Autonomous Surveillance Satellite
A Paper Presented To
AIAA/Utah State University Conference
On
Small Satellites

Gary S. Grisbeck - Control Data Space Systems Division

John G. Doleman - Control Data Space Systems Division

Robert G. Reed - Control Data Space Systems Division

Introduction

This paper describes the potential to enhance the performance and reliability of small satellites by using increased levels of digital processing on-board the spacecraft. It describes an architecture, developed by Control Data, that will satisfy its users' needs and provide mechanisms for high reliability. The term "SMARTSAT" means a small satellite that has significant "smarts" (processing capability) on-board.

Increasing The Capabilities of Small Satellites

Small satellites have two operationally driven requirements that differentiate their operation from large satellite systems:

(1) A need for operational autonomy, and
(2) A need to directly supply users with data.

These requirements relate to the fact that the small satellite will probably be in Low Earth Orbit (LEO) and will therefore have short time windows over a ground station. This condition impacts the total architecture and design of the spacecraft and provides the opportunity for significant improvements in capability by sophisticated on-board processing.

Need For Autonomy

A small satellite must be capable of maintaining its health and welfare for sustained periods of time without involvement of a ground support center. Depending on the orbit selected, a small satellite in a low earth orbit may only pass over a ground support facility once or twice per day. As a result, regular and timely communications are difficult, if not impossible. While the placement of several mission support centers around the world could alleviate this problem, the expense of installing and maintaining hardware and trained support personnel is prohibitive. In addition, small satellite economics will probably preclude the cost associated with a traditional mission support center.

A second factor emphasizing the benefit of autonomy is the vulnerability of fixed ground stations. This is especially true of any satellite that would be used during periods of intense conflict. At such times, the task of rendering specific ground stations inoperable is not difficult considering today's technology. In addition, experiences with natural events, such as seismic disturbances, impacting space operations serve to emphasize further the advantages of highly autonomous space systems.

A third factor related to autonomous operation is ground station workload. As small satellite constellations become more dense, the task of providing traditional levels of ground support will increase. Without enhanced levels of autonomy in the satellite, the complexity of managing these constellations will add cost and workload to a force structure that needs relief in both areas. New control and data processing centers would need to be constructed to support tactical satellites.

In addition, for tactical applications, a launch-on-demand capability requires very responsive launches. Associated with this capability is the urgent need for the spacecraft to quickly stabilize and start its assigned mission. The launch-on-demand capability will need self-aligning systems free of human dependence if timely response is to be realized.

The net result of these considerations is that a high degree of operational autonomy is very desirable for the success of the small satellite concepts being advanced today. This is especially true for tactical applications where long term stability, quick response, and error-free operation are essential for mission success.

Direct Downlinks to the User

To take full advantage of the opportunities offered by small satellite systems, mission related data must be downlinked directly to users on the ground. This eliminates the need for, and the delay associated with, having a large ground processing station in the distribution path. This is especially critical for applications that require timely data

distribution in which critical decisions must be made in real-time.

If small satellites are to be successful, the data distribution system must not be forced to travel the traditional path from space to the ultimate user. The existing system was not designed for rapid distribution of data to tactical users, although it is satisfactory for its intended use. The data distribution pipelines in place today are already stressed with data overload; it would be very expensive to expand to cope with increased data loads on a real-time basis.

The above considerations require a direct downlink to the end users in a form that can be readily exploited on the ground. Because of the "time over station" associated with a LEO SMALLSAT, the direct downlink also implies a relatively narrow bandwidth to accommodate the use of unsophisticated ground receiving equipment. This translates to a need for higher levels of on-board processing to deliver processed data over low bandwidth channels to users requiring fast, accurate, and relatively simple data capture.

Enabling Capabilities

During the past twenty five years there have been significant technology advances in digital electronics. While the processing power of computers has grown well beyond early planners' expectations, the space community has chosen to use a very conservative approach with minimal computing in space. The advent of Strategic Defense Initiative (SDI) forced the concept planners to break from this tradition. The primary cause is that the problem demands vast improvements in space computing for ballistic missile defense battle stations to be responsive.

A compounding problem on the technology side is that as the geometry size of electronic devices was reduced, the circuits became more vulnerable to cosmic radiation. Also, as circuit speed increased, there was a corresponding criticality introduced to the internal timing

sequences that were affected by radiation. SDI has funded research addressing these technology issues, and while there is much work yet to be done in this area, significant progress has been made.

Space Computer Systems

As a result of the SDI actions described above, and as a result of more than a decade of experience using increasingly sophisticated computing in spacecraft, industry now has key technology centers that are poised to construct space computer systems that are orders of magnitude more capable in terms of on-board processing power. These on-board systems focus on autonomy, fault tolerance, reliability, and modularization. They are designed to withstand the environmental stresses associated with the space environment, and to provide reliable, error-free operation.

Tailorable Concept

An important aspect of achieving the size, weight, and power envelope required for space operations is the tailoring of the spacecraft to a very specific suite of missions. A processing concept is required that can be easily and readily adapted to the range of requirements imposed by the various types of missions and sensors. This concept must allow for varying processor types and for degradation over the mission life. This approach provides the opportunity to maximize the synergy between the on-board processor and chip level processing capability that can be embedded into the sensor subsystems. This will provide levels of efficiency in data handling that are not feasible in today's environment.

Small Satellite Suitability

The advent of the small satellite is one of the developments that enables the advances in on-board space processing power. This is the result of several factors:

- Small satellites are mission specific. This provides the basis for advancing the

integration of the processing and the sensor system beyond traditional multi-mission space systems.

- The small satellite community is adopting a new philosophy for the design, assembly, and operation of space systems. They are attempting to avoid many of the very rigid cultural issues associated with larger satellite programs that must, by definition, use very conservative technologies.

- While the digital processing subsystem is an expensive component of present space systems, high levels of integration in the small satellite can result in processing that is more optimized as to function, and thus more system cost effective. By recognizing that the mission life will be of shorter duration, the digital processing suite onboard can be designed to take advantage of relaxed failure rate allocations. This will bring the cost of the on-board processing into line with small satellite expectations without compromising the probability of success of the mission.

Downlink Considerations

Simple Receivers

Since it is possible to do data processing onboard, downlinked data can be simple and comparatively low volume. Under these conditions, a relatively small transceiver adequately supports the user.

Minimal Training Levels

User-friendly on-board processing of data eliminates the need for data processing experts at the ground station. This allows a force commander to employ less specialized personnel in the receipt and display of pertinent data.

Portability

The simplicity of the forward located ground station enables the direct field usage of the mission related data. A forward control facility would have a workstation similar to today's commercial workstations, i.e., PC-based. The workstation would be used to receive formatted and image data from the SMARTSAT and to superimpose collateral locally stored imagery. This local imagery could include any (semi-)permanent data such as base maps, geological maps, or other icons representing "friendly" resources (or "hostiles" from other data sources). The required equipment could be transported by jeep for maximum portability.

Possible Uplink for Tasking

An uplink could be provided for the specific tasking of the SMARTSAT. Among other tasks, this might include moving the field of view to re-examine a target recently detected, performing an electronic zoom of a visible sensor to characterize a suspected target, converting the visible sensor to a UV sensor to examine a suspected target in the UV spectrum, or directing the SMARTSAT to send surveillance information via communication links to another commander. Redirection of the orbit would be beyond the permitted capability of the forward area located user, but would be accomplished by a CONUS ground station.

Applications

SMARTSATs have many military and commercial applications. These applications, in general, require sensing of information from earth and retransmitting it in the same or altered form back to the earth. The sensors to be used are a function of the information required, and the retransmitted data needed is a function of the desires of the operator(s). These applications are divided into the following four groups.

Communications

Voice and data communications between widely separate areas on the earth can be most easily accomplished by a "store and forward" system using relay stations that have "line of sight" access to the users. Geosynchronous satellites perform this function today. Under

various conditions, however, these communication links could very easily be saturated with traffic. At such times, supplementary links could be established through SMART-SATs.

Weather Monitoring

Existing weather observation systems often do not yield data with sufficient resolution for detailed (localized) planning and operations. A Low Earth Orbit (LEO) weather observation system could give a much more detailed picture of local conditions.

Mapping

Fine detail mapping of either geodetic features or resources could be accomplished by a LEO system. This mapping could provide detailed mapping of selected areas of the earth to provide previously unavailable or rapidly changing data.

Reconnaissance

Fine detail reconnaissance of tactical situations could be readily available to a field commander using small LEO SMARTSATs. Depending on the sensor suite on-board, threat warning and data could be available on a real-time basis. The concept of maintaining satellites on the ground so that they can be "launched-on-demand" in a very short time period has become very attractive to the users. Currently, military reconnaissance is the driving requirement for small satellites.

Processing Requirements

Data processing requirements for the various applications will vary widely. These variations will be both in type of processing required and in quantity.

Communications

Communications satellites need processing for switching and routing of messages, both voice and data. Short term storage will be required for messages for situations where simultaneous receipt and transmission is not practical, such as in the case where both stations are not "visible" and linking facilities are saturated. Estimates of the processing requirements for this type of system are roughly 10 MIPS (Million Instructions Per Second) and 1 MBIT (MegaBIT) of RAM (Random Access Memory).

Weather Monitoring

Weather monitoring is an image generation and storage problem, with processed images downlinked to the user. The processing requirements for this type of mission have been estimated at roughly 20 to 30 MIPS of processing and roughly 10 to 40 MBIT of RAM.

Mapping

Mapping is another image generation and processing problem, but is different in that the image data can be processed on-board and only data of interest is downlinked to the user. Data requirements for this type of system have been estimated at 20 to 40 MIPS and 20 to 30 MBIT of RAM.

Reconnaissance

Multiple Sensor Suites

Tactical reconnaissance is an area in which on-board processing can provide dramatic results. Today's technologies are producing sensor systems of such capability and compactness that it is conceivable to place multiple sensors on a given SMARTSAT to provide previously unavailable capabilities. This concept enables multi-sensor coverage of target signatures, thus increasing threat detection confidence. For example, an IR, a visible, and an electronic intelligence sensor could be placed on the same SMARTSAT. This sensor suite would provide broad spectrum collateral coverage of suspect areas. The processing requirements for these sensors are estimated to be roughly 2 to 5 MIPS and 1 to 10 MBIT of RAM for each sensor.

In the near future it will be possible to fly an adaptive Synthetic Aperture Radar (SAR) on a

SMARTSAT. This system could have the capability to both image ground areas and to detect and image moving targets. Moving target detection and imaging would be a great benefit to a force commander in the identification of incoming threats. The radar subsystem would require considerably more on-board processing than the previously mentioned sensors, i.e., approximately 2700 MIPS of processing and 1600 MBIT of RAM. Special purpose signal processors would perform the bulk of the processing load. In this application, the downlink would consist of prepared images of either ground areas or moving targets, or locations of identified targets.

A long standing goal of defense planners has been a "track from source" capability, i.e., detecting threats immediately after departure from their bases, rather than electronic fences, i.e., detecting at a perimeter. By using a SMARTSAT constellation to draw a fence around known staging bases, coverage of the threat is provided every 50 to 60 minutes (for a constellation of six satellites). This capability provides a very important aspect of early tactical warning and assessment.

A secondary capability that greatly extends the track from source capability is the use of the SAR mode on a companion SMARTSAT to provide an air order of battle update of threat airfields. This periodic monitoring of the source bases would provide critical intelligence of activities at the threat base. These capabilities have previously been infeasible because of the inability of present systems to see/track aircraft from geosynchronous altitudes.

A SMARTSAT constellation, launched-on-demand during times of heightened tension, would offer a show of deterrence before the onset of hostilities while providing the primary mission functions of warning and target detection.

Downlink Simplicity

With multi-sensor suite, it is possible to downlink only the earth coordinate location

and velocity vectors for individual threats. This would enable a force commander to counter such threats effectively with minimal wasted resources. Thus, the downlink required would be very small and the ground processing required would be only that needed to display the threats in whatever display mode would be most informative.

Encryption Issues

For a tactical SMARTSAT, communications security is a significant issue. Minimally, uplink security must be very tight. Downlink security may need to be tight or loose, depending on the sophistication of the on-board processing and the operational concept. Encryption will double the telemetry processing (up and down).

Housekeeping

Attitude control, orbital maintenance, and housekeeping functions would be required on all applications. This effort requires 1 to 2 MIPS of processing and 2 to 3 MBIT of RAM.

Processing System Concept

Introduction

SMARTSATs will require sufficient resources to perform basic payload control and housekeeping functions in addition to mission functions. Special telemetry related formatting not accomplished by the telemetry hardware can be included in the housekeeping function.

Housekeeping includes statusing of satellite resources/subsystems and control of basic satellite functions that are prerequisites to performing the mission, including:

- Power Distribution and Control. This includes sensing of battery charge state, (dis)charging rate, voltage level, load shedding, solar array health status, and solar array pointing control.

- Thermal Control. This includes monitoring thermal sensors to detect hot (or cold)

spots and to heat or cool subsystems or parts thereof when required. Some simple thermal management may be done by local resources such as thermostats.

- Attitude Determination And Control (ADAC). This is a necessity for all satellites, including SMARTSATs. In the past, ADAC has been effected through analog hardware. Recently, there has been impetus to implement the control laws digitally. This provides more flexibility and potentially more accuracy in that control law implementation (and sensor sampling rate control) may now be done via software. This standardization of hardware and specific mission tailoring by software eliminates the need for different hardwired solutions to each application. This reduces the total cost of the SMARTSAT family amortized over a variety of missions and mission requirements.

- Telemetry. This function includes formatting, time-tagging, and packetizing of data (mission or routine satellite health status), receipt and distribution of commands from the ground, encryption, and verification of command authenticity.

Housekeeping and mission support functions, i.e., not part of the applications code, e.g., updating time tagged stored commands or shutdown of sensors or other resources in a battery/power crisis, are easily handled with today's spaceborne processors. However, the significantly larger throughput and memory requirements associated with the missions described indicates that a single processor is inadequate, even state-of-the art flight computers. This condition is further exacerbated by the previously mentioned desire to maximize satellite autonomy, which requires significantly more on-board throughput, memory capacity, and possibly different processing element types. In response, Control Data has developed the Spaceborne Processing Array (SPA) concept.

Processing System Capabilities

The Spaceborne Processing Array (SPA) concept provides the following capabilities:

- SPA supports a variety of missions. The system is modular, expandable, reconfigurable, and uses a family of processing modules that can be mixed and matched.

- SPA, through its Module Interconnect Network (MIN), will allow data paths to be configured to match application data flow and a variety of data structures including rings, arrays, parallel or pipeline, and fanout. This provides maximal data path flexibility.

- Data path reconfiguration isolates failed modules and incorporates functional modules from a group of similar spares. This achieves maximum reliability, especially for longer life missions.

- SPA and MIN together provide support for a variety of strategies to effect fault detection, isolation, assessment, and recovery (FDIAR). This is true at the system, module, and module element level.

- MIN provides a high performance interconnect system to support signal and data processing at very high system throughput rates.

- The multi-modular characteristics of SPA, along with the configurable features of MIN, allow for a system with distributed control. This greatly reduces the likelihood of a system failure due to loss of a critical control function. Single point failure modes at the system level may be reduced or eliminated with this capability.

Processing System Detailed Description

The SPA concept is a system capable of providing high performance, high reliability digital signal and data processing systems. An SPA is a modular collection of processing and

external input/output (I/O) resources. The modules are interconnected by the MIN. This network has a high bandwidth, low overhead, is reconfigurable, and is fault tolerant. The SPA concept allows a heterogeneous mix of modules, achieving greater optimization than could be achieved by universal modules. The following modules can be attached to the MIN:

- General purpose computers - 16 and 32 bit
- Fixed and floating point signal processors
- Application specific I/O modules, i.e., high bandwidth sensor interfaces
- Large volume, non-volatile memory units.

Every module contains processing resources, instruction and data memory, and an interface to the MIN. Applications are generally, but not necessarily, data driven.

The MIN is the heart of the SPA system; it distributes both power and data. Its bidirectional, segmented ring approach eliminates fan-out problems while providing maximal path flexibility to support different missions. The MIN uses software configurable data paths to form logical data path structures on an interconnect that is a two dimensional physical mesh. The modules form a plane with each processor having connections only to its four nearest neighbors, thus forming a Manhattan geometry (orthogonal interconnect). The circuit switching capability is used both to establish the initial data path configuration and to reconfigure the SPA in order to maintain an active configuration. This is accomplished by isolating failed modules and switching spare units into the active configuration. The physical implementation of this scheme allows the plane of processors to be formed logically into a toroid. This minimizes the longest data path, which maximizes overall system operating speed.
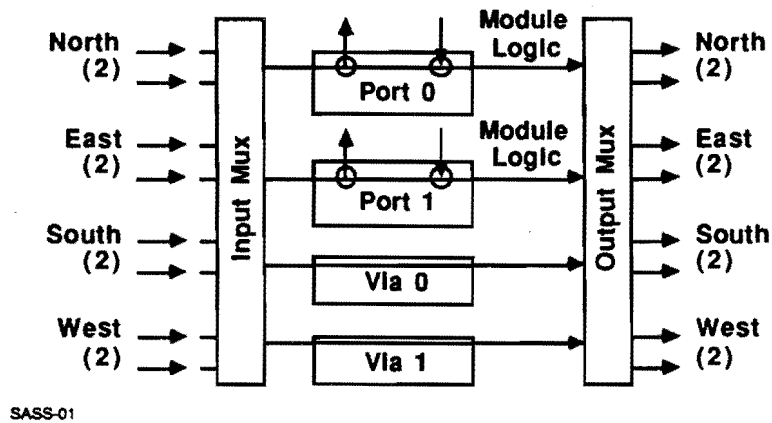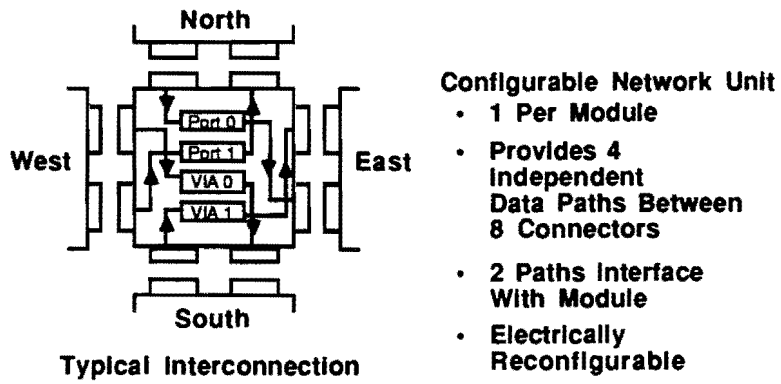
The MIN is a fully distributed system. The element of the MIN residing in each node is the Configurable Network Unit (CNU). The CNUs in each unit are identical. Each CNU has two data path connections to each of its four nearest neighbors, for a total of eight

connections. Each of the eight interconnects is called an edge connector. Each edge connector can be programmed to be an input or output, but not both simultaneously. Within the MIN, there are four internal data paths between the edge connectors and a system of multiplexers. The multiplexers allow any edge connector to provide the input to one of the internal data paths. Each edge connector that is programmed to be an output can take its data from any of the four internal data paths. In this way, fanout is achieved by programming several outputs from a given data path. Thus, up to four internal data paths can be established, with each data path consisting of an input edge connector, the internal path, and an output edge connector.

Two of the four internal data paths have address recognition logic. These are called Network Interface Units (NIU) and will extract any word of data that is addressed to them, sending it to the processor associated with that node on the MIN. The ports also can place data in any unused data slot in the data path. When a port removes a data word, it can place (output) a new data word into the same slot for transmission down the MIN.
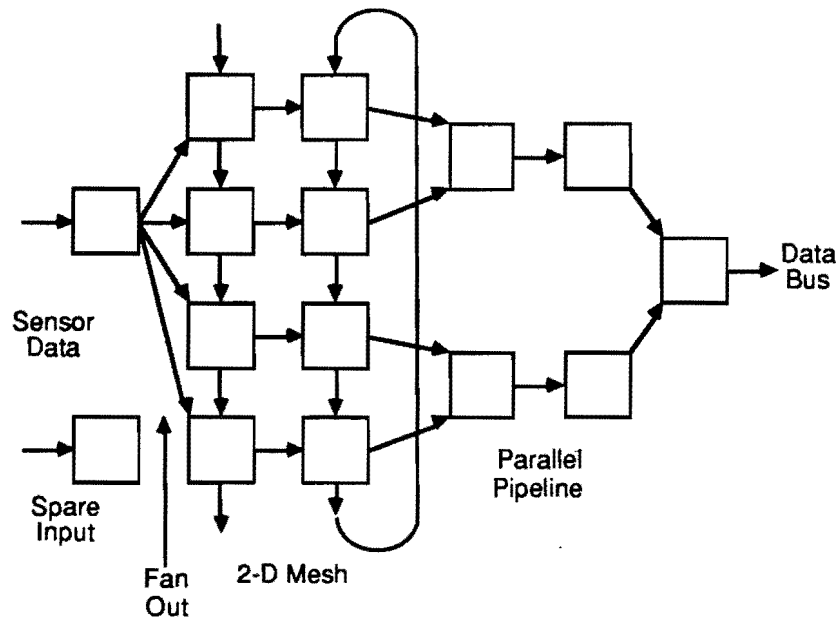
The other two internal data paths merely pass data through the CNU and to the next node in the MIN. These internal data paths are called vias. A port can act as a via by putting the same data word back on the data path and sending it downstream to the next node. An example of the internal connections of a CNU is shown in Figure 1. Examples of SPA network structures are shown in Figure 2. A complex example of an SPA network, with failed CNUs and with a mix of processor types, is shown in Figure 3.

Present MIN hardware has information transferred at a rate of up to 8 million data words per second. The 52 bit data words are comprised of 32 bits of data, an 8 bit destination code, and other control information. Each data word has its own destination address; this allows data from different sources and going to different destinations to be interleaved on a given data path on a word by word basis. A

Configurable Network Unit
- 1 Per Module
- Provides 4 Independent Data Paths Between 8 Connectors
- 2 Paths Interface With Module
- Electrically Reconfigurable

Typical Interconnection

SASS-01

Figure 1. CNU Connectivity

SASS-02

Figure 2. Examples of SPA Network Structures

Logical Design

Physical Array

2 Failed CNUs
and Processors

Failed
CNU
and
Processor

Failed
CNU
and
Processor

Signal Processing
Unit (SPU)

Input Output
Unit (IOU)

General
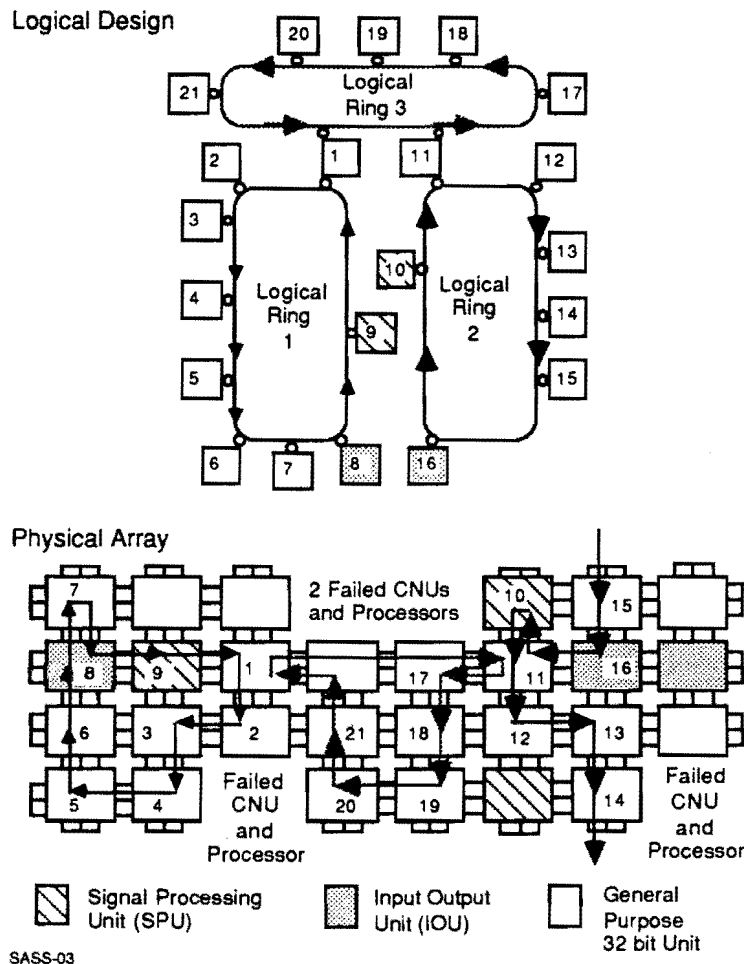Purpose
32 bit Unit

SASS-03

Figure 3. Complex SPA Network

data path does not have to be reserved for exclusive use by any given module, which greatly reduces system latency.

With MIN architecture,

• Programmable data paths and destination codes will minimize hardware and software overhead;

• No store and forward is required, as occurs in some other interconnect schemes;

• The port only handles data for the associated modules; and

• The application software is configuration independent.

Each data word carrying its own destination code means that no channel request protocol is required to send data. Data is queued for transfer and goes whenever room is available. Figures 1 and 2 show the rich interconnectivity of the MIN. The ability to reconfigure the MIN under software control makes an SPA system incredibly reliable and fault tolerant.

Reliability - Fault Avoidance, Fault Detection, and Fault Tolerance

The ability of an SPA system to achieve the required mission reliability is achieved through a combination of fault tolerance and fault avoidance at both the system and the module level.

Fault Avoidance

Fault avoidance is achieved by designing the hardware to be inherently reliable. At the SPA system level, the CNU elements in the MIN are designed to be one tenth as likely to suffer a failure as are their associated processing modules. They were thus designed because studies have shown that the total loss of a CNU at a MIN node is a more serious resource loss to the SPA system than the loss of a processor module. This is because of the attendant loss of connectivity between other processing elements at other nodes.

Fault avoidance at the processing module level is achieved by:

*   Minimizing component counts
*   Raceless design rules that eliminate failures due to changes in circuit timing and delays
*   Using Very Large Scale Integrated (VLSI) circuitry to minimize parts count
*   Time multiplexed buses to minimize the number of wires and solder joints
*   Physical/mechanical designs that keep the semiconductor junction temperatures low
*   Using CMOS/SOS technology to resist single event upsets due to radiation.

These fault avoidance features are used on the existing hardware.

Fault Detection

To use fault tolerance strategies, it is necessary to first detect the faults and then, as practical, to isolate them. Fault detection at the system level, with the MIN, and with individual modules will be discussed individually.

System Level

System level fault detection involves multiple modules. The following are common techniques used.

*   Heartbeat techniques - a heartbeat system has each module periodically send a message to the system monitor. This monitor may be a single processor, or a designated processor on each ring, or other modules in a fully distributed system. The monitors check that all active modules have reported within the appropriate time period. The method tests both the communications capabilities of the MIN and the health status of individual modules. If a message is not received from an individual module, there is probably a problem with that module. If, messages are not received from several modules, it is likely that there is a break in the MIN data path, since heartbeats from the modules downstream from the break will still be received. Heartbeat is an example of a fault detection technique that immediately provides fault isolation data. Heartbeat techniques are implemented with system level software, and will be discussed in the section on software.

*   Sample problems - where known data is injected into the data stream and the output is checked for the proper answer.

*   Watchdog timers - where application processes can be required to report their status periodically.

*   Load monitoring - in some applications, the processing load of an algorithm must be spread among several processors. However the data is distributed, by monitoring the backlog and throughput of individual modules, defective modules can be detected and their load distributed to others.

MIN Level

The MIN has hardware fault detection. The current MIN hardware uses four identical VLSI in a bit slice fashion to form the CNU. Each bit slice has a parity bit to detect errors. Parity is checked when data is removed from the data path and sent to the processor. Further, with 8 bits of data for the destination code of any data word, one of these bits can be used as parity. This allows 128 unique module destination addresses, none of which can be changed (without parity detecting it)

by a single bit flip. Software fault detection on the MIN is available as part of the SPA software operating system. Software controlled timeouts are a fundamental method of detecting total failures of ring data paths. Block counts and request/acknowledge protocols are also implemented in software to support block data transfers.

Module Level

Fault detection at the module level is achieved through use of the following mechanisms:

- Bus parity
- Microcode parity
- SECDED (single error correction, double error detection) on memory
- Parity on some internal registers
- Memory address parity
- Unimplemented memory detection
- Memory block write protection
- Privileged instructions
- Page access lock and key

Processor module diagnostics is another fault detection technique. On power up, microcoded tests check that the microinstruction sequencer is working and that special memory locations can be read without error.

Fault Tolerance

Fault tolerance provides the means for a system to operate without pathological behavior or without producing errors even though a fault has occurred at some system element. In general, fault tolerant activities should be as transparent as possible, occurring with minimal user awareness and minimal interference to the operation of the system. The use of SECDED protected memories is a good example of a user transparent fault tolerant hardware mechanism.

At the SPA system level, one of the more powerful fault tolerance strategies is the use of spares. The switching capabilities of the CNUs in the MIN allow reconfiguration of the system to eliminate failed modules from the network and to incorporate operational

modules from a pool of spares into the system's active processor suite. Pooled sparing provides significant reliability benefits over other sparing strategies (Figure 4). An example of a reconfiguration around a failed module using a spare is shown in Figure 5. Simulations have shown that in significant sized systems it is always possible to find a usable reconfiguration, even after multiple failures.
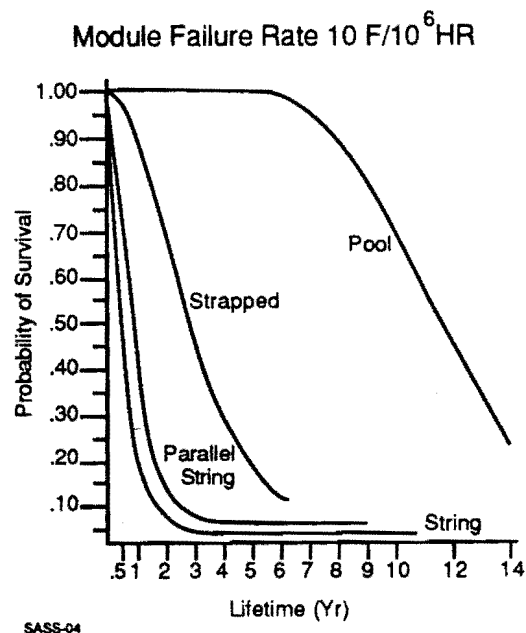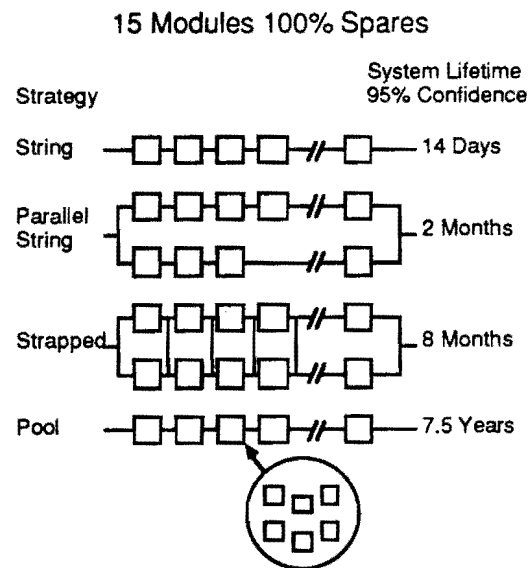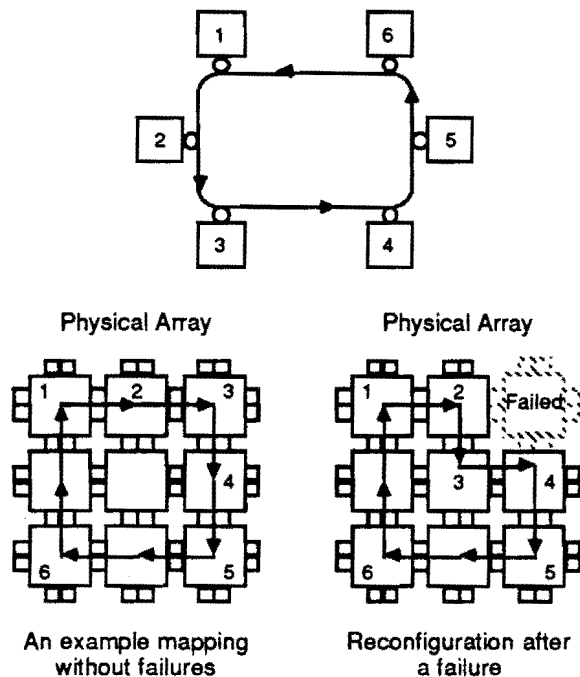

Figure 4. Advantages of Pooled Sparing

Physical Array   Physical Array

An example mapping without failures   Reconfiguration after a failure
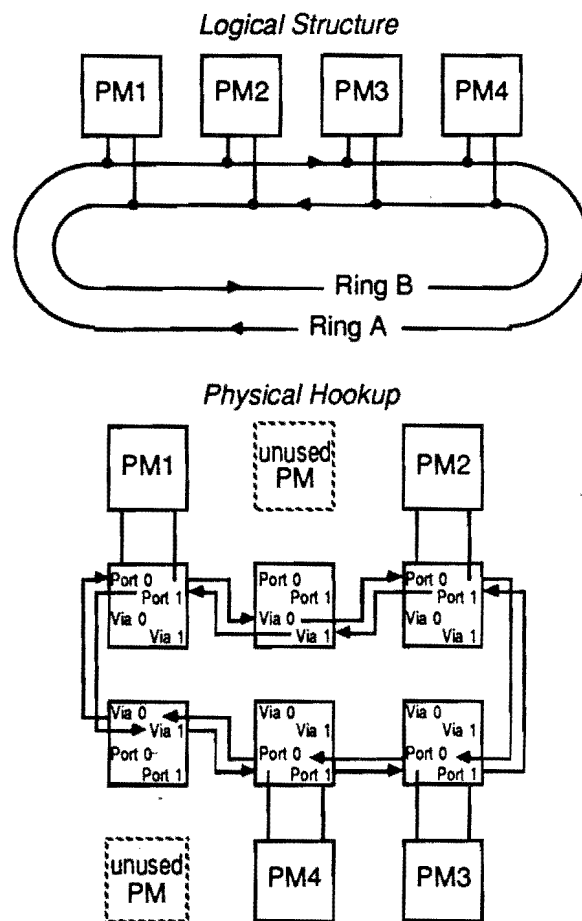
SASS-05

Figure 5. Single Reconfiguration

There are several different methods of maintaining spared resources. Spare processing resources can be powered down (cold), powered up and loaded but not in use (warm), or can be running applications (hot). Cold sparing uses minimum power; hot sparing minimizes the time required to switch in the spare in face of a failure. In addition, electronic components have a lower failure rate when powered down (dormancy factor), so that cold sparing enhances the overall life of the system. Most specifications attribute a value of ten to the dormancy factor, i.e., the failure rate for powered off hardware is one tenth of that which is powered on. Thus, for cases where the system or function can be stopped, cold sparing is preferred because it uses less power and can accrue reliability benefits from the dormancy factor. For time critical situations, where the system must operate through failures, hot sparing is preferred.

While the ultimate form of network fault tolerance is to reconfigure and replace failed processing elements with functional modules

from a spares pool, there are certain data path structures that can be programmed to allow fault tolerance without reconfiguration. One of these is the use of dual, counter-rotating rings. This gives every module two paths to every other module. In the case of the catastrophic failure of a ring element, the alternate path is used. An example of a counter-rotating ring structure is shown in Figure 6. The reaction to a ring failure in this situation is shown in Figure 7.

A last effort system level fault tolerant strategy is to move a processing task from a processor of one type to a processor of another type that has adequate, though not optimal, processing resources to perform the calculations. This option may have significant impact on system performance and is more of a "graceful

Logical Structure

Physical Hookup

SASS-06

Figure 6. Counter-Rotating Ring Structure

Failed CNU

PM1　PM2　PM3　PM4

XX

XX

Ring B
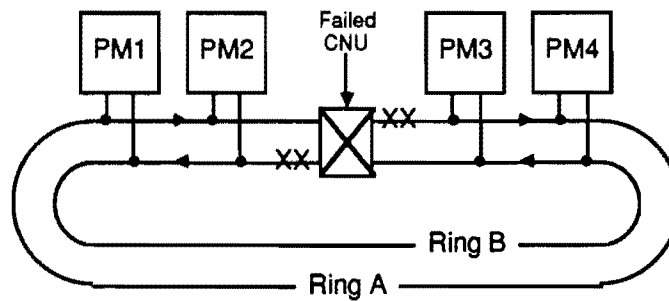
Ring A

		Destination Module			
		PM1	PM2	PM3	PM4
Source Module	PM1	X	A	B	B
	PM2	B	X	B	B
	PM3	A	A	X	A
	PM3	A	A	B	X

Usable Ring After CNU Failure

SASS-07

Figure 7. Data Path Failure with Counter-Rotating Rings

degradation" technique than a fault tolerant strategy. Similarly, the dropping (deletion and de-allocation) of processes when inadequate spares remain to perform the total mission function is an end-of-life strategy.

Another MIN fault tolerance method involves lost data. Since misaddressed data would circle a ring endlessly, there is a hardware mechanism to remove data that has circled a ring completely. To do this function, one of the CNUs in a ring would enable the lost data detector. When the lost data bit is set (upon one complete revolution of a ring) and the data encountered at the appointed node, the data is removed.

Fault tolerance at the module level is also possible. For example, SPA applications are generally structured into applications programs and data bases. The data bases are separate from the current or temporary program context, and updates to the data bases are performed at carefully selected intervals. Changes are accumulated in a local copy of

the data, and the update is not performed until the current processing is complete. This means that the time during which a data base is being updated is small, reducing the likelihood of an error during the update. Furthermore, any error that occurs before an update may be corrected by a restart, if the application allows. The data base serves as a checkpoint, in that a (checkpoint) restart (which does not completely reset the database) will "roll back" the application to its last data base update.

A simple fault tolerance mechanism is the use of one of a pair of dual CPUs within a processing module as a cold, warm, or hot spare.

The final fault tolerance technique discussion involves pairs of modules checkpointing each other's data bases, acting as partners for checkpointing purposes. In addition, one or more warm spares are configured into the system. These spares are loaded with the software, but do not process data. When a distributed system uses more than one algorithm, it is sometimes possible to have the warm

spare serve for more than one part of the total application. The warm spare in this case can contain two or more sets of software; when the spare is activated, it is notified which algorithm to run.

When a module fails, its checkpoint partner detects the failure (via the system level fault detection mechanism of heartbeat), and handles the recovery. The partner copies the checkpoint data to the appropriate spare, and sends notification to all modules in the system that the spare has begun processing for the failed module. Figure 8 illustrates pair-wise checkpointing.
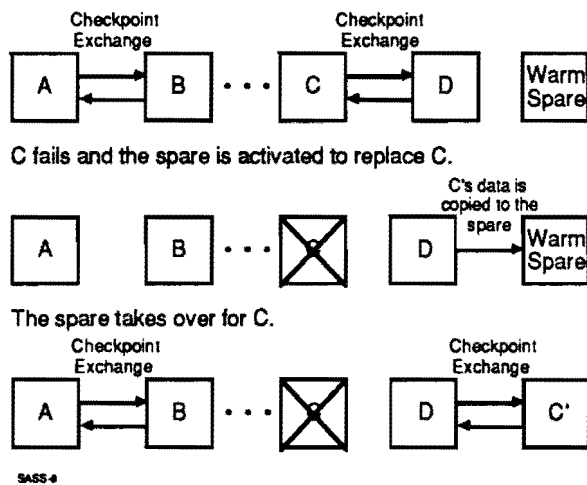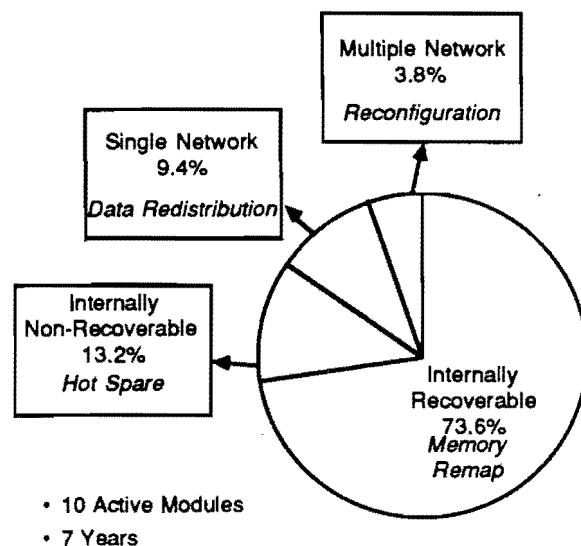

Figure 8. Pair-wise Checkpointing

The expected failure mode distribution of an SPA system with ten active modules over a seven year mission is shown in Figure 9. This data is based on a Failure Modes and Effects Analysis (FMEA) and reliability analysis of the existing system.

- 10 Active Modules
- 7 Years

1. Internal, Recoverable Failure
   Example: Memory Chip
   Recovery: Memory Remap Under Processor Control

2. Internal, Nonrecoverable Failure
   Example: CPU Fails
   Recovery: Shed Load, Shut Down, Schedule Repair

3. Single Network Failure
   Example: CNU Failure or Module Power Supply
   Recovery: Redirect Traffic on Counter-Rotating Rings, Schedule Repair

4. Multiple Network Failure
   Example: Trauma
   Recovery: On-Board Reconfiguration, Emergency Repair

5. On-Board Recovery Failure
   Example: Gremlin
   Recovery: Ground Directed

SASS-9

Figure 9. Expected Failure Mode Distribution of an SPA System

14

Existing Hardware and Software

Many of the basic elements of the data processing system concept described have already been implemented as part of development or flight programs. Among these are the Control Data 444R^2 spaceborne processor, and a nine node SPA system jointly developed by Control Data and Boeing. Part of that SPA program served to develop, test, and deliver executive/applications support software. A high speed signal processing unit (SPU) has been designed, simulated, and partly laid out and entered into a CAD data base. Plans have also been made to enable the interfacing of an existing 32-bit Reduced Instruction Set Computer (RISC) machine to an SPA.

444R^2 Spaceborne Processor

The 444R^2 ("444" because the computer is approximately a four-inch cube, R^2 because it is Rugged and Reliable) computer was designed for spacecraft control and sensor data processing. The 444R^2 is the next generation of computer to replace the Control Data 469RR, a previous generation machine with the enviable record of never having suffered a mission lifetime failure in a total flight history of over 608,000 operational hours (as of 6/15/1990), and having never delayed a launch. The new generation 444R^2 uses VLSI technology to increase performance, memory capacity, and reliability through new design, new fault tolerant features, and reduced parts count.

The 444R^2 has the following features:

- MIL-STD-1750A Instruction Set Architecture (ISA)
- Single or dual CPU, with a single processor throughput rate of 1.2 MIPS, DAIS mix, with an 8 MHz system clock
- Single Error Correction, Double Error Detection (SECDED) memory
- Single or dual serial I/O
- CMOS/SOS technology for intrinsic radiation hardness
- Small size, light weight, and low power consumption.

Figure 10 shows the 444R^2 architecture. The computer is configured internally with three functional units:

- Central Processing Unit (CPU)
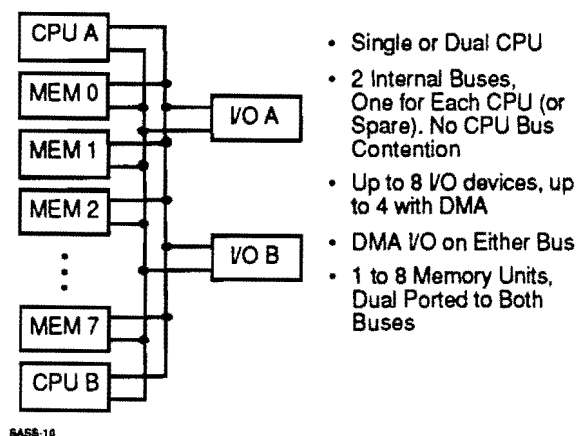- Memory Unit (MU)
- Input/output Unit (IOU)

- Single or Dual CPU
- 2 Internal Buses, One for Each CPU (or Spare). No CPU Bus Contention
- Up to 8 I/O devices, up to 4 with DMA
- DMA I/O on Either Bus
- 1 to 8 Memory Units, Dual Ported to Both Buses

Figure 10. 444R^2 Architecture

The 444R^2 can accommodate up to 12 functional units within a system, allowing the user to configure a 444R^2 to solve the mission's computing requirements. The system also allows single or dual CPUs and IOUs, and up to eight MUs. Typical 444R^2 configurations are shown in Figure 11. Data and commands are transferred between functional units on dual memory buses. A single control bus is used for interrupts, enables, clocks and other control signals.

Either CPU or the IOU can serve as the 444R^2 controller. All I/O and memory management are performed by the controller. The controller accomplishes these tasks by reading and writing registers in the units using one of the memory access buses (as shown in Figure 10).

If there are two active CPUs, one is normally the controller and the other acts as a "slave" or coprocessor. The master directs I/O and controls the page registers for both CPUs. The slave CPU contains a reduced version of

CPUs	MUs	Memory (KWords)	I/O	Size (W x D x H) (In)	Power (Watts Reg)	Weight (Pounds)
1	4	256	DSIOU	5 x 4.1 x 4.7	12	3.8
2	4	256	DSIOU	5 x 4.5 x 4.7	19	4.2
1	8	512	DSIOU	5 x 6.7 x 4.7	13	5.5
2	8	512	DSIOU	5 x 7.1 x 4.7	20	5.9

SASS-11

Figure 11. Typical 444R^2 Configurations

the executive software and interacts with the controller to obtain I/O and other control services. Additional partitioning of executive functions between CPUs is application specific and tailored to the mission requirements.

Engineering models (form, fit, and function) have been tested and delivered. The 444R^2 was designed to withstand vigorous flight environments at levels shown in Figure 12. The first flight module is scheduled for delivery in January, 1991.

SPA-1 Hardware

SPA-1, a nine module SPA system, was fabricated, tested, and delivered as part of a joint development program between Control Data and Boeing Electronics. SPA-1 provided integrated hardware and software to demonstrate:

* A functional MIN
* Multiple processor capability
* Reliability and reconfiguration
* Data movement across the network and between processing modules
* Fault detection and isolation

Parameter	Design Goal	Program A	Program B
Mechanical			
Temperature	-40°C to +75°C	-22.7°C to +60°C	-34°C to +71°C
Pressure	810 TORR to 10^{-10} TORR	10^{-5} TORR	10^{-5} TORR
Random Vibration	22 g RMS All Axis	15 g RMS Overall	6.5 g RMS + 6db Overall
Pyro Shock	3000 g	N/A	1000 g
Electrical			
EMC*			
Susceptability (Radiated, Elect)		1 V/M 14 KHz to 18 GHz	10 V/M 14 KHz to 30 MHz 5 V/M 30 KHz to 10 GHz 20 V/M 10 KHz to 40 GHz
Conducted Emissions (Broadband)		80 db (uA)/MHz 250 KHz, Rolloff to 40 db, Continuous 2.5 MHz to 400 MHz	80 db (uA)/MHz 15 KHz, Rolloff to 60 db, Continuous 50 MHz to 100 MHz
Radiated Emissions (Elect. Broadband)		20 db (uV)/M/MHz 2.5 MHz to 126 MHz Rolloff to 40 db 1.26 GHz to 18 GHz	120 db (uV)/M/MHz 40 KHz, Rolloff to 75 db 200 MHz, Rolloff to 90 db at 1 GHz

SASS-12   * Sample only, full range of susceptability, conducted and radiation test to be performed.

Figure 12. 444R^2 Environments

The delivered SPA-1 system is shown in Figure 13. A test set, and debugging and software development tools were also delivered as part of this contract.

Only three of the nine nodes in SPA-1 had processors. Two of these processors were smart I/O modules, i.e., I/O units with CPUs, a variation of the basic $444R^2$, and the other a basic $444R^2$. The other six nodes on the MIN were "stubs" that electronically terminated the MIN elements to allow the net to function with no processor attached to the CNUs at any of these nodes.

While the $444R^2$s used on SPA-1 are space qualifiable, no plans have been made to make this configuration of the SPA flightworthy. SPA-1, the MIN in particular, is, and always has been, a proof of concept vehicle that is reasonably close to the function of the first SPA flight hardware, but would be redesigned to miniaturize the hardware even further.

Signal Processing Module (SPM)

Control Data has developed a signal processing module design. The design has evolved to match the interface requirements of the SPA network and to satisfy the functional and performance requirements of anticipated radar, ELINT, and electro-optical missions. The SPM is a 16 bit fixed point/20 bit floating point, two to six stage, pipelined processor capable of over 380 MFLOPS (Million FLoating point Operations Per Second). The size and power constraints are the same as for any processor at an SPA node. The design is based on the same CMOS/SOS technology used in the Control Data/Boeing SPA-1 unit described above. This corresponds to a 1.25 micron feature size, two-layer metal design, providing packing densities of approximately 20,000 gates per chip (less than 1 cm square). Three new VLSI types have been designed: two to gate level (ZYCAD simulated) and one to register level.

The SPM is architecturally flexible in that different numbers of stages can be assembled to meet different mission dependent signal processing requirements. One card, the Signal Processing Section (SPS), has a maximum computation rate of 96 MFLOPS; thus a typical four unit SPM would offer a maximum rate of 384 MFLOPS. The SPM is not a
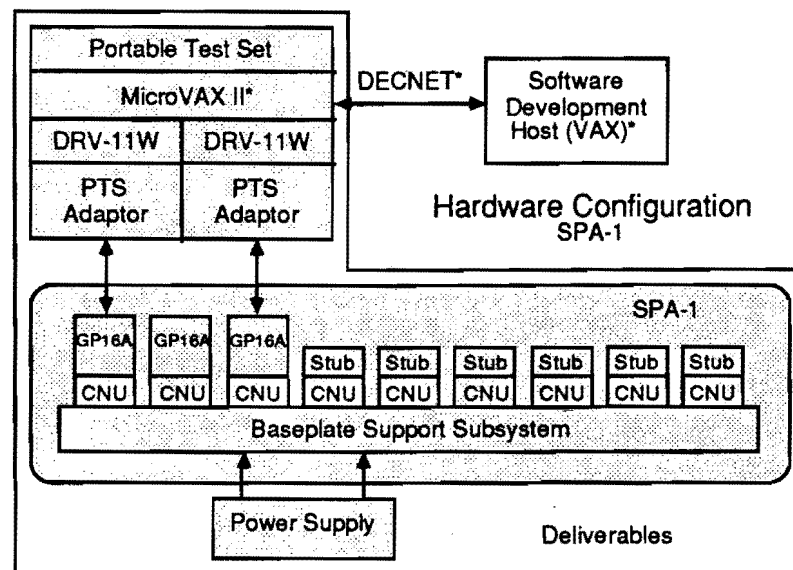
Figure 13. SPA-1 Configuration

hardwired unit; its programmability allows it to support a very wide variety of algorithms as shown in Figure 14. Figure 15 shows fast fourier transform (FFT) execution times as a function of FFT length. Filtering (fourier or FIR) is a frequently used technique in the area of signal processing.
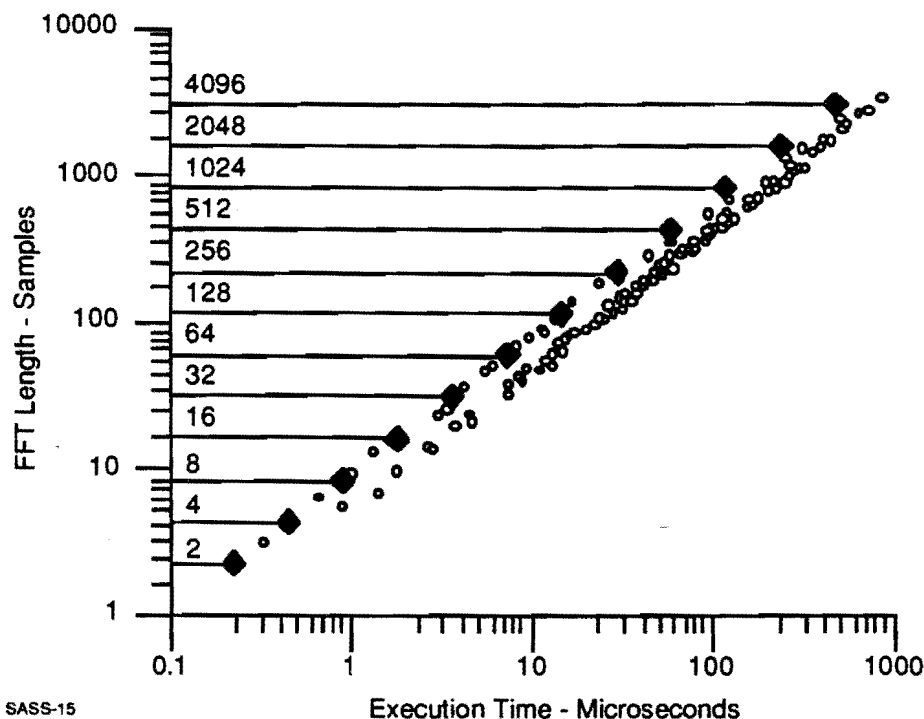
The SPM has fault tolerance features designed into it, including SECDED protection for constant memory and microcode memory. Single bit errors are automatically corrected when a value is read. A rewrite or "scrubbing" cycle can be enabled to avoid long term accumulation of single bit errors. This involves reading a constant on one cycle and writing the scrubbed constant back on the next cycle.

The SPM also has features protecting against synchronization errors. Should a corrupt data block or some undetected (multiple bit) upset cause an out-of-synch condition, the basic SPM array element automatically resynchronizes itself. An end-of-block flag is passed from SPM section to section, and each element in the array unconditionally resynchronizes itself when the flag is received. This guarantees that an out-of-synch condition cannot last longer than two blocks.

SPA Software

An SPA operating system was developed as part of the Control Data/Boeing development

- FFT & FFT $^{-1}$
- Power Spectral Density/Mean/Variance
- Aperture Weighting
- CFAR
- Band Pass Filtering
- Rectangular/Polar Conversion
- Matrix Inversion (Eigenvector Calculation)
- Sliding Window Averaging
- Block/Exponential Integration
- Thresholding
- Beam Splitting (Monopulse)
- Polar Processing
- FIR
- IIR

SASS-14

Figure 14. Signal Processing Unit Typical Algorithms

SASS-15

Figure 15. SPU FFT Execution Times

effort that resulted in the SPA-1 hardware described above. Demonstration software simulating a simple multi-module application was also coded. This software included multi-module processing, data passing around a ring on the MIN, detection and isolation of an induced hardware fault, and reconfiguration after failure.

The operating system, Operational Kernel (OK) was written in Ada, except for those routines that had to read or write hardware registers directly. Such code was written in 1750A assembly language.

Support for the execution of Ada programs is provided by the Ada run time executive (ARTE). The ARTE is included in every active processor module in the SPA and is compatible with the Ada compilation system. The ARTE used was supplied by the compilation system vendor. The OK consists of Ada packages that are built on top of the ARTE. This collection of packages provides operating system services not available from the ARTE, which provides basic services such as tasking, memory management, and exception handling.

The OK demonstrates that cooperative applications executing in multiple modules of the SPA can communicate through the MIN as configured by the users. The OK also provides support for applications in detecting and isolating faults in the SPA as well as in reconfiguring around such detected faults.

The control of the SPA/MIN by the OK is divided into lower level driver services and optional high level communication and control services. The lower level driver services involve the following interface devices:

- Port (network interface unit)
    Message buffers
    Input buffers
    Program control transfers
    Block transfers
    Interrupt handling
- CNUs
    Configuration commands
    Interrupt handling

- Serial I/Os
    Program control transfers
    Block transfers
    Interrupt handling

The higher level services include:

- MINIO - a block I/O facility for the MIN using protocols to ensure the integrity of data transfers between modules.
- HEARTBEAT - a healthcheck facility to support application controlled fault detection and isolation on the MIN.
- CONFIG - a configuration facility that provides a high level set of operations to configure and reconfigure the MIN.

The SPA application environment is shown in Figure 16. The OK/application tasking environment for a multi-module SPA system is shown in Figure 17.

Control Data wrote software for the Portable Test Set (PTS) that was delivered with SPA-1. The PTS software functions include the following:

- Control functions - upon start-up, prompt the user for data about the module(s) under test, and how the PTS itself is to operate, e.g., logging of inputs and outputs.
- Probes - allow the user to monitor and control the module under test by enabling the writing and reading of registers, memory, and the issuing of I/O commands, and the setting and monitoring of discrete lines, e.g., reset.
- Loader - allows the user to load 1750A object files into a $444R^2$ at a node on the MIN.
- Diagnostics - offer many user controllable session options, e.g., control of repetition of tests (number of times, number of errors, and time of day), control of disposition of output (to screen, logfiles, and hardcopy reports).

Control Data modified a standard vendor 1750A debugger/instruction level simulator to handle certain options and implementation specific characteristics of the $444R^2$'s 1750A hardware. The modified simulator also
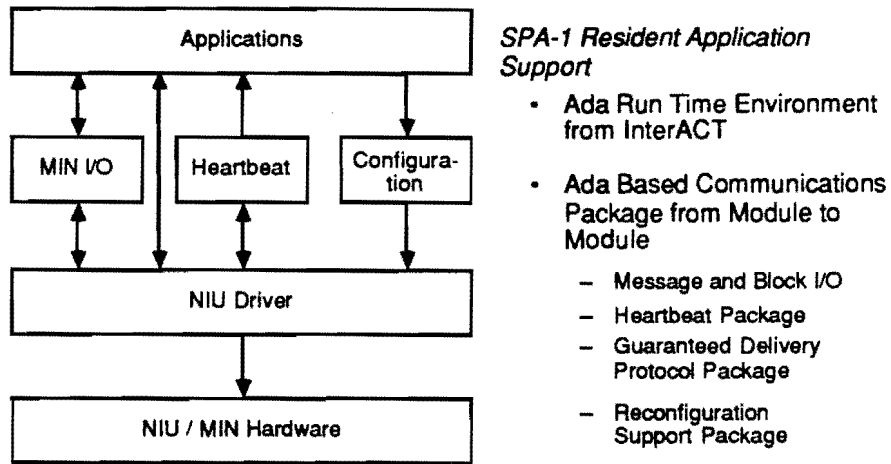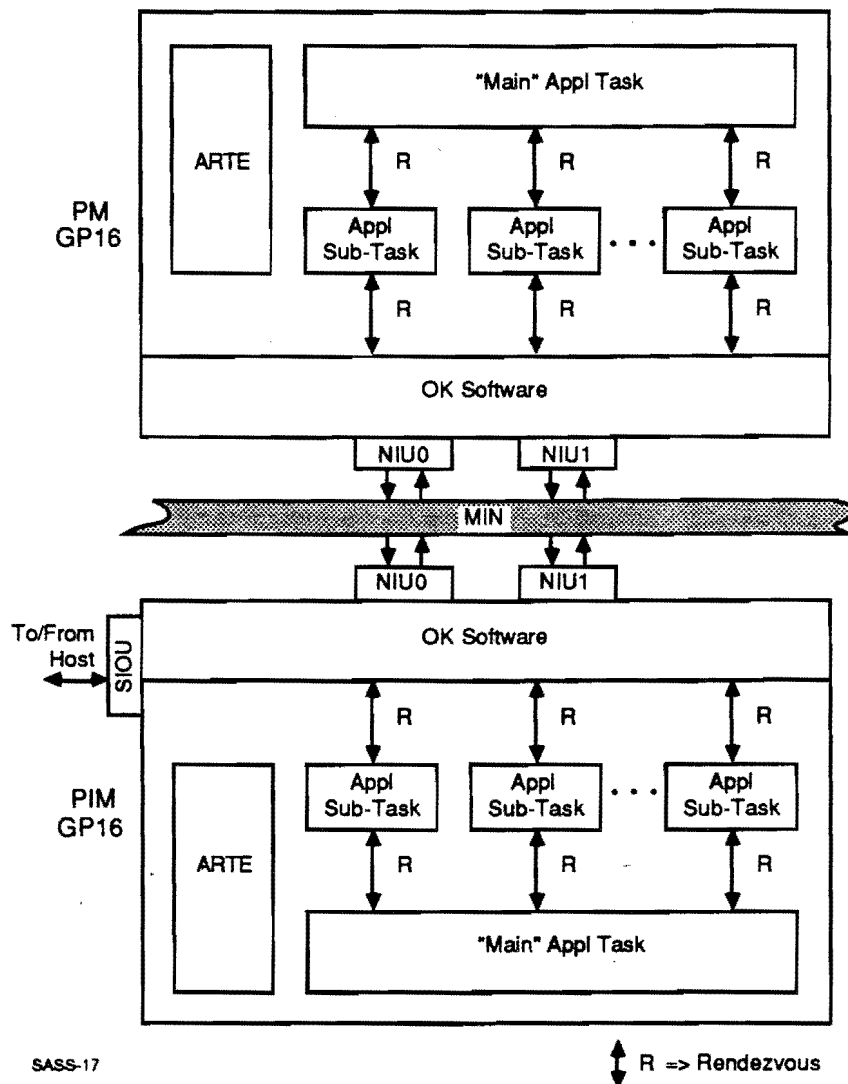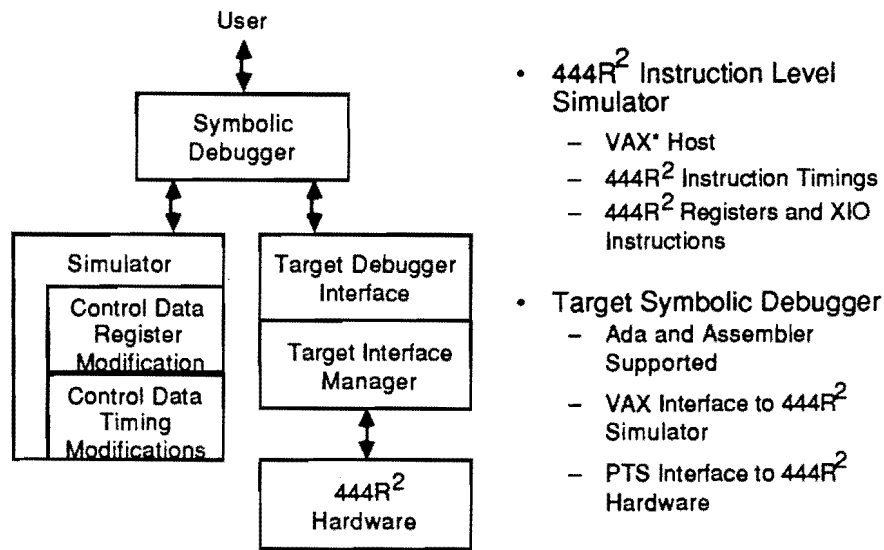
Figure 16. SPA Application Environment

Figure 17. OK/Application Multi-Module Tasking Environment

enables the debugging of processor modules within an SPA. The debugging environment is shown in Figure 18. Figure 19 shows the general software development support system.

- **444R^2 Instruction Level Simulator**
  - VAX* Host
  - 444R^2 Instruction Timings
  - 444R^2 Registers and XIO Instructions

- **Target Symbolic Debugger**
  - Ada and Assembler Supported
  - VAX Interface to 444R^2 Simulator
  - PTS Interface to 444R^2 Hardware

* VAX Registered Trademarks of Digital Equipment Corporation.

SASS-18

Figure 18. SPA/PTS Symbolic Debugging Environment
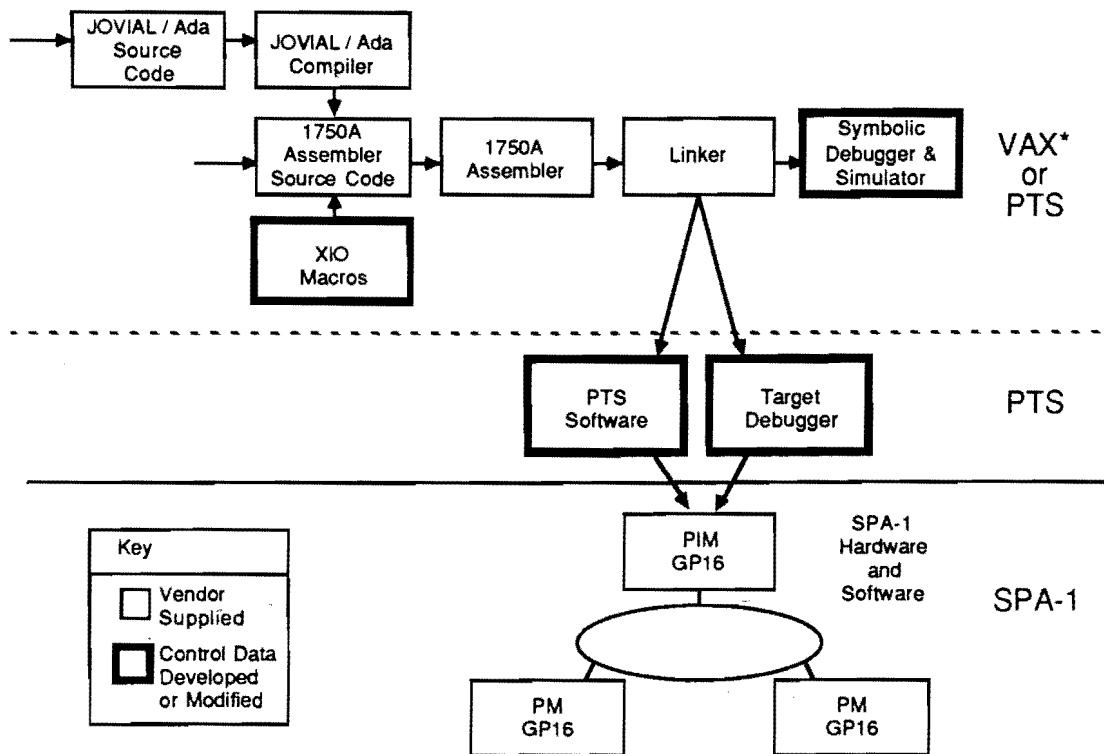
* VAX Registered Trademarks of Digital Equipment Corporation.

SASS-19

Figure 19. Software Development Support System

Future SMARTSAT Development

There are several areas of development in which Control Data is working to improve the SMARTSAT processing concept.

New Processors

New processor types will need to be developed to exploit Artificial Intelligence (AI) constructs and other specialized processing techniques fully. These new devices must be designed with the special environments of orbital operation as primary considerations (especially the radiation environment).

Existing Processors

Some existing processors, depending on "designed for space" suitability, can be interfaced to the configurable MIN to exploit prior development and to solve special problems.

Improved Packaging

New packaging techniques can be exploited to reduce the size, weight, and power of the processing net even further. While already low, these characteristics can be significantly reduced by the judicious use of today's packaging technology. An example is the packaging of two 1750A processors in a single two-inch square hybrid package. This results in a volume reduction factor of six and a weight reduction factor of 20, compared to the 444R[2]. The project was a cooperative effort between MCC and Control Data as a proof of concept for Tape Automated Bonding (TAB) techniques for a direct die to substrate packaging approach (Figure 20).

Disk Storage

Significant on-board processing produces a great need for on orbit non-volatile mass
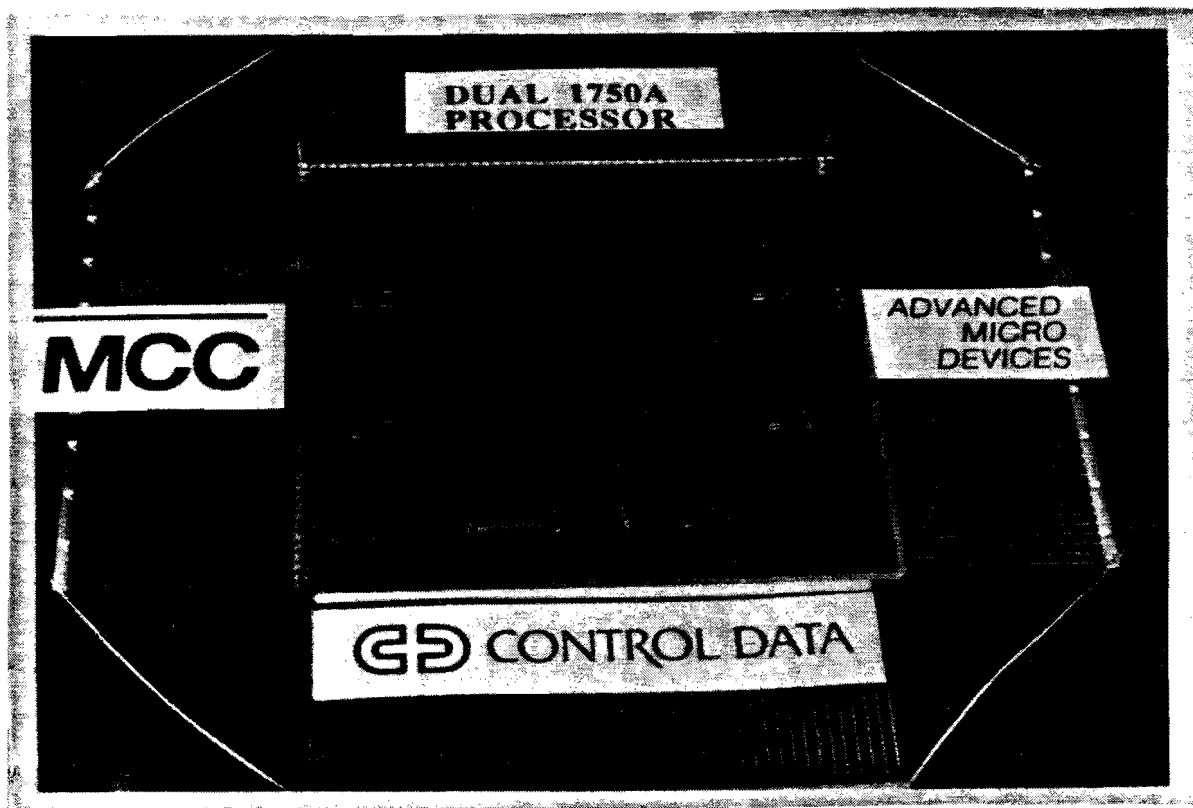

Figure 20. Dual 1750A Processor

storage of data and software. Development of a "space hardened" magnetic or rewritable optical disk drive for inclusion in the processing net will be relatively simple. Present designs exist for militarized disks which withstand shock and vibration, but need to be modified for radiation and zero-g.

Inductive Power Connect

NASA has demonstrated an inductive power interconnect across a .015 in. air gap that would provide power transfer into a module with no physical electrical connection. This provides the possibility of complete electrical isolation of a module and elimination of connectors (with all their reliability/maintainability problems).

Optical Data Interconnect

NASA has also demonstrated communication over a small air gap (.015 in.) of data at 50 MBIT rates using LEDs. This technique would allow a module to be merely placed in the correct position to provide complete functionality with no electrical connections.

Total System Integration

To properly use a processing net as described above, individual technologies that comprise the SMARTSAT must be integrated to a higher degree than has ever been accomplished before in a space system. This will require close cooperation of the integrator and all subsystem contractors to eliminate unnecessary hardware while maintaining full functionality.

This approach is a dramatic departure from the conventional space hardware program wherein each function is isolated with defined interfaces. The conventional approach leads to large segmented systems with significant interface hardware. In addition, the traditional approach is to ship all data to the ground for exhaustive analysis by a supporting team of experts and analysts. This approach has been successful, but very expensive, and makes timely distribution of results difficult.

Another area of potential cost savings lies in the usage of limited life satellites. Since low earth orbits decay relatively rapidly, the expected lifetime of the SMARTSAT would be comparatively short. Because of the low cost boosters available and the anticipated low cost of the satellites, potential cost savings abound in the "Hi-Rel" area. One of the principal cost drivers in conventional space programs is the extreme reliability measures. This is justified for large systems because of the expense of launch and replacement.

With the advent of small satellites, Control Data can address these costs and take calculated risks reliability-wise. One of the major tasks associated with the SMARTSAT efforts will be to establish which "Hi-Rel" processes can be modified while still providing high mission probability of success. In addition, properly designed fault tolerance and isolation features further reduce the need for extremely high reliability (expensive) components to achieve mission goals.

Control Data's SMARTSAT architecture will satisfy its users' needs and provide mechanisms for high reliability. When the traditional approach is contrasted with the real-time information availability of the SMARTSAT, the desirability of small, capable, intelligent satellite systems is obvious.

GLOSSARY

Ada	A DOD programming language
ADAC	Attitude Determination And Control
AI	Artificial Intelligence
ARTE	Ada Run Time Executive
CNU	Configurable Network Unit
CONUS	Continental United States
CMOS/SOS	Complementary Metal Oxide Semiconductor/Silicon On Sapphire
CPU	Central Processing Unit
DAIS	Name for a particular instruction mix
DMA	Direct Memory Access
DOD	Department of Defense
ELINT	Electronic Intelligence
FDIAR	Fault Detection, Isolation, Assessment, and Recovery
FFT	Fast Fourier Transform
FIR	Finite Impulse Response
FMEA	Failure Modes and Effects Analysis
IFF	Identification Friend or Foe
IOU	Input/Output Unit
IR	InfraRed
ISA	Instruction Set Architecture
LED	Light Emitting Diode
LEO	Low Earth Orbit
MBIT	Megabit (of memory) (1 million bits)
MCC	Microelectronic Computer Consortium
MFLOPS	Million FLoating point Operations Per Second
MIN	Module Interconnect Network
MIPS	Million Instructions Per Second
MU	Memory Unit
OK	Operational Kernel
PTS	Portable Test Set
RAM	Random Access Memory
RISC	Reduced Instruction Set Computer
SAR	Synthetic Aperture Radar
SDI	Strategic Defense Initiative
SECDED	Single Error Correction, Double Error Detection
SMARTSAT	SMART Satellite
SPA	Space Processing Array
SPM	Signal Processing Module
SPS	Signal Processing Section
SPU	Signal Processing Unit
UV	UltraViolet
VLSI	Very Large Scale Integrated (Circuitry)