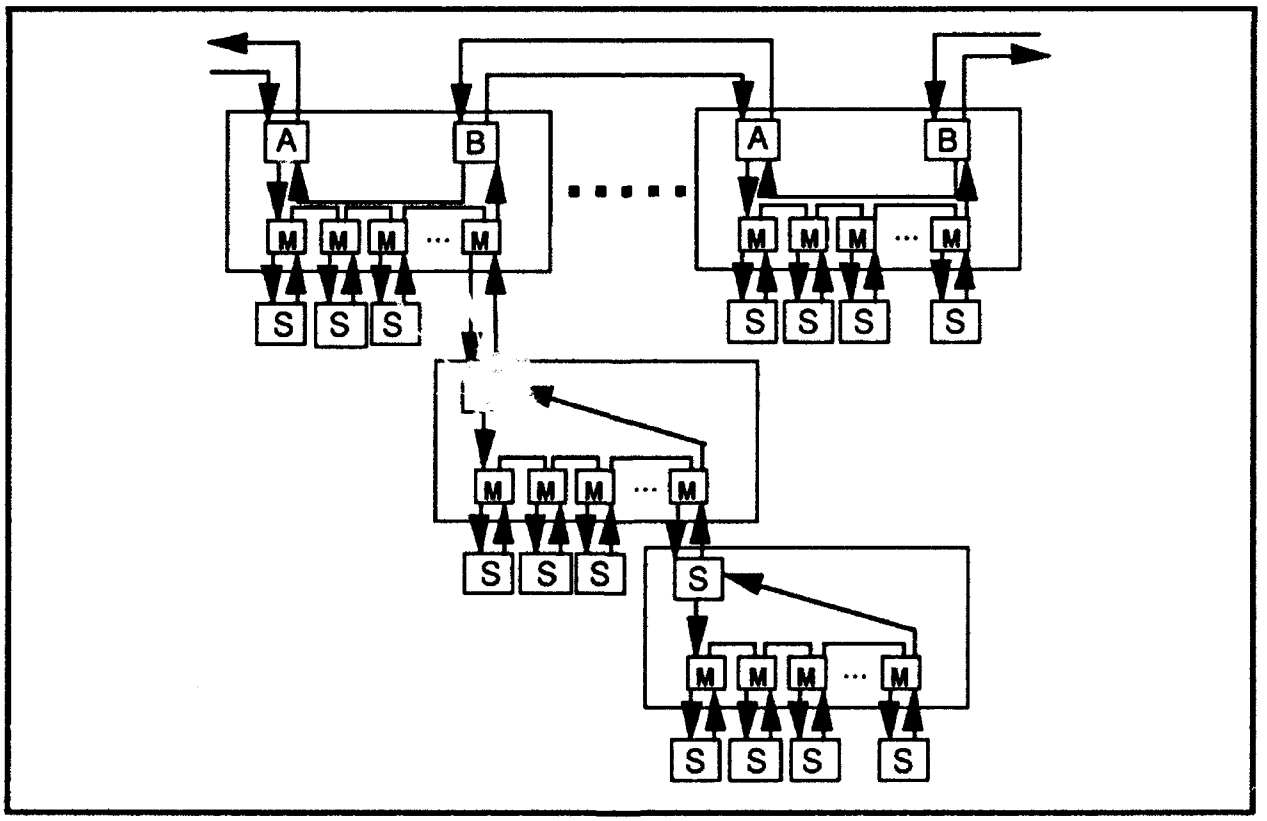

Figure 13. FDDI Concentrator Connections

One important function of the concentrator is its ability to remove the accessibility of the dual ring to every user, thus reducing the risk of affecting ring integrity. Also, with the use of a concentrator, only the station port is affected when its power is turned off while the rest of the network remains operable. The concentrator will furthermore provide a central point for monitoring and maintaining the FDDI network. Cascading a number of concentrators can form what is called the dual ring of trees topology, a configuration which allows the creation of large sophisticated networks, when future expansion requires it. In this topology, the concentrators electronically bypass ring disruptions. Concentrators usually require one Media Access Controller (MAC), the module that understands the FDDI packets which the concentrator will interpret and respond to network management packets. As a result, at least one of the plug-in boards must include one MAC. This board is connected to the next higher level in the structured cabling system.

The GBR FDDI communication design calls for a DECconcentrator 500. This hardware provides for the attachment of FDDI devices such as the program's workstations to the network ring. It accommodates up to three 4-port or 6-port modules and uses modular components for alternate topologies.

3.1.2.5 FDDI Bridges and Routers

When migrating from IEEE 802 networks such as 10Mbps Ethernet and 16 Mbps Token ring to LANs with 100Mbps capability, converting devices such as bridges and routers are required. Even though the features of bridges and routers overlap, they are still distinctly different. Use of bridges is required for LAN-to-LAN interconnection if the requirements call for low delay, high throughput and when non-routable protocols are used. Routers on the other hand, can be used if the LAN-to-LAN interconnection requires a high degree of isolation, or flow control is required. FDDI bridges are for the most part hybrid bridges that implement many of the same features available in multiprotocol LAN-to-LAN routers like protocol filtering, source address filtering, packet fragmentation, and translation.

3.1.3 Ethernet Communications

Ethernet networks perform data exchanges between digital devices such as the workstations and computers to be used by the program. COLSA's Ethernet configuration shown in Figure 14 will utilize the Physical Layer and the Data Link Layer. The Physical Layer will allow data transfer rate of 10Mb/s, at distances of up to 2.5 kilometers with a maximum station number of 1024. It will use shielded coaxial cable base band signaling and will support a star topology. The Data

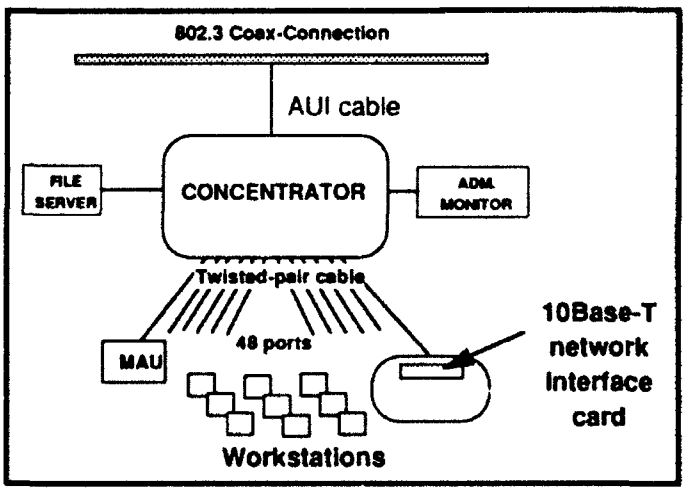

Figure 14. COLSA's Ethernet Layout

Link Layer will utilize a link control procedure fully distributed peer protocol and a massage protocol to support variable size frames.

3.1.4 High Performance Paralle. Interface (HIPPI)

The HIPPI Communications Protocol provides a workable, uncomplicated, relatively inexpensive standard for high-speed point-to-point communication over short distances. HIPPI has been endorsed by a large number of computer manufacturers and users (Table 2).

TABLE 2. COMPUTER MANUFACTURERS THAT ENDORSE HIPPI

MasPar	DEC	Silicon Graphics
Cray	IBM	Thinking Machines
Sun Microsystems	Intel	Convex

HIPPI has 80 to 160 times the speed of Ethernet communication (10 Mb/s). It supports data rates of 800 or 1600 Mb/s, when in 32 bit or 64 bit mode respectively. For comparison, ethernet data rates are 10 Mb/s and FDDI data rates are 100 Mb/s. HIPPI uses twisted pair copper cables for distances of up to 25 meters. Extensions allow connections of up to 10 km and switching can be used to manage communications among two or more computers.

3.1.4.1 HIPPI Layers

The HIPPI communications protocol contains several layers. Figure 15 describes four of the layers that have been specified for GBR use (e.g., the MasPar-HIPPI interface utilizes these four layers).

HIPPI-LE, the Link Encapsulation layer, is a higher level protocol which allows IEEE 802.2 and ISO 8802.2 protocols to be used with HIPPI in a way that is transparent to the implementor of the IEEE

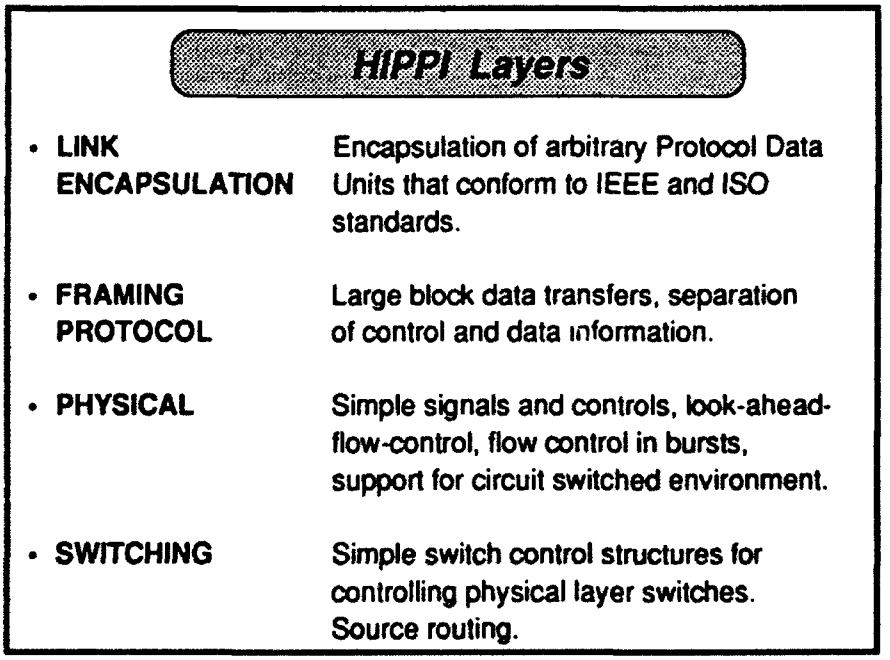

Figure 15. Four HIPPI Protocol Layers Specified For GBR Use

or ISO standard. The link encapsulator encompasses the extraneous protocol inside a HIPPI protocol for transmission on a HIPPI cable; the package is disencapsulated to IEEE/ISO when it completes the HIPPI leg of its journey.

HIPPI-FP, the Framing Protocol, specifies the format for data/information to be transferred on HIPPI lines. The FP is characterized by separation of control information from data for clarity, simplicity, and accuracy. Further, the protocol allows large block data transfers which decrease on source-destination communications overhead.

HIPPI-PH, the Physical layer, specifies the actual cable types to be used and pin connection diagrams. This layer also finishes the work of the framing protocol by specifying the lowest levels of data formatting (see Figure 16). The PH is characterized by simple signals and controls which allow for minimum "handshaking" between source and destination, and maximum data transfer. A key to the high data rates of HIPPI is the use of look-ahead flow control. This mechanism allows the source to fill the destination's buffer with several blocks of data without having to wait for a confirmation (an "okay I've got it, now send another block") from the destination. This cuts down on overhead.

HIPPI-SW, the Switching Layer, specifies the protocol for controlling physical layer switches. The switches provide HIPPI with multidrop capability; in other words a switch allows multiple copies of data/information to be "dropped" at multiple locations (specifically multiple MasPars, and the High Speed Recorder (HSR)). The SW is characterized by simple switch control structures for timely distribution of HIPPI data/information.

3.1.4.2 HIPPI Protocol Data Unit (PDU)

A HIPPI PDU consists of all control information, along with the data, that is necessary for a transmission. Figure 16 is a graphic representation of a HIPPI PDU. At the highest level, there is the *Disabled* state in which no logical connection exists. A prospective source will send an *I-Field* when he would like a connection. The I-Field contains information such as who the sender is, what the nature of the data is, and how much data is to be sent. If the destination is not busy he sends an acknowledgement and a *Connection Established* state is entered, now the source may begin to send *packets*. HIPPI protocol requires *wait* states between certain packets; the states last 1-3 40ns clock cycles at a minimum; however, the HIPPI implementor most likely will use more than the minimum. Packets contain *wait* states, *bursts* and Length Longitudinal Check Words (LLRCs). The wait states are required to last 1-3 40ns clock cycles as with the packets. The bursts contain 256 words of 32 bit length (64 is not applicable to GBR). The LLRC is a data integrity check implemented at the completion of each burst transmission.

An outstanding feature of the PDU is the ability to send an immense amount of data within a single packet; 4 GBytes of data, which is the equivalent of 16 "megabursts." Sending large amounts of data in a single packet minimizes overhead and increases throughput.

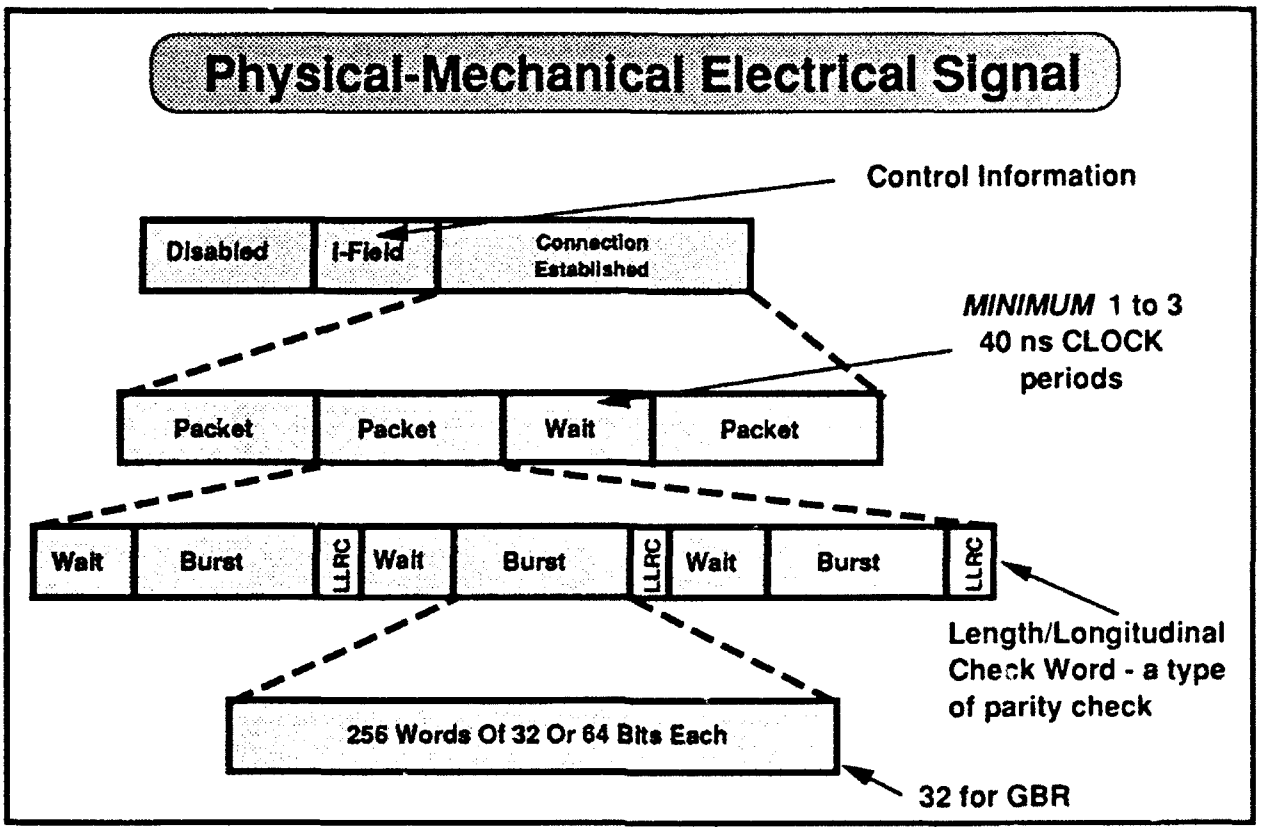

Figure 16. HIPPI PDU (Protocol Data Unit)

3.1.4.3 HIPPI Communications in GBR

The HIPPI Protocol plays a critical role in the GBR program. In this section, several aspects of HIPPI are analyzed including interfaces, a device used to test interfaces for ANSI compliance, a device used to analyze efficiency, a device used to drive tests, and techniques for optimizing implementations.

The three main conclusions of the analysis are as follows:

(1) There will be three relatively unproven HIPPI interfaces in the GBR loop, all designed and manufactured by different sources: Raytheon, MasPar, and either Triplex or Input/Output Systems Corporation (IOSC) depending on the HSR choice. The interfaces are not very complex; however, the immaturity of these products may produce reliability concerns.

(2) The most crucial aspect of HIPPI interfaces in GBR seems to be reliability, rather than performance optimization.

(3) There are three products that are needed to enhance COLSA's ability to perform HIPPI testing on the testbed:

- The HIPPI Tester (IOSC) - verifies compliance, reliability.
- The HILDA (IOSC) - allows performance analysis of connections.
- The HIPPI "Blaster" (IOSC) - source for driving tests.

An overall block diagram of the interfaces is shown in Figure 17.

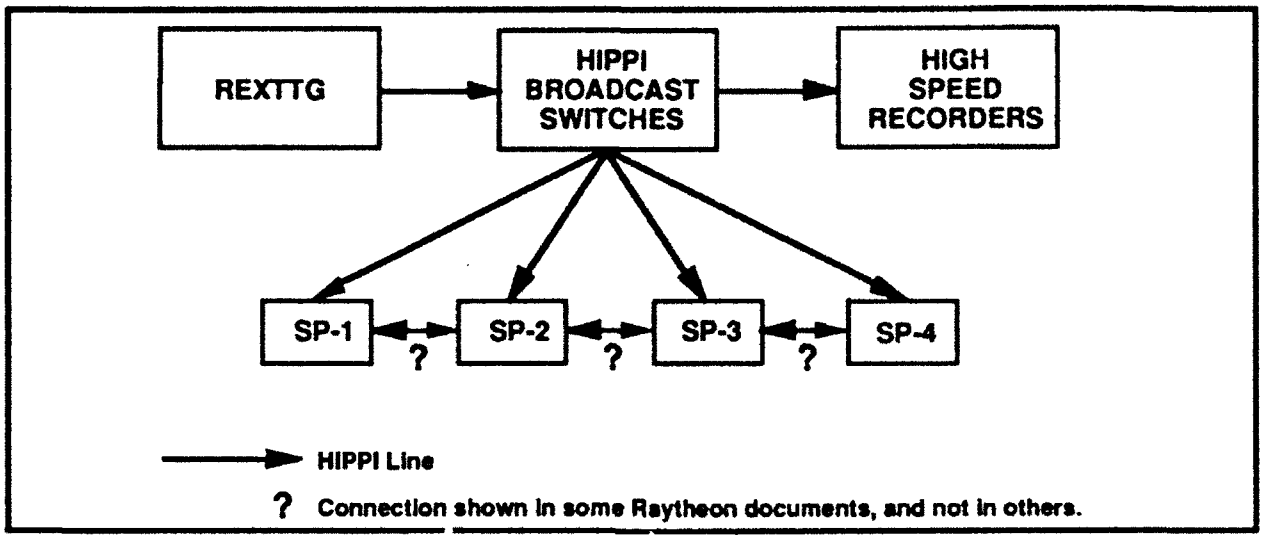

Figure 17. Block Diagram of HIPPI Interfaces

The REXTTG-HIPPI interface is being designed by Raytheon, the HIPPI-Signal Processor interface is being designed by MasPar, and the HIPPI-HSR interfaces are being designed by IOSC and Triplex for Datatape and Sony, respectively.

3.1.4.3.1 Reliability. There are two concerns within the reliability issue that have been examined. The first is the reliability of the HIPPI Protocol itself, due to the fact that it is relatively new. While it is a new standard, contacts at intensive HIPPI centers such as Cal-Tech and Los Alamos have stated that it is very reliable. In addition, it has been implemented as an ANSI standard and, therefore, has been subject to intense review by their committees. COLSA's opinion is that the relative immaturity of HIPPI as a standard is a low risk consideration for GBR.

The second concern is the design and manufacture of three new HIPPI interface cards by three separate companies. The interfaces are not particularly difficult in terms of logic design. However, there are reasons for concern: (1) A logic error in hardware could go undetected during in-house testing of an interface, or (2) A logic error in software could go undetected during in-house testing of an interface. In either case, the functionality of the GBR system would be threatened. Therefore, COLSA's opinion is that there should be additional testing of the HIPPI interfaces at the testbed.

3.1.4.3.2 Optimization of Data Rates. The degree to which HIPPI data rates are optimized will probably not be significant to GBR. The Hardware Configuration Item (CI) Preliminary Design Review (PDR) documents give an average data rate of 26 MB/s coming from the REXTTG. This is only 26% of HIPPI's specified 100 MB/s capability.

3.1.4.3.3 HIPPI Extender. Broadband Communication Products makes an acceptable HIPPI extender. It is totally transparent at distances up to 10 km. Extensions up to 30 km will be available soon. A typical dual-simplex, fiber optic HIPPI extender model 1200 is shown in Figure 18.

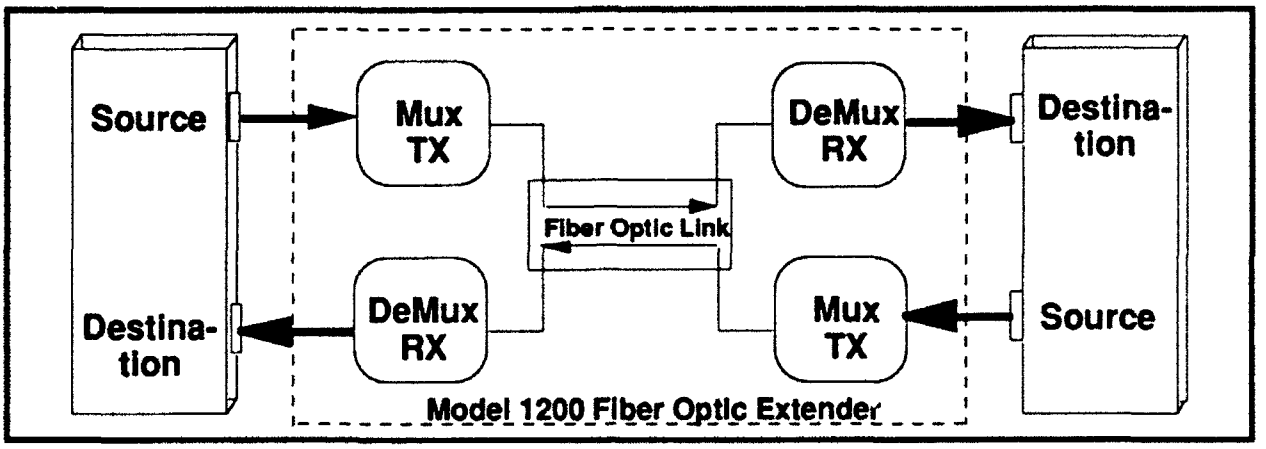

Figure 18. Model 1200 Fiber Optic Extender

Optical fiber, mentioned before, is increasingly the medium of choice for communication systems. Plug -and- play operation is the logical thing to do whenever possible, and fiber optics can provide that capability. One such GBR system that incorporates HIPPI and fiber optics is depicted in Figure 19.

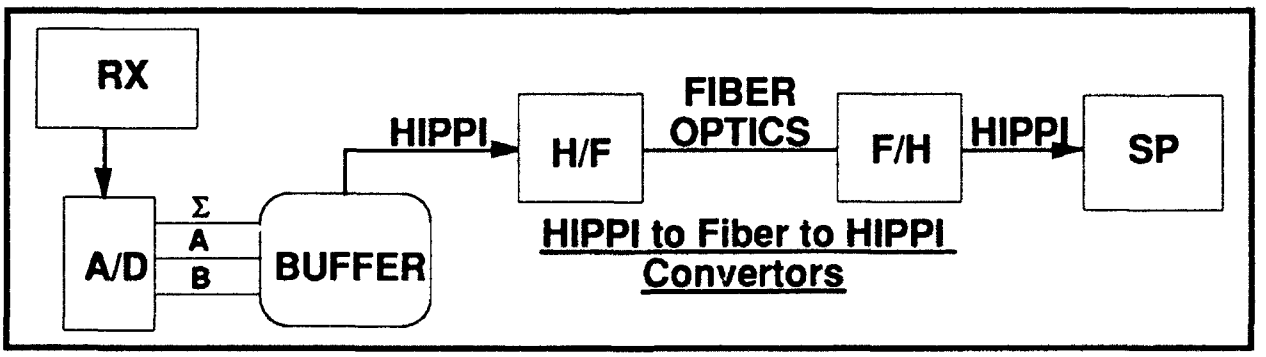

Figure 19. Possible HIPPI/Fiber Optic System

3.1.4.3.4 HIPPI Testbeds. A number of testbeds have been surfacing in the U.S. which are approaching gigabit transmission rates in WAN formats. The testbed shown in Figure 20, known as CASA, tests high speed networks for reliable operations over long distances. All computing systems have HIPPI interfaces with HIPPI SONET gateways being presently developed at Los Alamos. SONET may become available through telephone connection.

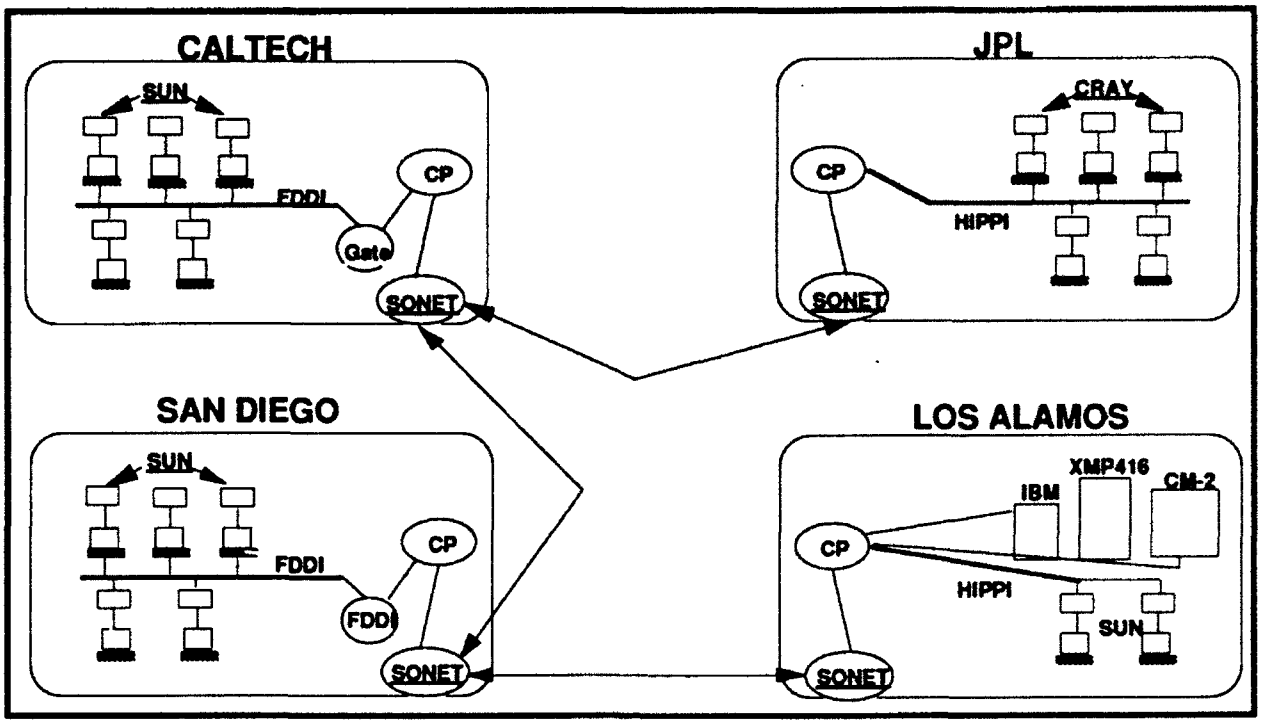

Figure 20. CASA Testbed for High Speed WANs

Another point of interest, is the joint venture between Cray Research and the Regional Computer Center of Stuttgart Germany; they demonstrated a transmission rate of 140 Mb/s over a 600km range (from Stuttgart to Berlin).

3.1.4.4 HIPPI Testing Devices

There are three HIPPI testing devices that are manufactured by IOSC, all of which would be very valuable for testing HIPPI components.

The **HIPPI Tester** is a device designed to support debug, compliance, and reliability of HIPPI-based computer peripherals. This device is capable of generating an ANSI compliant suite of signals, and it is capable of generating non-compliant signals to test robustness of implementations.

The Tester has several other capabilities. The Tester includes HIPPI Tester hardware, HIPPI Tester software, and the Los Alamos National Laboratory Test Suite that is used to verify ANSI compliance. The base price is $7,500.

The **HILDA** HIPPI Network Analyzer is designed to perform analysis on HIPPI networks. It measures link utilization, packet length distribution, packet interarrival time distribution, and many other parameters. It also provides the capability of inserting errors into data streams to test the reaction of interfaces to non-compliant formats.

The HILDA includes either one or two VME boards and HILDA Network Analysis software. A Sun Sparcstation is needed to run the software; the GBR program at COLSA does have a Sun Sparcstation available. With one VME board, the base price of HILDA is $9,000 and with two boards the price is $18,995.

The **HIPPI Blaster** is a device that generates and sends large amounts of "canned" HIPPI data. It is not clear yet whether the same task could be accomplished by the above mentioned Tester, but it seems as though the Tester may have limited capabilities for sending large amounts of data. If a Blaster is purchased for the testbed, COLSA must devise a way to extract results from Blaster testing. This means that there must be a way of comparing the data "blasted" into the HIPPI interface to the data that ends up in the device (MasPar or HSR), to make sure they are the same. The Blaster is not a commercial item. However, IOSC has sold Blasters to companies that have expressed an interest. Raytheon may be purchasing a Blaster. The price of the Blaster is approximately $7,500.

3.1.4.5 HIPPI Interfaces

In this section each of the HIPPI Interfaces is analyzed. The analysis will include:

- The company producing the interface board
- Performance and maturity of the board
- A block diagram of the board where available

3.1.4.5.1 REXTTG-HIPPI. Designed by: Raytheon. **Maturity:** The board is a brand new product that has not been shipped. It is a Raytheon proprietary design.

A block diagram of the REXTTG-HIPPI interface card is shown in Figure 21. Details on the operation of this card have not yet been given. Through a memo to the project office COLSA has requested that Raytheon provide an explanation on how this board operates.

3.1.4.5.2 HIPPI Broadcast Switch. Designed by: IOSC. **Maturity:** HIPPI switches have been shipping for 1 year now. Both Cal-Tech and The Jet Propulsion Laboratory are using the switches. According to contacts at both of these institutions, the switches have worked "flawlessly." Further, contacts at Cal-Tech stated that IOSC produces the highest quality HIPPI products on the market.

3.1.4.5.3 HIPPI-MasPar Interface. Designed by: MasPar. **Maturity:** The Channel-based HIPPI board is a brand new product that has not been shipped.

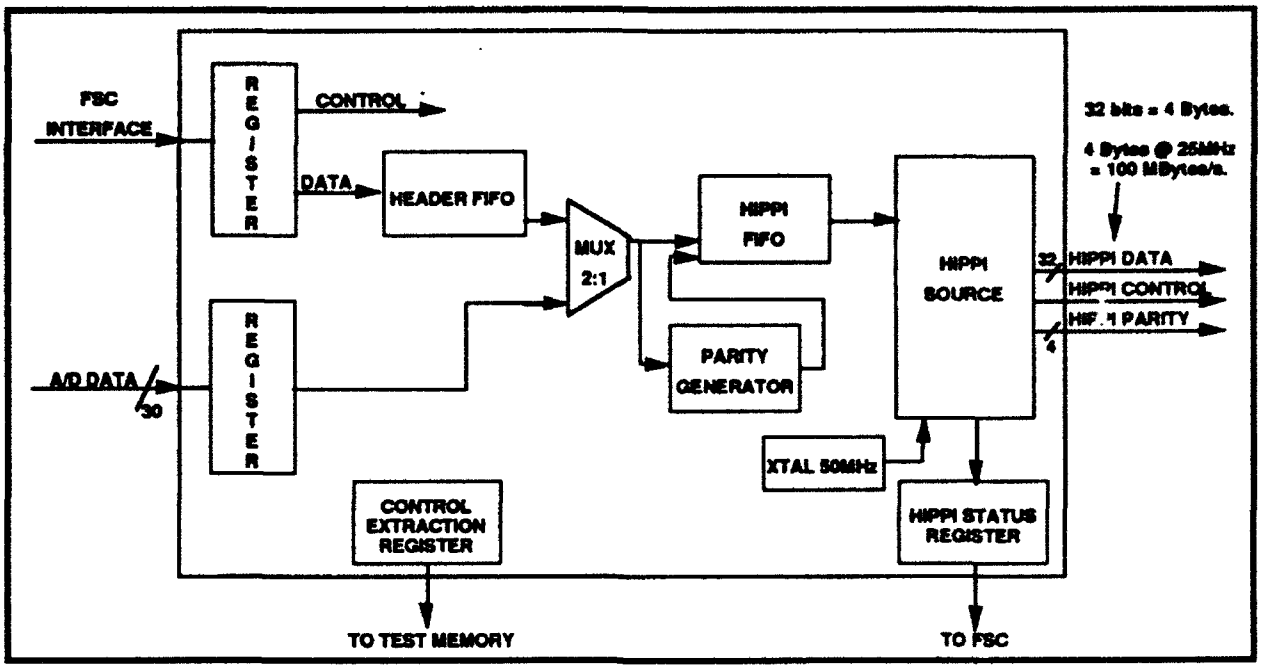

Figure 21. Block Diagram of Raytheon's REXTTG-HIPPI Interface

Before analyzing the operation of the board itself it is helpful to look at the board in relation to the rest of the architecture. Figure 22 shows a block diagram of the board integrated into the MasPar architecture. The MasPar will be a destination for both HIPPI data and HIPPI control and status information. The critical data path is shown with dashed lines, and the rates are specified for each path segment. The path is: HIPPI-board——>MasPar Input/Output (I/O) Channel——>I/O RAM——>Processor Elements; this path supports full HIPPI rates.

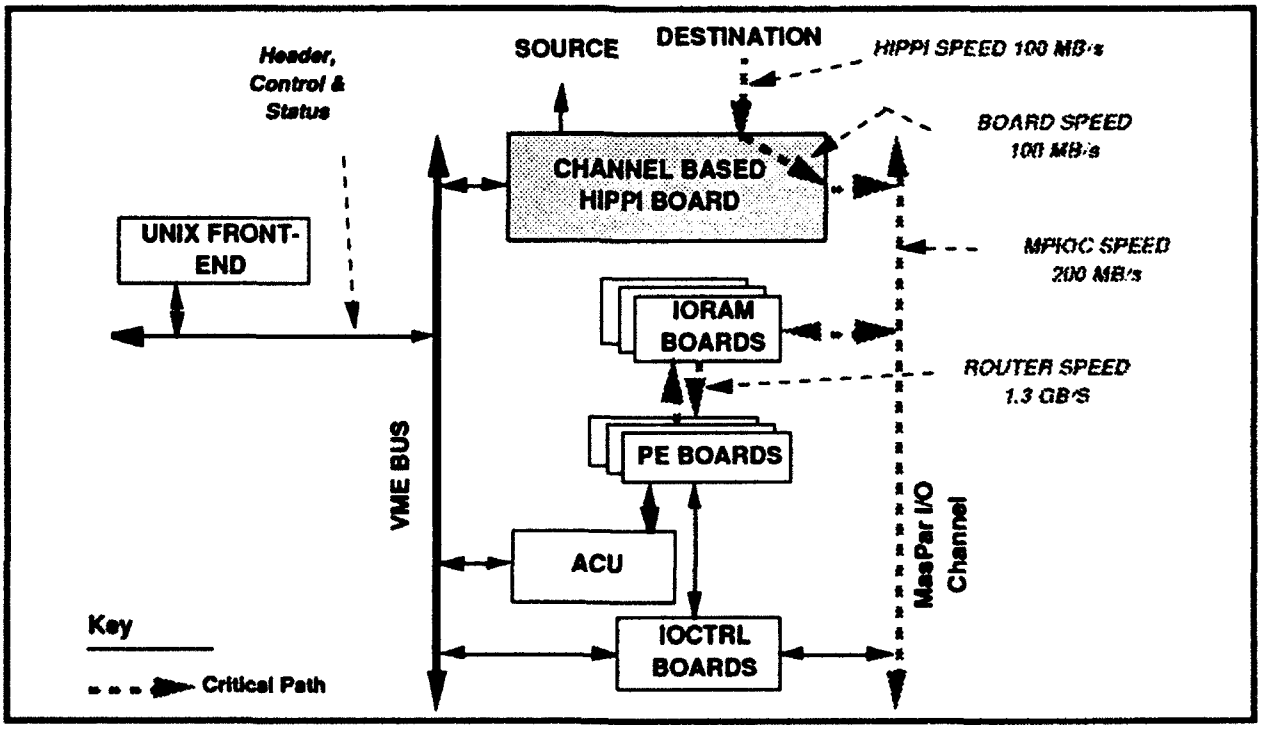

Figure 22. The HIPPI MasPar Interface

Control and status information follows a path to the front-end, through the VME bus. The VME bus has a rate of approximately 14 MB/s, which is sufficient for the control and status information since it consists of a small number of bytes.

The HIPPI Channel Board itself seems to be well designed. A block diagram of the board is shown in Figure 23.

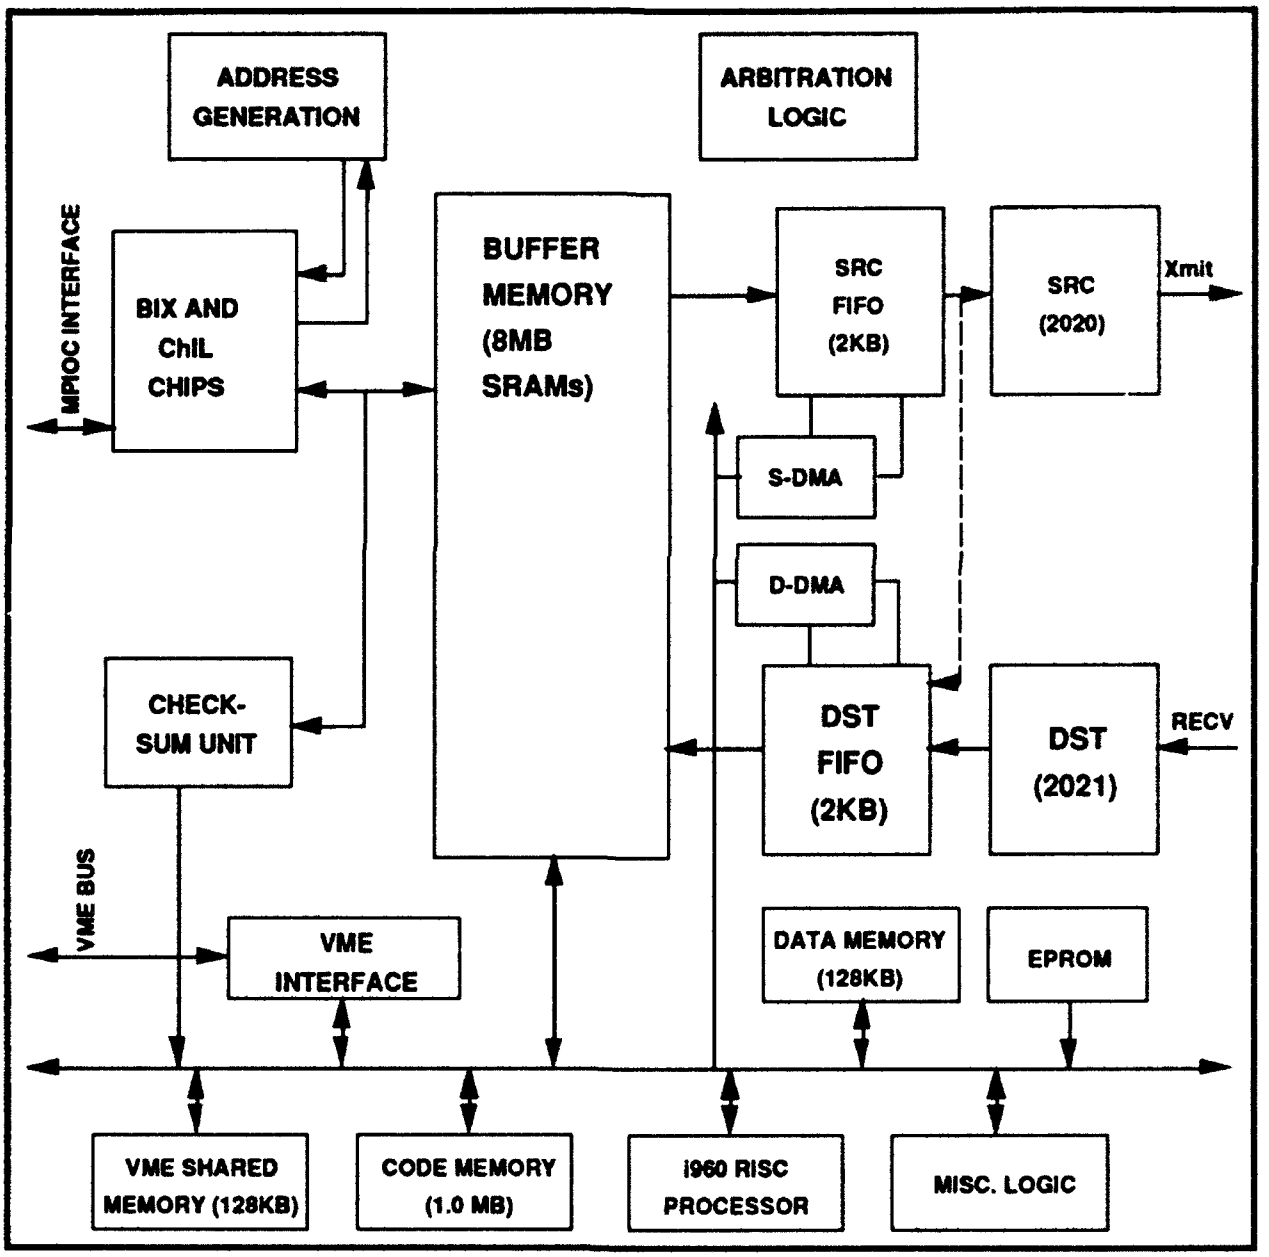

Figure 23. Block Diagram of the HIPPI-MasPar Interface

On the right side of the diagram one can see that the RECV (receive) channel is input into a "2021" chip. The 2021 chip implements a lot of the HIPPI-PH (lowest level protocol) in circuitry; this hardwiring of logic gives the 2021 high speed and low overhead, as compared to a software implementation. The 2021 chip is designed by Network Systems Corporation and manufactured by AMCC Corporation, the only known manufacturer of a HIPPI chip set. The chips are used by IOSC as well as MasPar. IOSC has had AMCC-based products in the field for a year without any problems.

Upon exiting the 2021 chip, the data passes through the DST FIFO buffer under the control of the DMA engine. From there the data reaches the buffer memory, and is operated on by the i960 processor. The processor examines the I-Field (information field) of the HIPPI packet to determine the destination address. If it does not match the destination device code the data is discarded. If it does match, processing continues with the "chopping-up" of the packet. The D1-Field of the packet, which contains control information as per HIPPI protocol, is separated and sent over the VME bus to the front-end. The D2-Field, which contains the data, is sent to the MasPar I/O channel via the BIX (Bus Interface Transfer RAMS) and CHiL (Channel Interface Logic) chips. From the I/O channel the data is transferred to I/O RAM and then to the processor elements via the global router. At this point, the processing elements may begin pulse compression calculations.

The BIX and CHiL chips are MasPar proprietary, but have been field proven in the MasPar MP-1 and MP-2 Computers.

3.1.4.5.4 HIPPI-Datatape Interface (HIPPI-VRB Adapter (HUA)). Designed by: IOSC. **Maturity:** The HVA is a brand new product that has not been shipped. It has been designed, built, and in-house tested via emulators. It is currently being integration tested with Datatape's LP400 high speed recorder.

Low level details of HVA operation were not available at this time. A block diagram was available, and is shown in Figure 24. The board can operate solely through HIPPI control commands which are fed from the HIPPI source to the control processor. Additional fine-tune control of the board is possible through the optional ethernet interface.

3.1.4.5.5 HIPPI-Sony Interface. Designed by: Triplex Corporation. **Maturity:** The Sony Interface is a brand new product that has not been shipped. Triplex claims that they have achieved HIPPI rates during in-house testing. They have promised delivery of their interface in the late second quarter.

Triplex feels that the HIPPI interface was rather simple, and that it was not their main concern with regard to the GBR program. The difficult part for them is the military hardening of the Sony recorder itself. The hardening capability or lack of, may be a factor in choosing a HSR vendor for GBR. COLSA has not received a specification sheet or block diagram from Triplex yet, but documentation should be received soon.

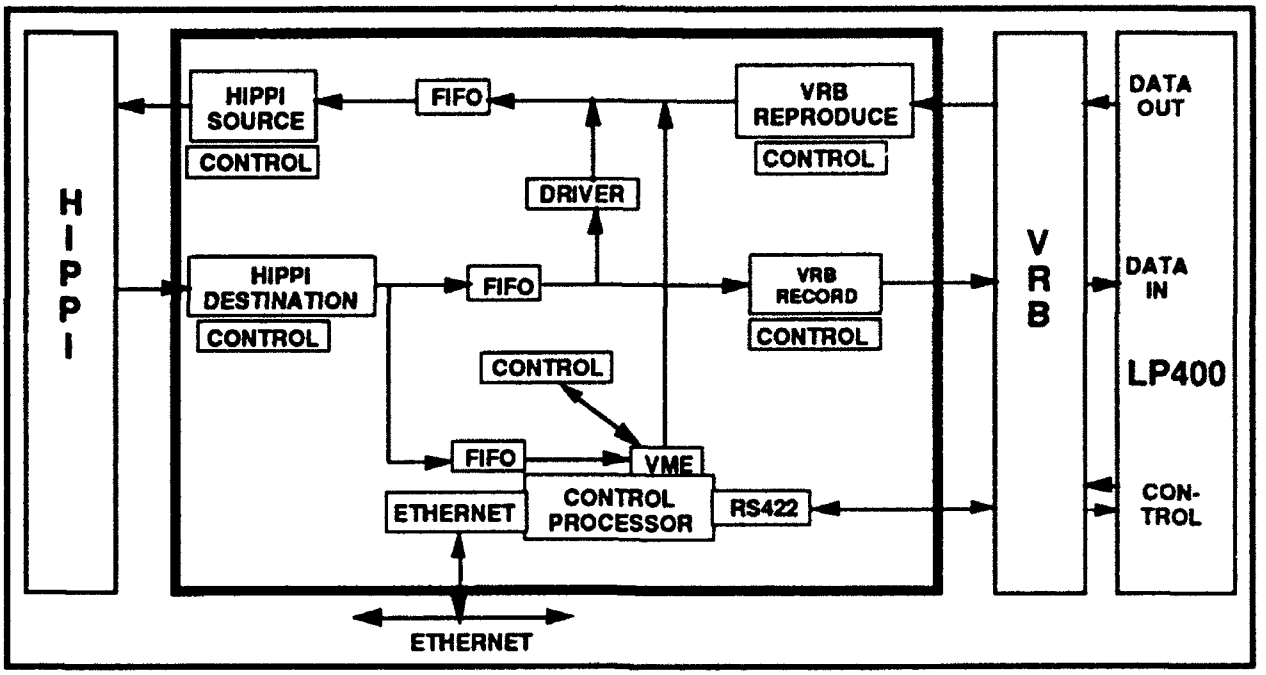

Figure 24. Block Diagram of the HIPPI-Datatape Interface

3.1.4.6 Interface Summary

Because the maturity of the interfaces has been the major part of this analysis, the results have bee summarized in Table 3 for quick reference.

TABLE 3. MATURITY OF HIPPI INTERFACES

INTERFACE	DESIGNER/ MANUFACTURER	MATURITY
REXTTG	RAYTHEON	HAS NOT BEEN PROVEN. RAYTHEON PROPRIETARY DESIGN.
BROADCAST SWITCH	IOSC	HAS BEEN PROVEN FOR ONE YEAR IN INDUSTRY.
MASPAR	MASPAR	HAS NOT BEEN PROVEN. MASPAR PROPRIETARY DESIGN.
DATATAPE	IOSC	HAS NOT BEEN PROVEN. NEW PRODUCT.
SONY	TRIPLEX	HAS NOT BEEN PROVEN. NEW PRODUCT.

Note that while IOSC's Datatape interface is a new product, IOSC has established themselves with other HIPPI peripherals that have been very successful in industry and academia alike.

3.1.4.7 Optimization Methods for HIPPI Communications

There are several techniques for implementing HIPPI communications that result in optimum data rates. These techniques may not be too important to GBR, since the REXTTG data rate will probably utilize only 25% of the specified data rate. However, if the channels do approach saturation from REXTTG data these techniques may become important. In the following, two basic techniques for fast HIPPI implementations are described.

First, HIPPI rates may be increased if the protocol is implemented in hardware as much as possible, rather than software. This is simply because hardware is faster than software. Two of the companies that are building interfaces, IOSC and MasPar, have stated that they are using AMCC HIPPI chips to implement lower levels of the protocol. The other companies have not indicated whether or not they are using the AMCC chips.

Second, rates may be increased if the connection is maintained over multiple packets. This can be accomplished by leaving the CONNECTION signal, one of the HIPPI interface lines, asserted between packets. If the CONNECTION is deasserted, then overhead is introduced in reestablishing the connection.

There are several other techniques that may be used to implement an efficient HIPPI implementation. If this becomes an issue in GBR such techniques can be studied and implemented to improve data rates.

3.1.4.8 HIPPI Conclusions

COLSA has completed a significant amount of "on paper" analysis of GBR HIPPI communications. It is COLSA's opinion that thorough HIPPI testing should be performed at the testbed to confirm reliability of HIPPI interfaces, and to verify throughput rates.

If a decision is made that HIPPI testing should be performed, then the purchase of the above mentioned HIPPI testing products should be seriously considered.

3.1.5 Technology Update On System Buses

Intel has defined a bus, the LocalBus, that connects device controllers directly to Central Processing Unit (CPU) control lines. The LocalBus will provide an intermediate I/O speed enhancement; ultimately "PC on a chip" technology will cut I/O distance requirements to a minimum, and that will greatly increase speeds.

Sun Microsystems has a high performance I/O bus called the SBus. The SBus has a transfer rate of 160 MB/s, four times the 40 MB/s transfer rate of the VME bus. SBus speeds may be valuable in GBR digital signal processing and FDDI areas due to its high transfer rates.

A new communications standard called the "Fiber Channel" (supported by ANSI) claims to have transfer rates of 1.0625 Gb/s, as compared to HIPPI's 800 Mb/s (32 bit mode). Fiber Channel will employ fiber-optic technology, and its protocol will be standardized by ANSI. It will not be available on the market until 1994, and it will be a superset of HIPPI.

3.2 Computer Analysis

During 1992 many off-the-shelf computers were evaluated for suitability in the GBR system. This section gives analyses of candidates for the Signal Processor, the Data Processor and the Workstations.

3.2.1 Signal Processor (SP) Analysis

During the proposal evaluation period, COLSA personnel researched the market in order to determine which manufacturers had computers suitable to GBR's signal processing requirements. The following analysis will help clarify this effort.

3.2.1.1 SP Candidate Comparisons

A signal processor comparison of three candidate manufacturers, MasPar, Cray and Control Data Corporation (CDC) was performed in order to determine which computers might be suitable for the GBR program. After the evaluation it became apparent that the CDC machine was a strong signal processor candidate. Cray's S-MP/APP was analyzed as much as possible and seems to be a strong candidate; however, COLSA feels that more in-depth analysis would be beneficial The MasPar MP-1 was borderline due to questions of system efficiency. The MasPar MP-2 system, introduced to the market after the above comparison, seems to be a much more capable computer. Table 4 describes some of the criteria that was used in the comparison.

TABLE 4. A COMPARISON OF SIGNAL PROCESSOR CANDIDATES

	MasPar MP-1	*Cray S-MP/APP*	*CDC APP*
Architecture	Parallel SIMD	Parallel MIMD	Vector
MFLOPS Per Processor	0.073	80	720
Processors	16,384	84	8
Peak MFLOPS	1200	6720	5760
16 K FFT Batch	10.4 msec.	5.6 msec.	5.7 msec.
Flexibility	One batch at a time	5 batches at a time	8 batches at a time
Study Performed	Yes(Too slow)	No(Need one)	Yes(Good choice)

3.2.1.2 MasPar MP-2

MasPar computers are now being sold by DEC as well as MasPar; DEC has bought MasPar computers, placed them inside of DEC cabinets, and now sells them as the DEC MPP 12000.

Maspar introduced its next generation computer, MP-2, which is highlighted in Figure 25.

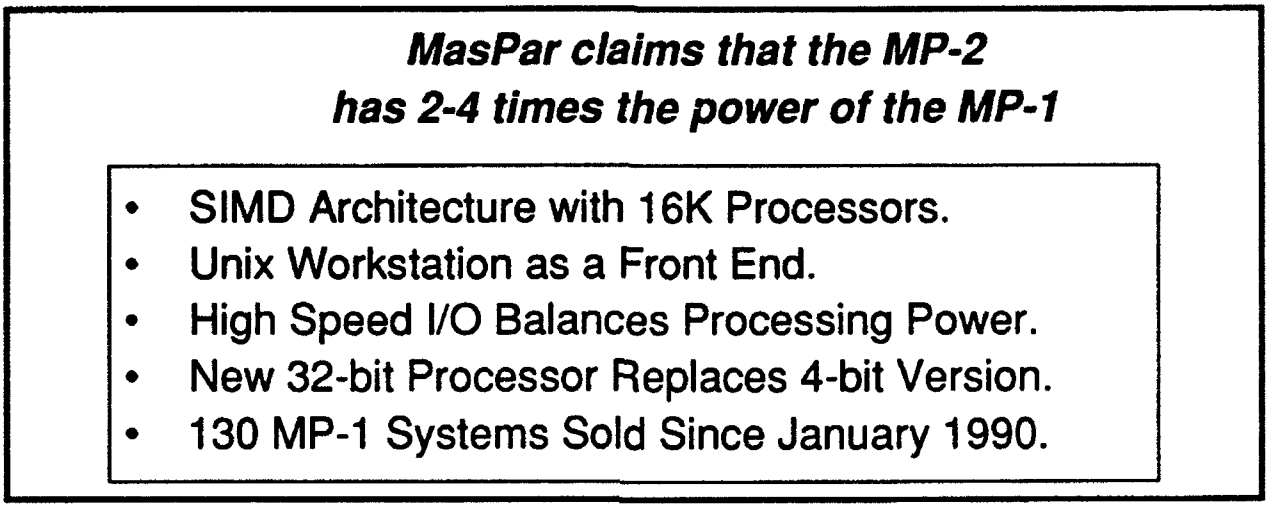

Figure 25. MasPar MP-2 Significantly Different From MP-1

3.2.1.3 *GBR Proposed Design*

Raytheon decided to use a string of MasPar MP-2 massively parallel processors, (three for TMD and five for GBR-T) as the GBR signal processor. COLSA established two meetings and a training session, with MasPar's Mr. John Dutton and Mr. Mark Holt, in order to understand the computer's architecture, capabilities, and software tool performance. After these meetings, it is COLSA's opinion that Raytheon's signal processor choice is of sound engineering judgement and should perform its radar tasks well.

3.2.1.4 MP-2 Footprint

The MP-2 has a relatively small footprint, as shown in Figure 26. Its dimensions are 58" H x 23" W x 32" D, and its weight is less than 800 pounds. It is an air-cooled machine and requires 220V/ 60Hz.

3.2.1.5 MP-2 Specifications

As shown in Figure 25, the MP-2 has tremendous processing power, two to four times that of the MP-1. It is a Single Instruction Multiple Data (SIMD) architecture. With the full configuration of 16K processors the MP-2 operates at 5000 MFLOPS. A new 32-bit processor architecture has been utilized with the MP-2 to provide the improvement over the MP-1.

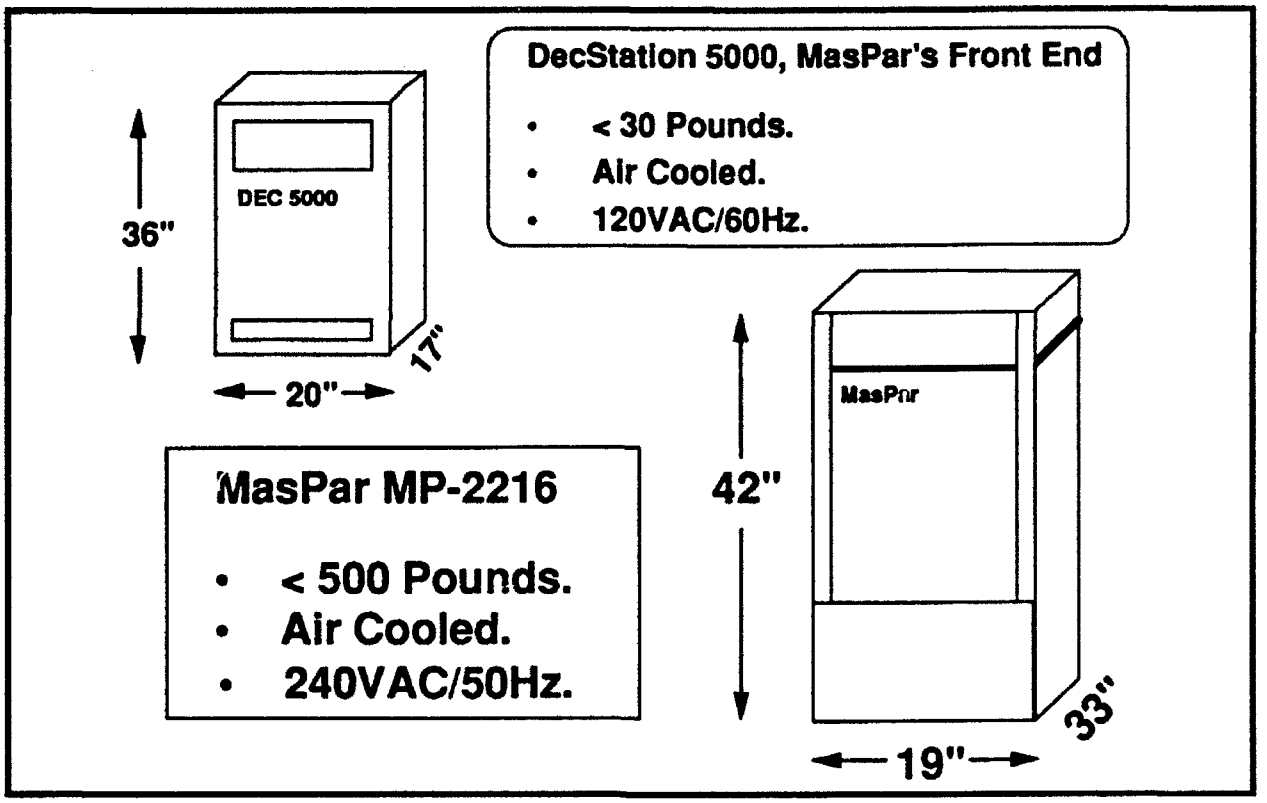

Figure 26. The MP-2's Footprint

3.2.1.6 MP-2 Architecture

Figure 27 is a graphic representation of MasPar's MP-2 architecture. The architecture can be broken down into four sections: the front-end, the ACU (Array Control Unit), the PE (Processor Elements) and the Communications System.

The Front-End (FE) is a DEC 5000/240 Unix Workstation. MasPar's intention

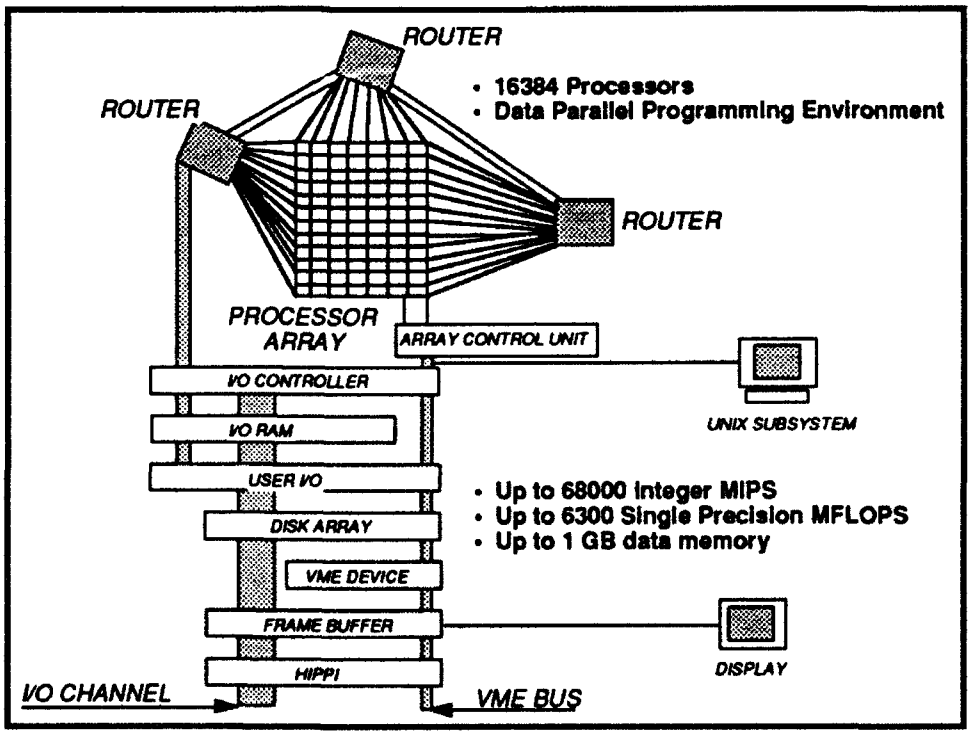

Figure 27. MasPar MP-2 Architecture

for the front-end was to make the massively parallel architecture of their machine as transparent to the programmer as possible. By successfully doing this, MasPar has mitigated one of the problems with high performance massively parallel machines, which are difficult to interface with and develop code for.

The ACU is the liason between the front-end and the PE. The ACU consists of a simple processor that can perform scalar computations, an instruction decoder and a system whereby it can communicate decoded instructions to the elements in the PE.

The simple MasPar proprietary processors that make up the parallel array, receive all of their instructions from the ACU. Upon receiving the instructions, they execute them in lock-step fashion and then compress the results and communicate them back to the ACU.

3.2.1.7 MP-2 Communications

The communications system utilizes three separate methods of interprocessor communictions, as shown in Figure 28. The first is simple broadcasting of data and instructions via the ACU. The second is through the Global Router which exchanges data between the processor elements and IORAM at a rate of 1.3 GB/s.

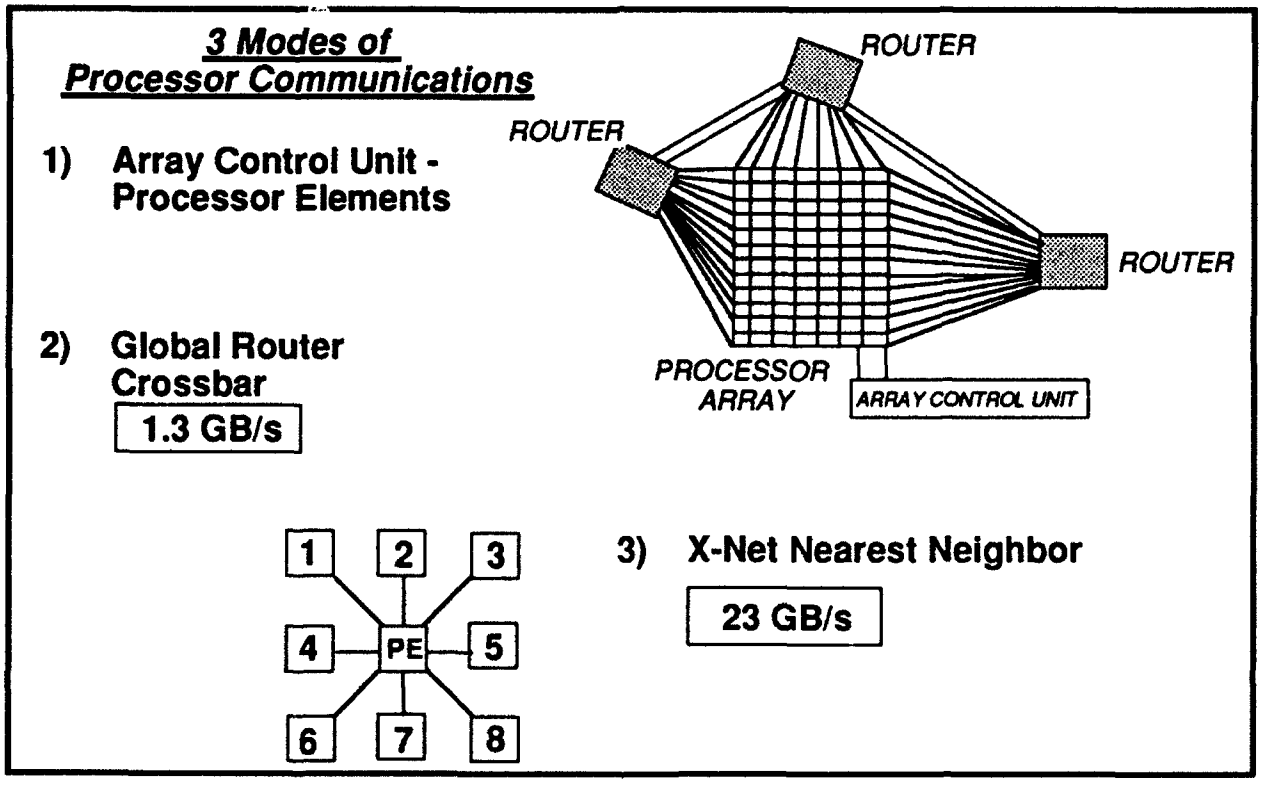

Figure 28. Three Modes of MasPar Communications

The final communications mode is the X-Net Nearest Neighbor system. X-Net can communicate data between neighbors (North, NorthEast, East, etc.) in the array at the rate of 23 GB/s. It is prudent to choose carefully between the X-Net and the Router depending on the application; and to avoid the use of ACU data broadcasting as much as possible. X-Net rates are inversely proportional to word size and distance between processors and, therefore, are effective for small word sizes over small distances. When communicating over greater distances or when using larger data words the Router may provide better throughput.

3.2.1.8 MP-2 Memory System

The MP-2 has up to 1 Gbyte of memory distributed among the processing elements and an additional 1 Gbyte is available in IORAM. Also, there is 1 Mbyte of physical instruction memory in the ACU.

Memory for the MP-2 exists in four areas: the front-end contains its own memory as any DecStation 5000 would, the ACU contains both instruction and data memory, the processor elements each contain 64 Kbytes of personal memory, and the IORAM boards contain 1 Gbyte of storage in the maximum configuration.

3.2.1.9 MasPar MP-2 Benchmarks

The benchmarks shown in the following figures (Figures 29a and 29b) are performance numbers, in response to a program office request pertaining to the capabilities of the new MasPar MP-2. According to MasPar, the MP-2 is still officially in beta testing, so these numbers are preliminary.

PRELIMINARY BENCHMARK RESULTS ON THE MP-2

BENCHMARKS ON 4,096-PE MACHINES

Benchmark: NAS (APPBT - block tridiagonal – 64-bit)

MP-1 Performance	-	75 MFLOPS
MP-2 Performance	-	225 MFLOPS
Speedup	-	3

Benchmark: NAS (APPSP - scalar pentadiagonal – 64-bit)

MP-1 Performance	-	58 MFLOPS
MP-2 Performance	-	153 MFLOPS
Speedup	-	2.64

Benchmark: NAS (APPLU - lower upper, SOR – 64-bit)

MP-1 Performance	-	36 MFLOPS
MP-2 Performance	-	115 MFLOPS
Speedup	-	3.1

Figure 29. MasPar MP-2 Preliminary Benchmark Results (a) Benchmarks for the 4,096 Processor Configuration

3.2.1.10 MasPar Hardware/ Software Relationship

Since a massively parallel computer has a different architecture from a uniprocessor system, it is beneficial to have programming languages that utilize the strengths of the architecture. MasPar provides a programming language called MPL (Massively Parallel Language - a version of ANSI 'C') for its MP series. MPL has extensions that are designed to take advantage of the massively parallel architecture and to make parallel coding more natural for the programmer. Some of these extensions were examined in order to understand how the massively parallel architecture is effectively utilized by software.

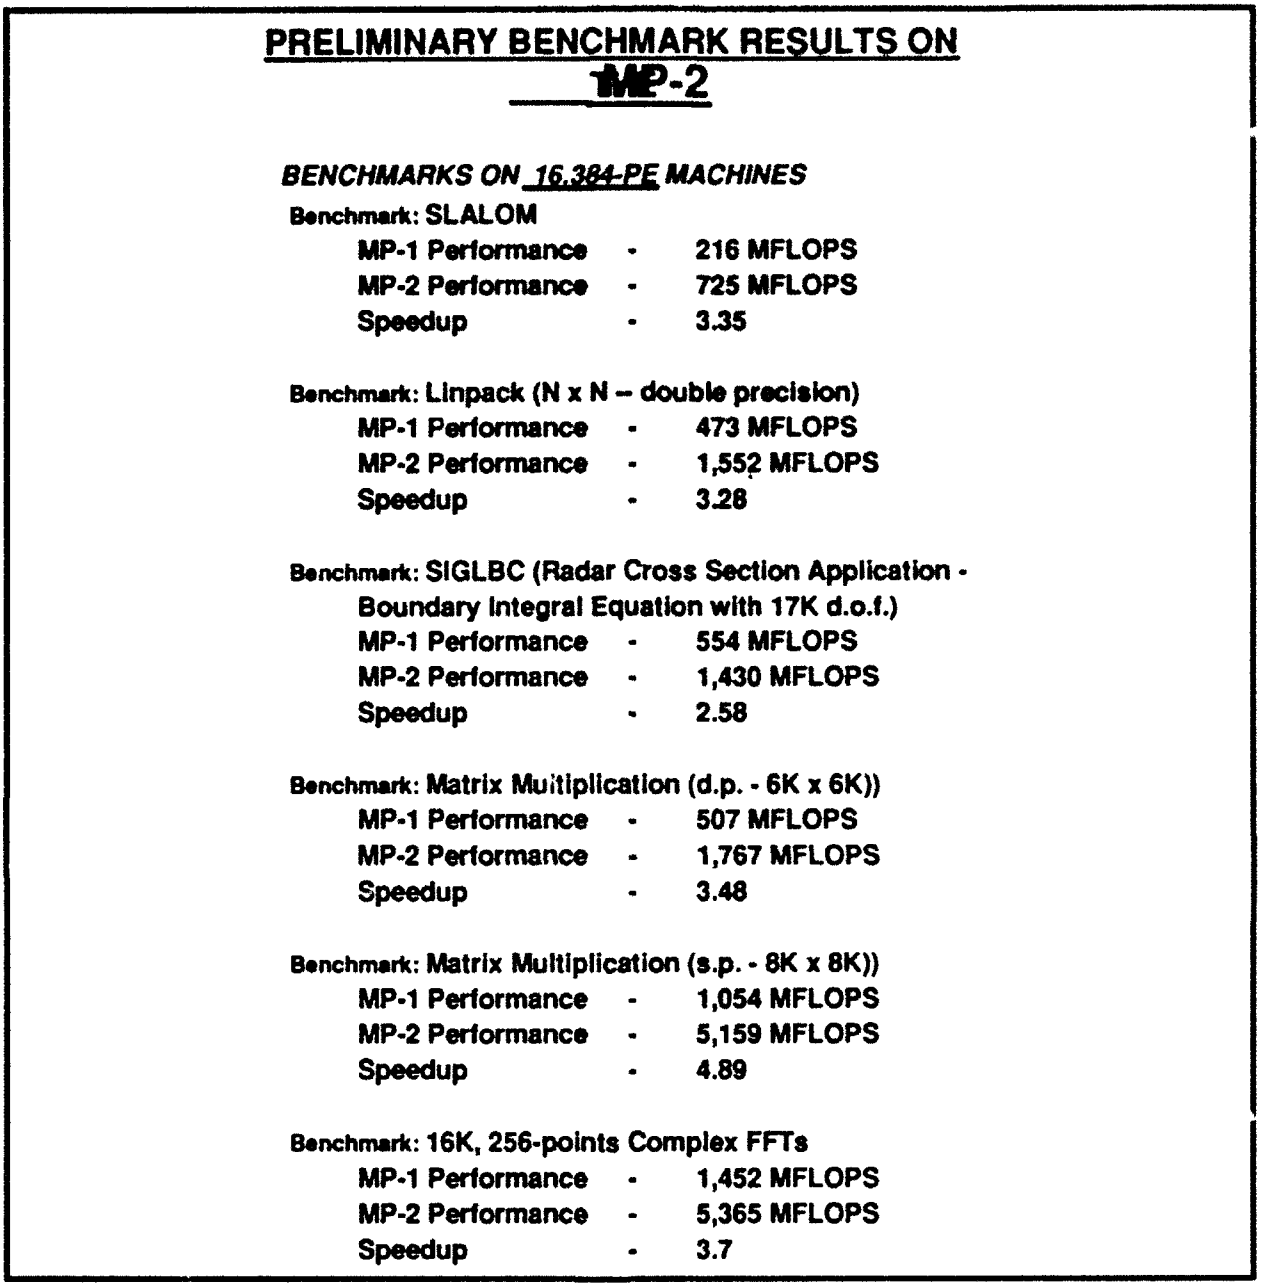

Figure 29. MasPar MP-2 Preliminary Benchmark Results
(b) Benchmarks for the 16,384 Configuration

3.2.1.11 MasPar Software Statements

PLURAL is used to declare a variable in each of the processors in the array. For example, PLURAL INT X, results in the variable X being declared identically in all processors. The following statements show the use of PLURAL in some simple arithmetic. First, variables X, Y, and Z are declared in each processor

PLURAL INT X, Y, Z;

Then, X and Y are added together in each processor giving a Z result in each processor.

Z = X + Y;

Thus, the programmer can command all processors in the array with very simple statements.

Control Statements such as IF, WHILE, and DO may be used to control which processors are active for a given operation. For example, to divide X by Y in each processor, while avoiding a division by zero error

IF (Y not equal to zero) then
Z = X / Y;
ELSE
do nothing;

The IF clause deactivates all processors in which Y = 0 for the divide operation.

The *SPREAD* extension is a very useful function for programming in the parallel environment and is efficiently implemented at the hardware level. When implemented, the function copies a data element across a row or column for a specified distance. The actual code for the function is given in Figure 30, part a; where DIM specifies the dimension (row or column) and M is an integer specifying how many rows (columns) to "spread" the datum across. The result of calling a SPREAD is shown graphically is Figure 30, part b.

The way the function is implemented in hardware is through an XNET (nearest neighbor) pipelined copy, as shown in Figure 30, part c. Since XNET communication lines are only 1 bit wide and because it takes one

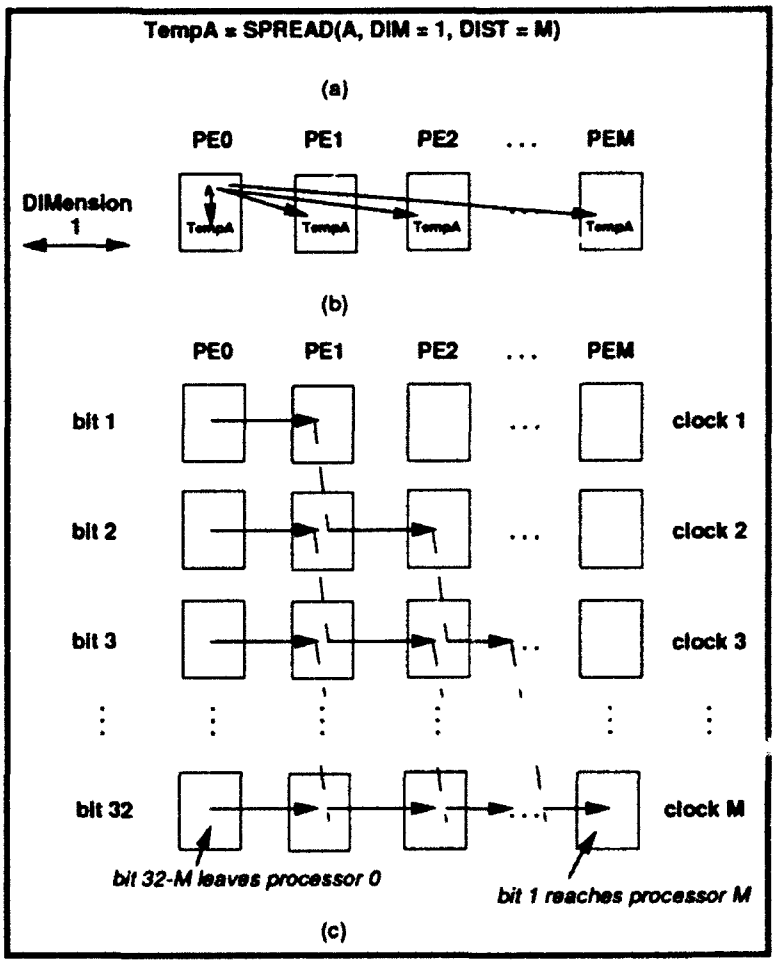

Figure 30. MasPar's Spread Function

clock cycle to transmit each bit, a non-pipelined approach would require 32 clocks to copy from processor 1 to processor 2, then 32 more cycles to copy from processor 2 to processor 3, etc. The result would be M x 32 clock cycles for the M-distance copy. With pipelining however, as each processor receives a bit, it retains a copy of the bit and transmits that bit to the next processor during the proceeding clock cycle. The result is that the copy takes only M + 32 clocks which saves a considerable amount of time. This process is analogous to instruction pipelining in a processor. This is an example of XNET dependency on operand size and distance ; note the "M + 32" signifies distance, M processors, and operand size, 32 bits.

The global router network is implemented using four communications primitives:

ROPEN - opens a connection between source and destination processor elements.
RSEND - moves data from destination processor element to source processor element.
RFETCH - moves data from source processor element to destination processor element.
RCLOSE - terminates connection.

Through the use of these primitives, a programmer can connect any two processors. Each processor can be involved in only one connection at a time; in the case of contention, connections are established on a "*first come-first serve*" *basis.*

Figure 31 depicts the architecture of the router for a single board (1024 processors). It is a 3-stage hierarchical connection that employs crossbar switches. As an illustration, let us examine the possible connections between processor #0, point A in the upper left, and processor #1023, point B in the lower right. One possibility is to route from S1#0—>S2#0—>S3#F, a second possibility is S1#0—> S2#1 —> S3#F, a third is S1#0—> S2#2 —> S3#F, etc. There are 16 possible paths. Since multiple paths exist, contention is decreased and the sustained router bandwidth is increased.

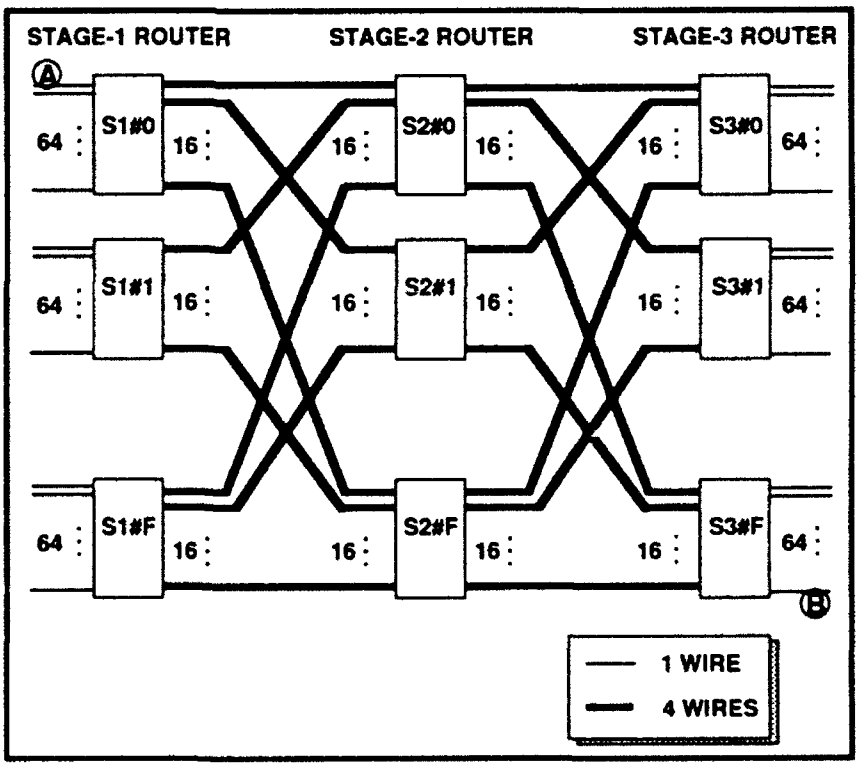

Figure 31. Multistage Interconnection Network

In a 16K MasPar MP-2 which will be the GBR signal processor, the same redundancy of paths exists. The stage two router of each board is connected with all others, thus providing for 16 paths between any 2 processors in the system. Further information on MasPar programming and hardware applications will be given in the future as they become available through system use by COLSA personnel and Raytheon alike.

3.2.2 Data Processor Analysis

As with the signal processor analysis, COLSA personnel compared as many pertinent computer systems as possible. IBM, CONVEX, Cray, and other manufacturers were evaluated. The systems most capable of meeting the GBR Data Processor requirements were the CDC computers, which have been discussed, and the VAX-7000.

In a number of periodic reports and quarterly presentations, COLSA personnel described in detail the capabilities of the VAX 7000 computers. Raytheon's choice of the VAX 7000 as the GBR data processor is reasonable, especially since it can be upgraded to include the high performance Alpha processor which was recently announced by Digital. The only problem that can be foreseen with going to an Alpha design, is the timely transition of software between the two different CPU designs. We believe that such hardware changes should be considered by Raytheon design engineers because it could provide a reasonable means for a significant upgrade for future GBR deployments. COLSA is continually trying to accumulate all the available information by meeting with Digital, in order to plan for the installation of a VAX 7000 within the GBR testbed located at the ARC.

3.2.2.1 DEC VAX 7000 Footprint

The VAX 7000 system is a medium-sized mainframe the size of a refrigerator (see Figure 32). It is an air cooled machine that weighs approximately 900 pounds when fully configured. The cabinet provides room for a tape drive, a battery backup unit, extra hard disks and high speed XMI buses.

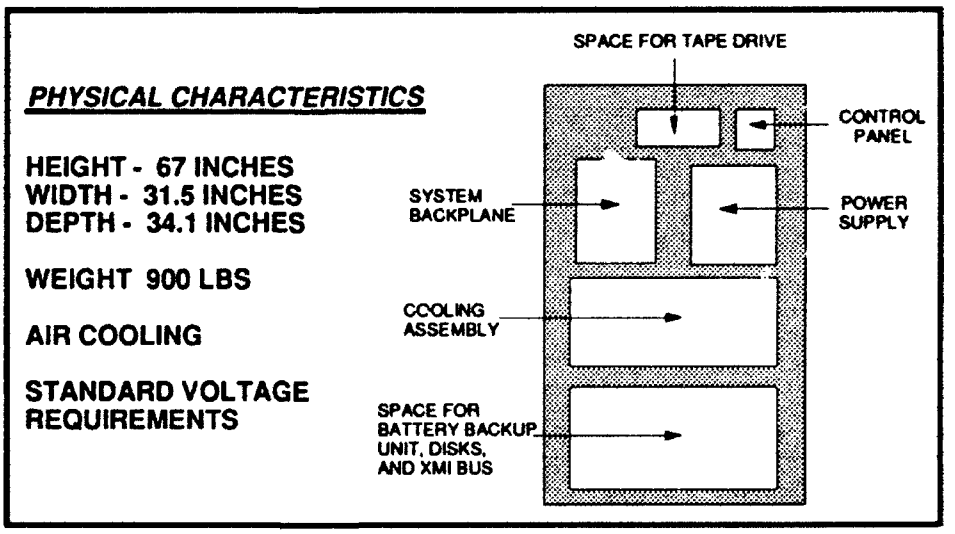

Figure 32. Physical Characteristics of the VAX 7000

3.2.2.2 DEC VAX 7000 Specifications

The VAX 7000 is a powerful computer that utilizes one to four VAX processors and that with special hardware is upgradeable to six processors. The VAX CPU is a 4-stage RISC pipeline that operates at 91 MHz. The system is board-upgradeable to the Alpha configuration for increased computing capability.

Strong computing performance is balanced with high I/O performance (see Figure 33). The system bus that connects memory and processors operates at 640 MB/s. A write-back cache makes processor-memory communications a fast process. For peripheral device use the 7000 provides up to 4 XMI 100 MB/s buses, providing an aggregate bandwidth of 400 MB/s.

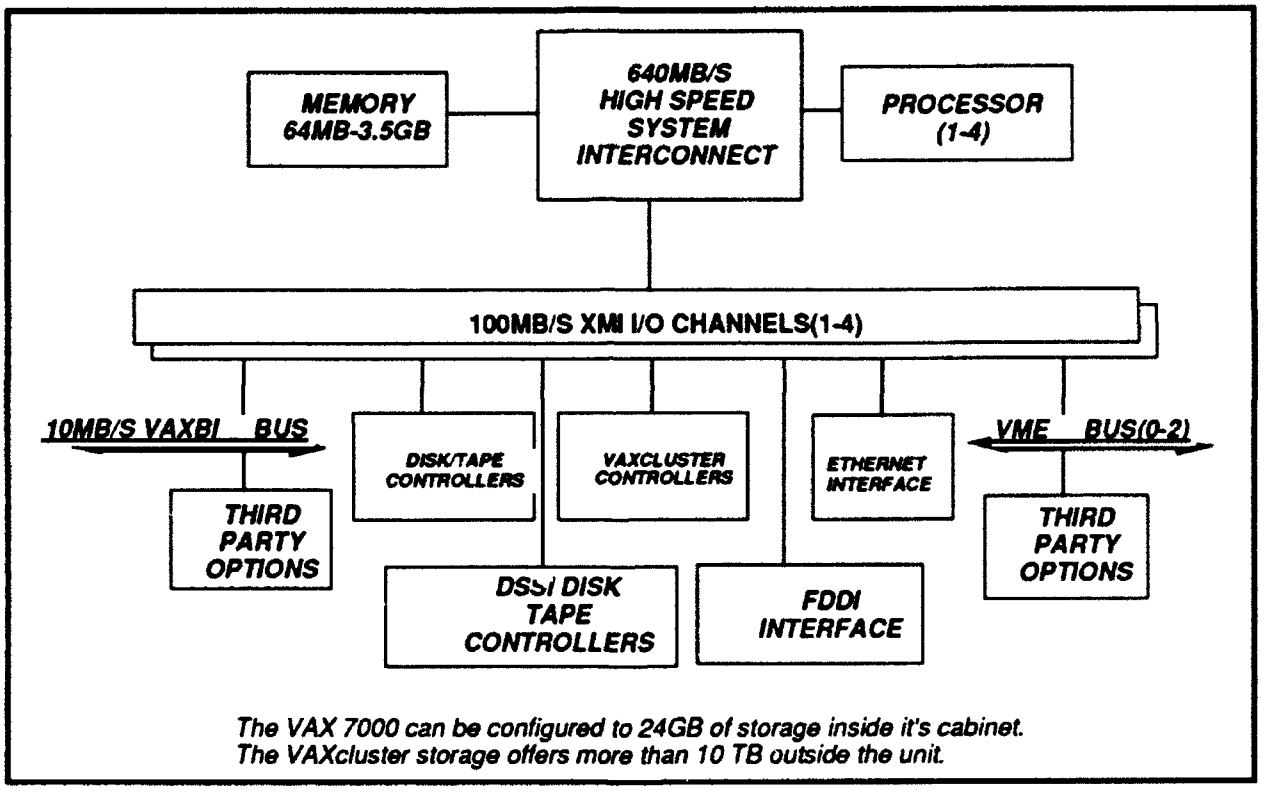

Figure 33. DEC VAX 7000 Architecture

Main memory capacity is presently at 512 Mbytes and will be increased to 3.5 Gbytes with the next release of the Open Virtual Operating System (VMS) operating system. Up to 24 Gbytes of secondary memory can be configured in-cabinet. For additional storage an I/O expansion cabinet is available, but at the cost of a significant footprint increase.

3.2.3 GBR Workstation Candidates

Several companies in the market today offer alternatives for use by the GBR program. The GBR console (TSEC) of the old Raytheon design was a Silicon Graphics (SG)/CDC configuration called the CYBER 910B-340 workstation. Since then, the relevant industry has produced many good new

products that should be considered by the program office. The SG IRIS Indigo, the IBM Powerstation 320, and the HP Apollo series 700 are all examples of reasonable performers in both number crunching and rendering, the two basic components of a good graphics solution.

Computing in the '90s is not going to be more of the same. We are moving from word and data processing to visual computing. By using pictures we can quickly comprehend complex information and communicate results, which is an important factor in the GBR design and functionality.

3.2.3.1 DECstation 5000/240

The DECstation 5000/240 will be used by the MasPar MP-2 computers as a front-end.

The workstations are an integral part of the MasPar system and have been working well with MasPars since their union. They are Alpha ready workstations that require only simple board level changes for the transition. Figure 34 describes the system's capabilities and performance values.

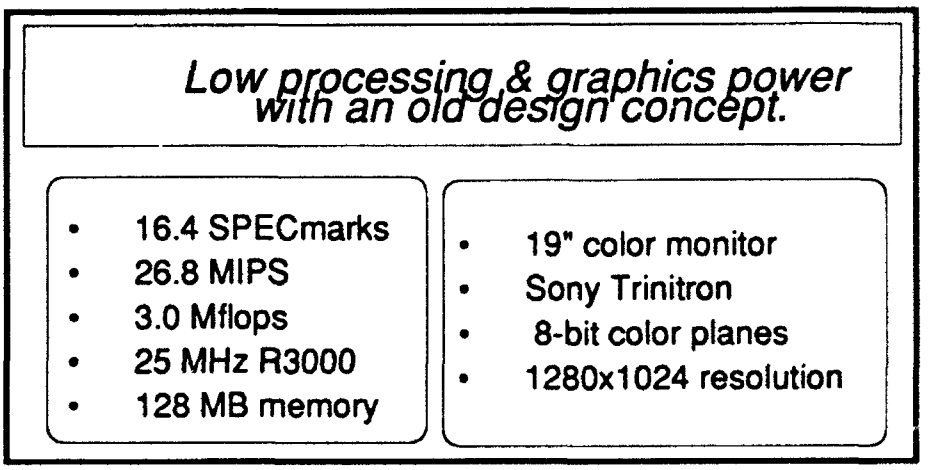

Figure 34. DECstation 5000 Specifications

3.2.3.2 DEC VAX 4000 Workstations

The VAX 4000 workstations are a high-performance yet cost-effective workstation series, and are upgradeable to the new Alpha processor configuration. The model 90 is the most powerful of the systems. It can deliver 32.8 SPECmarks, 128 MB of memory capacity, a disk space of 8.7 GB, and maximum I/O throughput of 5 MB/s in SCSI form. The model 90 also supports an optional TURBO channel adapter, providing connections to TURBO channel devices at speeds of 50MB/s. The workstation is completely compatible to the Open VMS system family. It offers accelerated two-dimensional (2D) graphics as well as three-dimensional (3D) graphics which are ideal for solids modeling. The VAX 4000 systems support many workstation applications, including those based on DECwindows/Motif, Uisx, OpenVMS and X-Windows applications. The system also has access to industry-standard DOS applications through DEC softPC for OpenVMS software.

3.2.3.3 HP-Apollo 9000-720CRX

The PA-RISC HP Apollo Series 700 delivers high price/performance in an expandable desktop workstation. Integrated graphics include leadership X11, 2D/3D vector performance, and outstanding

3D color modelling and rendering. The system utilizes a 50MHz PA-RISC processor which provides 57.9 MIPS integer, 17.9 MFLOPS floating point, and 59.5 SPECmarks. It has 128KBytes instruction cache, 256KBytes data cache, and up to 128MBytes of Experimental Control Center (ECC) Random Access Memory (RAM). The system utilizes a 2GB to 8GB 4mm DDS(DAT), and has up to 840MB internal disk storage capability. Its I/O connectivity includes SCSI-2, RS 232(2), Centronics, HP-HIL, and optional connections of FDDI, X.25, IEEE 802.5 token ring, HP-IB, fast differential SCSI, and Extended Instruction Set Architecture (EISA).

3.2.3.4 IBM Powerstation 320

The IBM RISC System/6000 7235 POWERgraphics GTO is an externally attached powerful graphics subsystem that can be connected via a single Micro Channel adapter slot to IBM's POWERstations 320, 320H, 520, 530, 540, and 550 systems. The GTO is the functional equivalent of the graphics subsystem of the IBM RISC powerstation 730. The GTO is available in two models, an 8-bit model 001 with dual frame buffers, area fill, and 256 colors, and a 24-bit model 002 with shading 24-bit Z-buffer, anti-aliasing, and 16.7 million direct colors. The POWERgraphics GTO provides excellent performance for PHIGS-based applications. It also supports IRIS GL-based applications.

Both GTO versions have a rated hardware performance capability for 2D and/or 3D vectors per second of 990,000. Additionally the 24-bit GTO has a rated hardware performance capability of 120,000 3D triangular strips per second. Peek hardware performance will be affected by the application and the IBM RISC System/6000 POWERstation processor attached.

Table 5 is a quick reference of processor capabilities for the IBM POWERstation POWERserver 220 and 300series.

TABLE 5. POWERSERVER SERIES 300 SPECIFICATIONS

Model No.	220	320	320H	340	350
Clock speed	33MHz	20MHz	25MHz	33MHz	42MHz
SPECmarks	25.9	32.8	43.4	56.6	71.4
Integer SPECmarks	17.5	15.9	21.8	28.8	36.2
Floating-Point SPECmarks	33.7	53.1	68.8	88.7	112.3
Data/Instruction cache	8kB	32kB/8kB	32kB/8kB	32kB/8kB	32kB/8kB
Bus width	64-bit	64-bit	64	64-bit	64-bit

3.2.3.5 Silicon Graphics IRIS Indigo

The Indigo uses the MIPS RISC processor, which provides quick, balanced performance: Integer 30 Million Instructions Per Second (MIPS), floating point 4.2 MFLOPS, total system performance 26 SPECmarks (see Figure 35). The system also delivers high speed 2D color vectors, fast pixel operations and interactive 3D graphics. Users are free to buy just as much graphics power as they need, and can protect software investments even as upgrades are made. The system also utilizes video connections for a VCR as well as printer and audio capability.

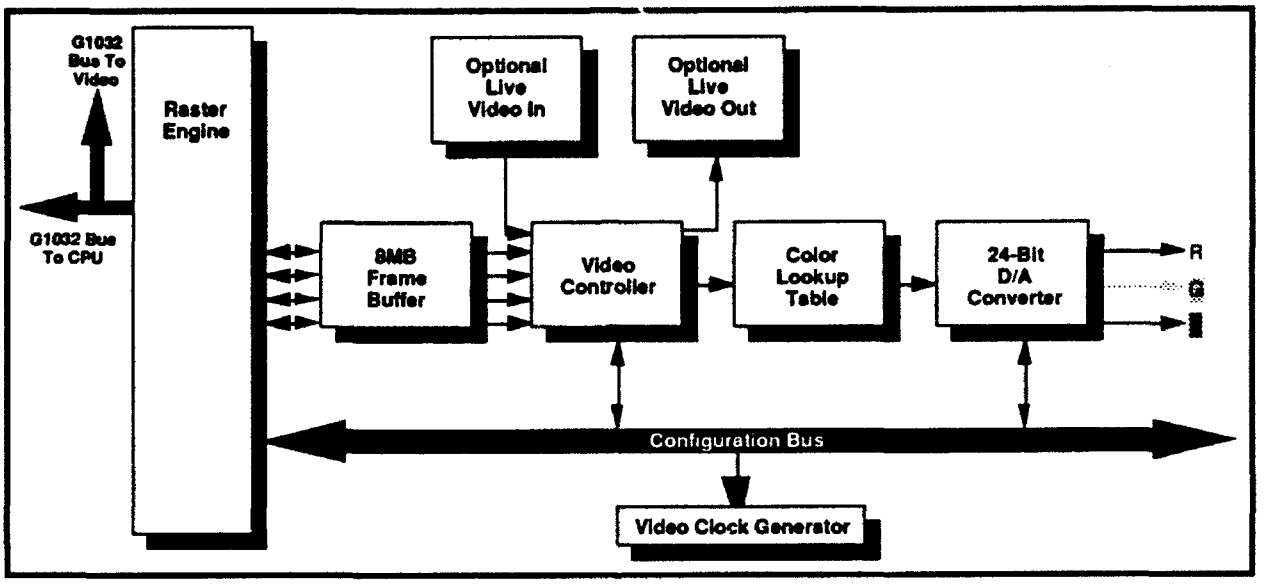

Figure 35. Silicon Graphics Indigo Dedicated Graphics Board

3.2.3.5.1 SG Indigo Graphics Board. The IRIS Indigo performs graphics computations at a speed that equals or surpasses many machines with special-purpose graphics processors. The CPU board contains four functional sections: the processor core, which contains the CPU and the FPU; the main memory which contains Dynamic Random-Access Memory (DRAM) and supporting circuitry; the I/O system which contains peripheral ports and hardware to read incoming data, and to manage incoming and outgoing data; and the audio system which contains audio ports and digital signal processing hardware.

3.2.3.5.2 SG Indigo CPU Board. The R3000A CPU, Figure 36, which is used by the system is the latest generation RISC processor from MIPS. It is assisted by 32 KB data and instruction caches, and a MIPS R3000A floating point unit. Main memory capacity is 96 MB of DRAM, augmented by on board interleave controllers that greatly increase the efficiency of memory access. The IRIS Indigo uses a CPU bus clocked at 33MHz, and a data-transfer rate of up to 133 MB/s. To take advantage of improving technology, the system clocks the CPU and system buses independently. This not only allows both the system bus and the CPU to run at maximum performance, but also allows upgrades to the CPU and graphics to occur independently. The IRIS Indigo graphics architecture uses general-purpose processing and special-purpose circuitry. It runs a version of SG IRIS Graphics Library that uses Indigo's CPU and FPU to perform most of the graphics calculations previously performed by

proprietary graphics processors, and can execute graphics calculations at speeds comparable to those of other IRIS systems with specialized graphics hardware. Future upgrades to a faster CPU will give a corresponding increase in graphics speed without the need to replace existing graphics circuitry. A high quality display provides a framebuffer that supports 1024 by 768 pixels and includes 8 bitplanes for screen images, 2 bitplane for overlays, and 2 window-clipping bitplanes.

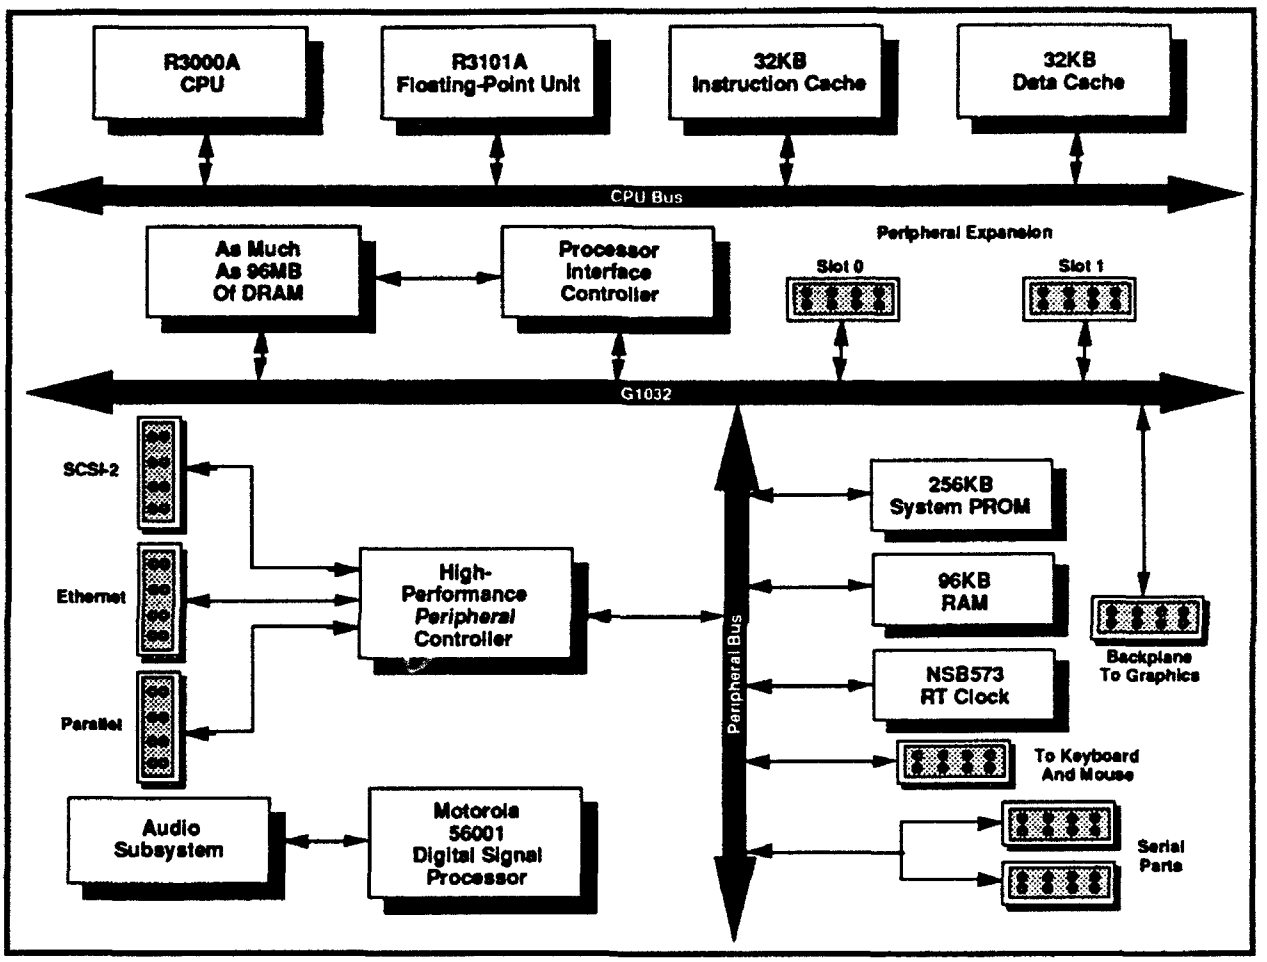

Figure 36. Silicon Graphics Indigo CPU Board

The system also utilizes a 3 button ergonomic design mechanical mouse that requires no mouse pad. Its I/O capabilities include a 4MB/s SCSI-II interface, a high-throughput Ethernet interface, a bi-directional 8-bit parallel port implementation implementing a superset of the Centronics specifications, two serial ports with transfer rates of 38.4 kBaud, and two serial ports dedicated to keyboard and mouse.

3.2.3.5.3 SG MIPS R4000. SG is the first manufacturer to use the R4000 CPU in their computers, clearly because of the merger between SGI and MIPS. The R4000 is significantly different from the older R3000. It is a full 64-bit processor and uses superpipelining and a twice as fast internal clock. On the Crimson, the CPU appears to the bus as a 50MHz chip, but internally, it operates at 100MHz. Superpipelining allows faster CPU processing while maintaining a clock speed that memory and

other devices can live with. Because the R4000 will run the 32-bit R3000 binaries without modification, a performance gain is achieved even without recompiling. The R4000 has three major logical units integrated on the same chip: the integer execution unit, the FPU, and the Memory Management Unit (MMU). There are 8KB of instruction cache memory and 8KB of data cache memory integrated on the R4000. It is SG intent to incorporate this CPU to most of its computers. Since the R3000 code can run on the R4000 without modification and with significant improvement in performance, this is an easy upgrade. The real performance gains for programs that use large data sets will be realized as the applications and operating system use the full 64-bit characteristics of the R4000. The excellent engineering of the IRIS CRIMSON and the IRIS Indigo are responsible for SGI's being recognized for more than 3-D graphics workstations and are machines for very serious work. The company is now carving out a piece of the market for general engineering and office workstations. The merger with MIPS has ensured SGI's foothold.

3.2.3.5.4 Silicon Graphics 3-D Graphics Engines. Two new 3-D graphics accelerators from SG and Matrox Electronics use hardware and software muscle to make personal computers perform almost like high-end graphics workstations. For now, this means users like GBR can spend less to display animations and speed display-lost processing and zoom/pan redraw functions. In the future when Windows 3.0 and other drivers become available, these boards could be integral to the work of graphics artists, designers etc. To achieve their performance, Matrox's MG-3D series and SG's IrisVision boards take over graphics-intensive algorithms from the main system. Both adapters have massive amounts of on-board video RAM for display-list storage. The hardware on these adapters are more commonly found on UNIX-based Scalable Processor Architecture (SPARC) and Iris workstations. The Matrox boards use the Texas Instruments DSP-30 chip. The DSP-30 accelerates the drawing of primitive functions to around 20,000 Gouraud-shaded polygons per second. This is a measurement based on adding color and light shading attributes to a wireframe model. IrisVision's geometry uses triple 8-bit D/A converters, a TMS-340C30 processor, and a Weitek floating-point coprocessor to speed graphics operations. SG rates the adapter at 14,000 Gouraud-shaded polygons per second. Both boards support standard AT-bus platforms, but each vendor has staked out different territory when it comes to bus architectures. IrisVision comes in ISA and Micro Channel versions. For ISA-based machines the IrisVision/8 ($3495) displays 256 colors and the IrisVision/24 ($4995) displays 16.7 million colors (24-bit). The Micro Channel version costs 43995 and supports 256-color output. Matrox's MG-3D series includes ISA and EISA boards: The ISA-based MG-3D which offers 12-bit dithered color images, costs $3495. The EISA-based MG-3D Ultra sells for $5995. All the products require 386 or 486 systems. IrisVision requires a math coprocessor. The systems support AutoCAD 386 release 1.1 including the Micro Channel adapters.

3.2.3.6 3-D Graphics Software Applications

For users purchasing 3-D graphics applications from software vendors, or developing their own, the choice of graphics standard will directly influence the performance of systems, networks and applications, as well as purchasing and development cost. The traditional 3-D graphics choice between X windows and Programmer's Hierarchical Interactive Graphics Standard (PHIGS),

changed with the announcement from SG of an open systems version of its Iris GL graphics library. The OpenGL application programming interface which will run on X windows, is being billed by the company as an emerging multiplatform industry standard for high-performance graphics. The differences between OpenGL and PHIGS are substantive. In general, PEX-based graphics are able to handle many of the mainstream 3-D applications quite easily. OpenGL on the other hand, is well suited to the handling of those leading-edge applications designed to make the most of photorealistic rendering or video applications. GL is also well equipped for immediate mode graphics, where a command from the keyboard or mouse immediately becomes a change on the screen. X windows requires extensions to handle 3-D graphics, because the protocol supports only primitives consisting of x, y, and integer coordinates, rather than floating point coordinates, which provide the detailed image description needed to generate the more complex 3-D graphics without unduly loading the network with pixel-by-pixel description of the image. With an extension such as PEX or OpenGL, the X client can create 3-D primitives, which are sent to the server. The server translates these primitives into 3-D displays. But the method by which the two systems support this process are substantially different.

3.2.3.6.1 The Role of PEX. PEX simply extends the X protocol so that it can support the functions of a higher-level Application Programming Interface (API) such as PHIGS. The PHIGS API provides graphics subroutine libraries for the creation of graphics within an application. In the case of PEX and PEX library (PEXlib) , it is PHIGS that provides the actual graphics library and subroutines that an application calls. The role of PEXlib is similar to that of Xlib. PEXlib, a low-level API, is triggered by subroutine calls from a higher level API or application. PEXlib produces the appropriate PEX protocol packets and sends them to the appropriate graphics server. PEXlib simplifies application development by using a higher level API such as PHIGS, which provides the routines necessary to build PEX protocol packets and send them across the network.

3.2.3.6.2 Opengl Functions. Open GL, both extends the X protocol to support 3-D primitives and provides a high-level API with graphics library and subroutine calls. Overall, the OpenGL API has more functions available to it than does PHIGS, the API for which PEX was designed. For example, OpenGL supports textured shading, which allows a rough surface to be displayed; PHIGS has no comparable function. Similarly, OpenGL has a number of functions required for video-type display, which could be used in the preparation of special effects for videos.

3.2.3.6.3 Immediate vs. Display. There are two broad approaches to rendering graphics: immediate mode and display list mode. In immediate mode, the application specifies a primitive using a series of commands, and that primitive is immediately rendered on the display. Display list mode however consists of a list of commands whose execution has been deferred. This list of commands constitutes a structure which may be stored in a graphics library. These two types of operations are best suited for different types of applications. Display list mode is suited for applications in which the overall geometry of a drawing may change while its elements do not. Immediate mode rendering is more efficient in situations where the geometry of the whole, as well as the shape of the individual elements is changing.

OpenGL and PEX support both rendering modes. However OpenGL is usually seen as the better performer in immediate mode, while PEX with its Phigs history is regarded as the superior performer with display lists.

3.2.3.6.4 OpenGL vs. PEX. From a simple features count, OpenGL has the advantage over PEX. With capabilities such as anti-aliasing, texture mapping and the ability to blend displayed objects into one another, OpenGL is clearly closer to the leading edge of technology for certain applications. PHIGS/PEX is concentrated on the set of features that were commonly available when the PHIGS graphics standard walked onto the set a few years ago. However this set is sufficient for most commercial CAD/CAM, GIS and similar applications.

OpenGL is a new standard, large parts of which remain untested. According to SG large amounts of IRIS GL had to be rewritten to create OpenGL. PEX on the other hand has been out for over a year since Sun released its PEX sample implementation. And PEXlib, although a recent entry, is largely composed of PEXlib versions placed into the public domain by both DEC and Sony. From a software developers point of view openGL may be the easier system to work with since applications written with GL flow easily.

3.2.4 Microprocessor Evaluations

In the process of evaluating computer systems it became apparent that microprocessor evaluations would be necessary in order to determine which architecture would be most profitable to the program. COLSA presented a number of microprocessors during quarterly reviews and monthly reports. The following is a brief description of different chips available on the market. Table 6 is a quick reference for chip specifications as provided by their corresponding manufacturers.

TABLE 6. MICROPROCESSOR COMPARISON

	SPEC MARKS	CLOCK SPEED	WORD SIZE	ARRIVAL	PRICE
DEC ALPHA	150	200MHz	64 BIT	4Q 92	$1650
MIPS R4000	70	100MHz	64 BIT	1Q 92	$1000
TI SUPER SPARC	75	100MHz	64 BIT	N/A	N/A
MOTOROLA 88110	63.7	50MHz	64 BIT	2Q 92	N/A
INTEL I860 XP	42	50MHz	64 BIT	1Q 92	N/A
IBM RS 6000	25.9	33MHz	64 BIT	N/A	N/A
HP 7100	120	100MHz	64 BIT	4Q 92	N/A
SUN SUPERSPARC	75	50MHz	32 BIT	4Q 92	N/A

Note: N/A = Information Not Available

3.2.4.1 DEC Alpha Chip

The DEC Alpha chip's distinguishing feature is its tremendous speed. At 200MHz, 150 SPECMARKS, and 300 MIPS it truly is an impressive processor. It has both instruction and data cache on board as shown in Figure 37, with 64-bit integer and floating point execution units, and 32-bit integer and

floating point register files. Its closest competitor is supposed to be the HP-7100 at 120 SPECMARKS which is due for release at the end of 1992.

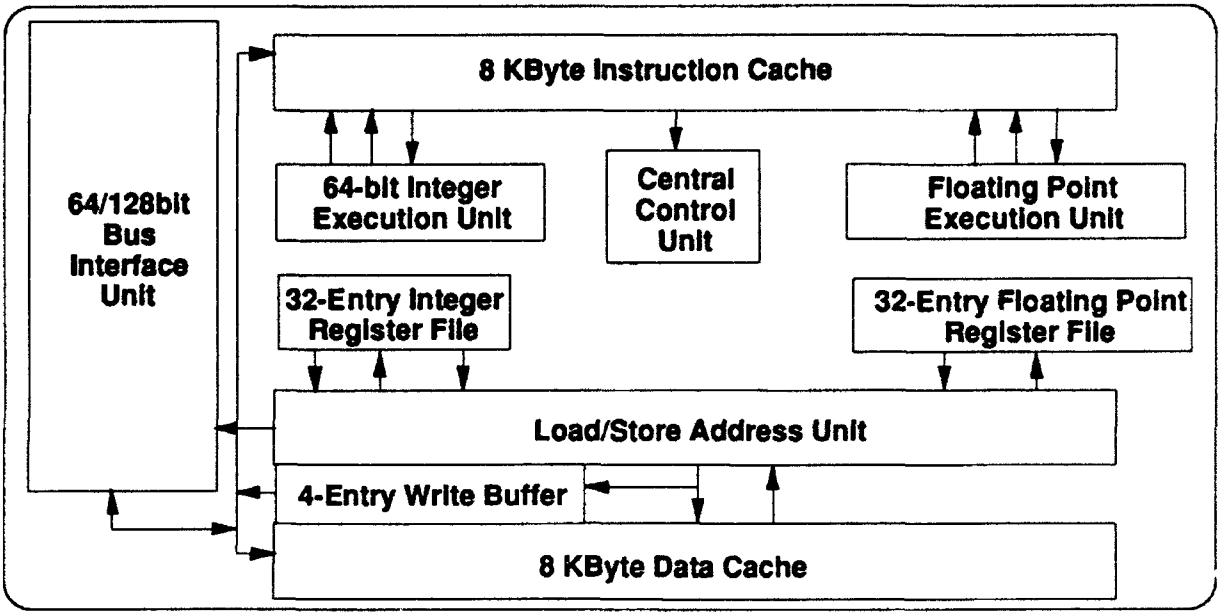

Figure 37. DEC's Alpha Chip

3.2.4.2 INTEL 486DX2

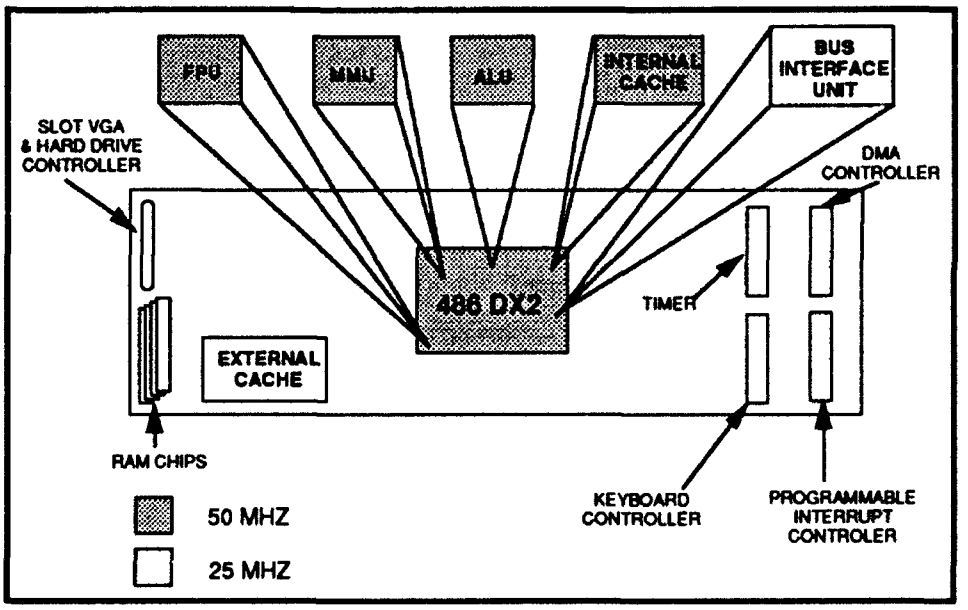

Figure 38. Intel's 486DX2 Chip and a Sample Motherboard

Intel has introduced the 486DX2, a chip that is quickly and easily integrated onto a motherboard. In the past, integration has been slow and difficult since designers had to make the remainder of the motherboard (see lightly shaded regions of Figure 38), fast enough to keep up with the microprocessor (dark shaded regions of Figure 38), a difficult task. The 486DX2 does run at high clock speeds internally (50-66 MHz), but at only half those speeds externally (25-33 MHz). With freedom to work at lower speeds, motherboard designers can produce market ready boards faster, more easily, and with more functional reliability.

3.2.5 Storage Device Forecast

The following tables (7, 8 and 9) show some of the capabilities of storage devices as of January 1991, with predicted capabilities for the end of 1992. Included is a comparison of 3.5" magnetic storage disks versus 3.5" optical storage disks. Industry experts predict that a 5.25" optical disk may hold 40 GB by the end of the 1990s.

TABLE 7. HARD DRIVE FORECAST

	JAN '92	DEC '92
MAXIMUM CAPACITY FOR 3.5" DISKS	550 MB	0.8 - 1.2 GB
PHYSICAL SIZE	3.5" MOST COMMON	2.5" MORE COMMON, 1.8" INTRODUCED

TABLE 8. FLOPPY DRIVE FORECAST

	JAN '92	DEC '92
CAPACITY FOR 3.5 " FLOPPIES	1.44 MB	4 MB or a leapfrog to 20 MB

TABLE 9. MAGNETIC VERSUS OPTICAL STORAGE DEVICES

	MAGNETIC	OPTICAL
COST PER DRIVE	$ 75	$ 800
COST OF MEDIUM PER MEGABYTE	$ 1.25	$ 0.55
CAPACITY PER 3.5" DISK	1.44 MB	128 MB
ACCESS TIME	16-20 ms	45-80 ms

The 1325VM Floptical Drive is a new product. It is a type of floppy drive that is a hybrid floppy-hard-optical disk. It is categorized as a very high density (VHD) disk because it can store 20 MB rather than the 1.44 MB that floppies usually store. The drive takes 3.5" magnetic disks very similar to what we now use. The floppy has two advantages over a magnetic tape drive; it is mechanically faster, and it allows direct rather than sequential access.

3.2.6 High Speed Data Recorders

A part of the requirements for the GBR system is to record data that is received during missions. The probable configuration of HSRs within the GBR loop is shown in Figure 39. HSR(1) will record digitized radar signature data which consists of the Sum, Alpha and Beta channels (i and q components for each). HSR(2) will record event data which will consist of several items such as waveform data, tracking data, and operator action data.

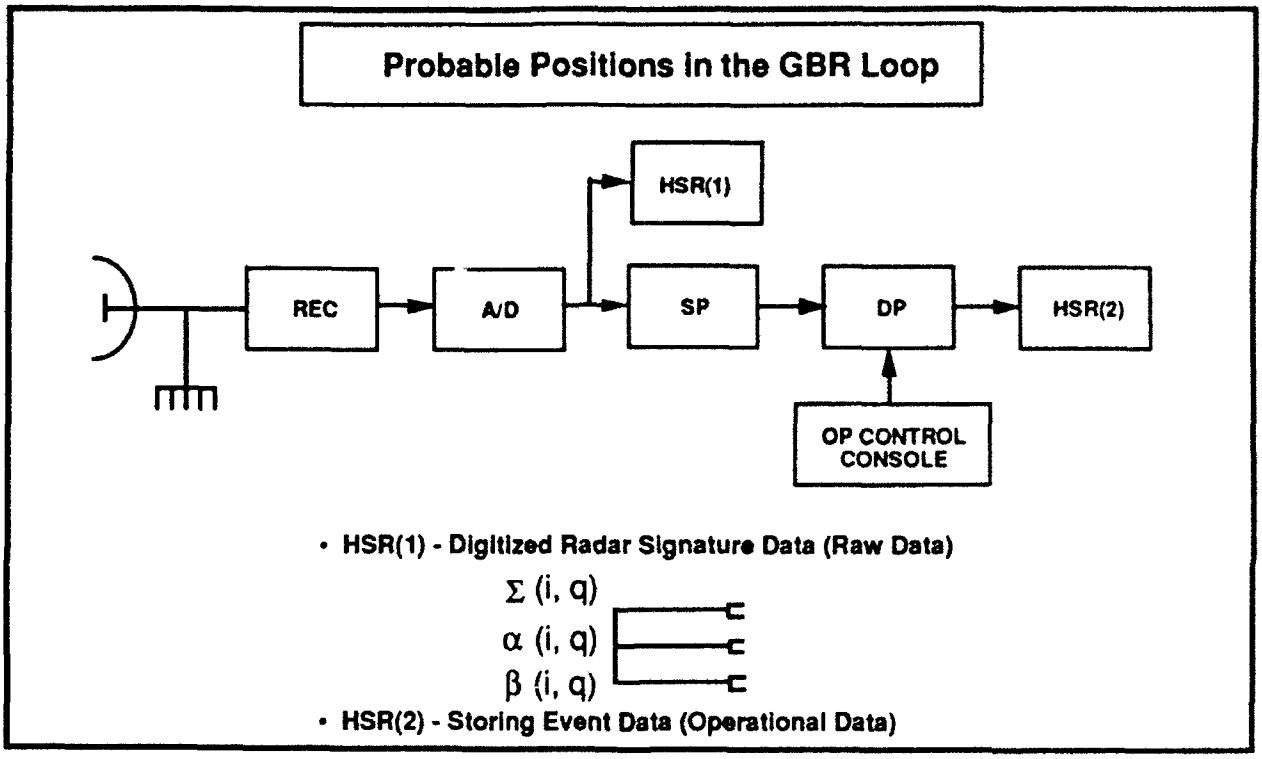

Figure 39. Probable HSR Positions and Roles in the GBR Loop

During the past contract period, COLSA personnel evaluated four different types of HSRs. All of them use magnetic mediums; optical medium devices were considered, but the transfer rates seem insufficient for GBR. The systems are listed in Figure 40. The Sony system is a hard drive system, and the other three are digital cassette systems. The Ampex unit is a robotic system that utilizes 256 cartridges at 25 gigabytes per cartridge to provide a capacity of 6.4 terabytes.

- Storage Concepts RAID 51 System
  - * Redundant Array of InexpensiveDisks system
- Sony Dir-1000 Tape Drive System
  - * High-speed, High-capacity Digital Cassette Tapes
- Datatape DCTR LP4000 System
  - * Digital Cassette Tape Recorder System
- Ampex DST 800 ACL
  - * Automated Cartridge Library System (Robotic)

Figure 40. Four Leading HSR Systems

The specifications for each of the systems are given in Table 10. It is important to note that with some of the systems, multiple machine synchronization is possible, i.e., multiple units can be combined to boost performance by nearly N times for N units. For example, Datatape produces a Cluster Synchronization Module that allows a combination of 4 LP400s. In the table, transfer rates and capacity generally increase from left to right. The Ampex machine has impressive numbers, especially for capacity, but it is a very large and heavy machine compared to the others.

TABLE 10. A COMPARISON OF FOUR LEADING HSR SYSTEMS

	CONCEPT 51 RAID	SONY DIR-1000	DATATAPE LP-400	AMPEX DST 800
X-Fer Rate	20 MBs	32 MBs	50 MBs	60 MBs
Capacity	108 GB	96 GB	100 GB	6,400 GB
Record Time	90 min.	50 min	33 min	1776 min
System Cost	$540 k	$275 k	$300 k	$875 k
Physical Volume	20.78 ft^3	4.31 ft^3	4.6 ft^3	139.3 ft^3
Weight	1,272 lbs	147 lbs	170 lbs	2,850 lbs
Operation	Download to Secondary	Remove Cartridge	Remove Cartridge	Remove Cartridges
Availability	Now	Now	Now	Now

In Table 11 the systems are compared once again, this time on the basis of transfer rate usage and capacity usage on a per mission basis. It was assumed that a mission 20 minutes long and a data rate of 20 megabytes per second[1]. The transfer rates of the Concept and the Sony are marginal, while the Datatape and the Ampex have more storage capacity.

TABLE 11. A COMPARISON OF LEADING HSRs BY MISSION

	CONCEPT 51 RAID	SONY DIR-1000	DATATAPE LP-400	AMPEX DST 800
Percentage of Throughput Use	100 %	62 %	40 %	33 %
Percentage of Capacity Used	22 %	25 %	2 %	0.4 %

In summary, the most appropriate HSRs for the GBR seem to be digital cassette recorders. Also, the transfer rates and capacities necessary for the GBR seem to be available with 1992 technology.

3.2.6.1 Datatape/HIPPI Interface

It is known that Datatape has not been officially chosen as the vendor for the HSR, and that Sony and GE are still being considered. However, since information on Datatape products was available to COLSA, their hardware was evaluated. The HIPPI interface may be an area of concern because it is still under development, and, therefore, does not have the maturity and field-proveness desired of GBR components.

[1] A briefing given by Ralph McManus of Raytheon in March 1989 listed a data rate of 45.62 megabytes per second. Since two of the machines considered do not have transfer rates that fast, the data rate was normalize to the slowest of the four machines.

Datatape's fastest HSR is the DCTR-LP 400. The 400 provides variable recording rates in an 8:1 range, beginning at 50 Mb/s and peaking at 400 Mb/s. A variable rate buffer allows for any continuous rate between 0-400 Mb/s. It can record in burst mode at 480 Mb/s for a period of about 15 seconds. The bit error rate is 1010 surpassing the GBR requirement of 107. At 400 Mb/s the record time for a large D1 cassette is 33 minutes. The reproduce rate (playback) can be set in the same way as the record rate, i.e. an 8:1 ratio with 400 Mb/s maximum. Large D1 Cassettes have a 95 Gigabyte capacity and are light, compact, and easily removable. The Datatape system proposed for GBR will be composed of 3 parts: a HIPPI-VRB Adapter, a Variable Rate Buffer (VRB) and the Recorder itself. Price and footprint for each component are given in Table 12.

TABLE 12. PRICE AND FOOTPRINT FOR THE DATATAPE LP-400 COMPLETE SYSTEM

Part	PRICE	W	H	D	LBS
Recorder	300 K	19	14	30	170
VRB	110 K	19	7	24	30
HIPPI Adapter	20 K	19	3.5	20	15
Total	430 K	19	24.5	30	215

The LP400 and VRB are mature products, but the adapter is still under development.

As shown in Figure 41, the adapter is governed by an on-board controller which in-turn is controlled through an ethernet interface. HIPPI communications are used on the HIPPI side. On the VRB side, RS-422 protocol is used for control signals and VME protocol for the data. Throughout the adapter, first-in-first-out (FIFO) queues are used to buffer I/O data.

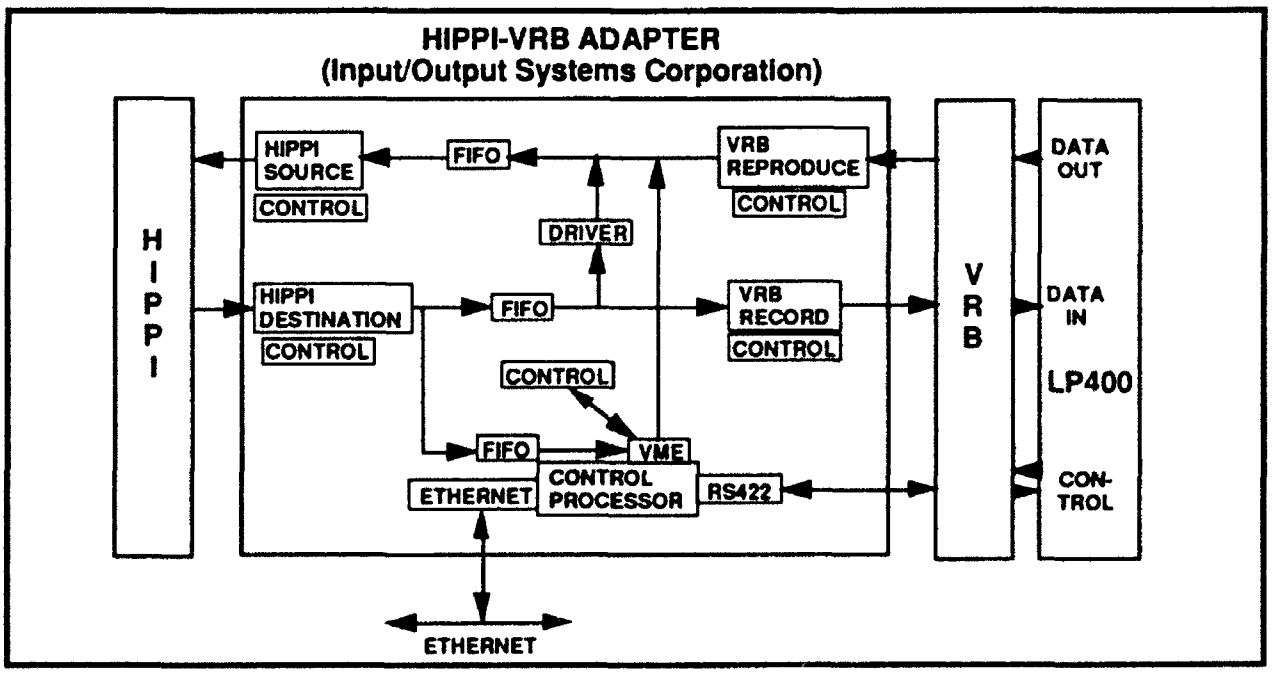

Figure 41. Datatape's (IOSC) HIPPI-VRB Adapter

3.2.6.2 Datatape LP400 Critical Path

One of the data recording paths for GBR is shown in Figure 42. A preliminary analysis indicates that a Datatape LP400 will do the job, but with only a small margin of safety. It may be significant to note that the Sony DIR-1000 recorder does not meet throughput requirements, and is not a legitimate candidate unless multiple Sonys are used.

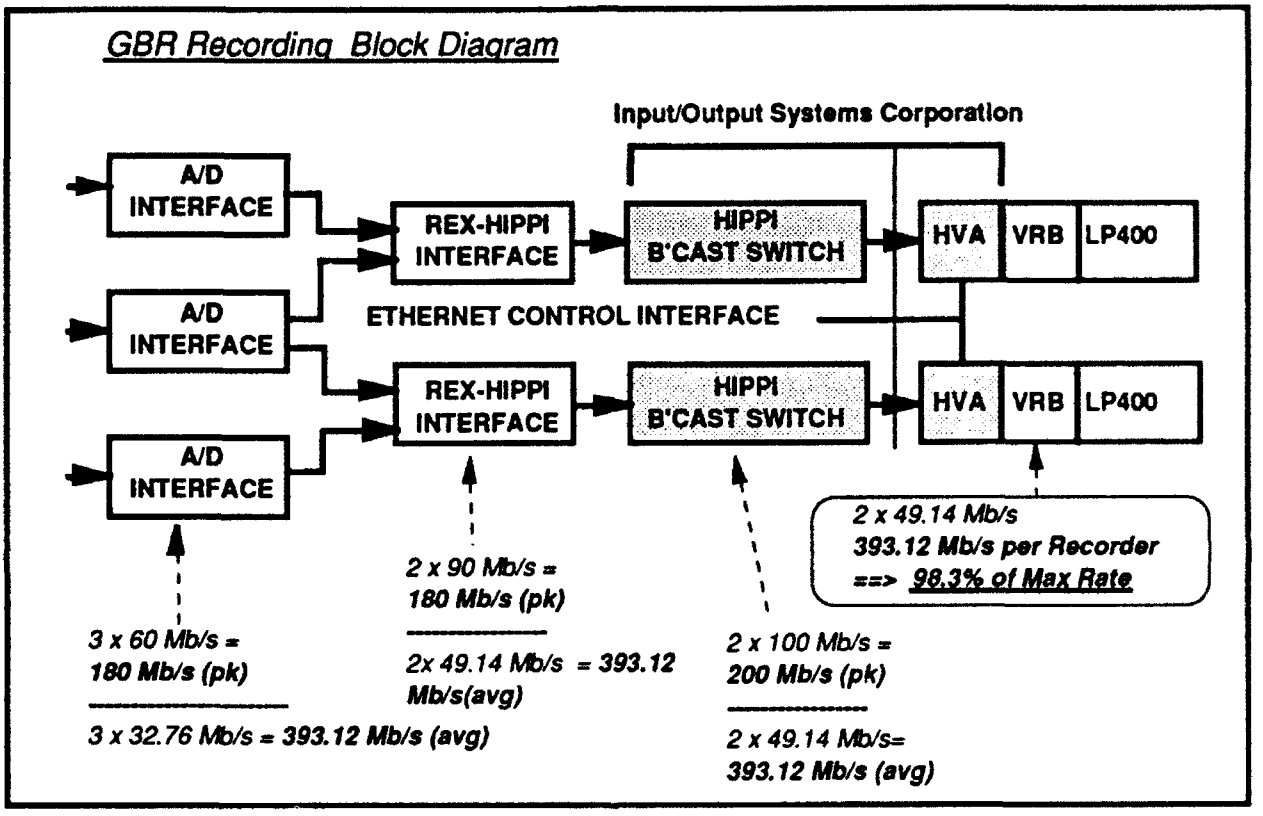

Figure 42. GBR Recording Block Diagram

The figure shows four legs of the high speed recording path. The analog to digital boxes each have a peak capability of 60 Mb/s, multiplying by 3 gives an aggregate rate of 180 Mb/s peak. The actual sustained rate that is being given by Raytheon is 54.6% of peak. The result is a sustained aggregate rate of 393.12 Mb/s. The two REX-HIPPI boxes maintain the average aggregate of 393.12 Mb/s. The HIPPI broadcast switches each should sustain at least the 49.14 rate, i.e., a 393.12 aggregate.

Figure 42 above, which was obtained from Raytheon, through Datatape representatives shows two HSRs being used. If this were the case each recorder would be responsible for one-half of the 393.12 average, i.e., 196.56 Mbits/s; because the LP400 accepts 400 Mbits/s this would leave a margin for safety of slightly over 50%. However, if we consider peak rates, which may be in effect at times, the margin of safety becomes very slim. The recorders would each be responsible for 393.12 Mb/s, which is 98.3% of the recorder's maximum capability.

If, as suggested in other Raytheon documents, there is only one recorder in this section than the margin of safety becomes very slim, even at average rates. An aggregate of 393.12 Mb/s going to a single 400 Mb/s recorder would use 98.3 % of the recorder's capability. At peak rates one recorder would only be be able to handle approximately half of the throughput, and would, therefore, be inadequate. Given the above numbers, it is clear that two LP400s are needed.

4. TRIPS

During the third week of October 1992, COLSA personnel attended Raytheon's kick-off meeting, in Wayland, Massachusetts. Many interesting ideas were exchanged, with Raytheon's design engineers presenting their ideology and views on the direction of the program. Subjects of technical interest included the digital hardware designs proposed for both the National Missile Defense (NMD)-GBR and TMD-GBR. Raytheon's selections of hardware include two VAX-7000 computers as the GBR data processor, and five MasPar MP-2 massively parallel computers (three for the TMD-GBR) as the GBR signal processor. The communications portion of the system design is shown to utilize HIPPI lines (100MB/s data transfer capability) between the REXTTG and Signal Processor (SP), and FDDI lines (10MB/s data transfer capability) between the SP, BSG, Transmitter, Display Consoles, and the Data Processor. The transfer rates expected for the radar communications are well within the capabilities of both communication schemes and should meet program requirements.

During December COLSA personnel traveled to Raytheon's facilities in Wayland, Massachusetts, in order to attend a number of preliminary hardware design presentations. Hardware designs on subjects such as the SP, the data processor, the REX, the BSG and the HSRs were introduced. The details of these presentations can be found in the written material Raytheon furnished to government personnel and subcontractors alike.

During the month of February COLSA personnel attended Raytheon's PDR of the TMD/Demonstration/Validation (DEM/VAL) and User Operational Evaluation (UOE) GBR systems. Many topics of interest were covered during the three day presentation. Topics such as REX/TTG design configurations, Signal Processing CIs, Data Processing CIs, BSG CIs, hardware testability, Electromagnetic Interference (EMI), Electromagnetic Compatibility (EMC), Electromagnetic Pulse (EMP), grounding and logistics support were all covered. It seems that even though Raytheon engineers have not finalized all of the design, they have made great strides in developing sound digital hardware designs.

5. CONCLUSIONS AND RECOMMENDATIONS

The GBR program is still in the very early design phases, as a result, COLSA can not yet determine weather the Raytheon approach is the most suitable way to design and build GBR radars. It is important to emphasize that Raytheon's approach to date, seems to be a very reasonable one, especially since a number of off-the-shelf digital hardware components are implemented in the once totally proprietary digital hardware design.

Raytheon's engineers seem to have taken into consideration the majority of critical factors, which will determine the success or the failure of the GBR program. Raytheon engineers have also produced GBR hardware configurations with a wide margin of safety built into their designs (FDDI uses only half of its maximum capability), again a very important factor for a successful GBR system.

Also, the use of the latest hardware available on the market (MasPars, DataTapes etc.) which were proven to work reliably, will be an invaluable asset to the GBR program. As a result,the system will not be prematurely outdated when fielded.

Some of the weaker aspects of the known Raytheon hardware designs seem to be on the grounding configuration that will be used when testing the TMD systems at Shite Sands Missile Range (WSMR). Judgement by COLSA will be reserved until final designs are available and can be closely examined

It is also COLSA's intention to follow the communication/bus designs very closely in order to insure that the needed data transfer speeds will be achieved with reliability, by both the computers and data recorders. A number of the HIPPI and FDDI interface cards are new Raytheon designs that could be risky if they are not completed and tested within a reasonable amount of time.

Finally it is COLSA's intent to set up, as soon as possible, the known GBR hardware configuration of the SP, Data Processor, Data Recorders and HIPPI/FDDI networks, in order to perform actual hardware and software tests, which should help prove or disprove the soundness of Raytheon's designs.