AIAA 79-1948R

Fault-Tolerant Communications Processor

J. J. Stiffler*
Raytheon Company, Sudbury, Mass.

An architecture for a fault-tolerant spaceborne communications processor is described. The processor is designed to demodulate, deinterleave, and decode low-rate data received over a number of channels and to encode and interleave an appropriately edited version of the received data for subsequent retransmission. Emphasis is placed on the combination of hardware and software techniques used to reconfigure the processor in order to circumvent its malfunctioning elements. The reliability of the fault-tolerant processor is evaluated and compared to that of a non-fault-tolerant, but otherwise equivalent, processor.

Introduction

THE Wide-Band Signal Processor (WBSP) is a fault-tolerant spaceborne communications processor designed to survive in space for long periods of time. It is specifically configured to perform the various onboard processing tasks (demodulation, decoding, deinterleaving, interleaving, encoding, formatting, and routing) needed to regenerate multichannel, low-data-rate (e.g., 75 or 2400 baud) signals. It is intended to operate as a peripheral to the Fault-Tolerant Spaceborne Computer (FTSC) currently under development. While the WBSP handles all the basic signal processing functions, the FTSC acts as master control for the entire satellite system. In this role, the FTSC performs spacecraft-housekeeping functions (navigation, attitude control, etc.), while ensuring the health of the entire system, including that of the WBSP. It is this latter aspect of the FTSC/WBSP system, the techniques by which the FTSC monitors the operational status of the WBSP and effects needed repairs, that is of primary concern here. In order to put the ensuing discussion in its proper context, however, it is necessary to describe briefly the architecture of the WBSP itself. (A more detailed description can be found in Ref. 1.)

WBSP Architecture

The WBSP architecture is highly modular and, as a consequence, can be tailored to the signal-processing requirements specific to the application of interest. Both the number and the types of processing elements can be adjusted, over a wide latitude, to accommodate the number of communication channels, demodulator requirements, and interleaver/deinterleaver and encoder/decoder algorithms postulated for a given situation.

To illustrate some of the available options, two WBSP configurations are described. Both are designed to process 25 frequency-multiplexed, 75-baud, 8-ary frequency-shift-keyed (FSK) channels and two time-multiplexed, 3750-baud quadrature-phase-shift-keyed (QPSK) channels. Their differences are due to different constraints placed on the synchronism of the arriving signal: In the "synchronous" case, all channels are required to arrive at the satellite at the correct frequency and with correct timing (at the chip level). This requires that the signals be precorrected in both time and frequency at the transmitter. In the "asynchronous" case, this constraint is removed from 15 of the FSK channels. Thus, in this latter case, the WBSP must itself establish the necessary time and frequency synchronism in order to regenerate these signals. The details concerning the various demodulation, interleaving/deinterleaving and encoding/decoding operations to be performed are omitted here for the sake of brevity. These details are, in any case, largely irrelevant to the ensuing discussion.

The asynchronous WBSP configuration is shown in Fig. 1. Each of the five buses provides a data path for one or two 8-bit bytes, one or two parity bits, and includes a spare section that can be used to replace any defective data or parity segment. There is no address bus; each element accepts and supplies information at prescheduled times. The data buffer in each element's bus interface contains a timing circuit that identifies the time slots during which the element is to send or receive data. Counters in each data buffer increment with each bus clock pulse. Two pairs of address and mask registers identify all counts at which the element is to transmit or receive information. These registers are programmed by the FTSC during system configuration.

Three status/control lines are used to coordinate system initialization and reconfiguration. A "reconfiguration" line distinguishes between normal and reconfiguration modes, a "clock enable" line is used to synchronize data transfers during reconfiguration, and a "hard/soft" line indicates whether the data on a given bus are to be interpreted as a hard address or as ordinary data. Six other status/control lines identify the bus- and timing-module configurations.

The asynchronous WBSP comprises nine types of elements. Dedicated spares are used for both the analog-to-digital converters and the modulators, as these are hardwired to analog inputs or outputs. The bus controller, power supply, and timing generator are also provided with dedicated spares since each of these elements is unique. All other elements, the stage-1 demodulator, stage-2 demodulator, interleaver/deinterleaver, and encoder/decoder processors use pooled spares. It should be noted that pooled spares can be used for each of these processor types even though different processors within a given group may be required to execute different algorithms. This is possible because the control storage, which specifies the algorithm a processor is executing, is implemented with random-access memories (RAMs) rather than read-only memories (ROMs). This enables the FTSC to change or reassign algorithms when necessary to maintain the functional integrity of the system.

The bus controller provides the interface between the FTSC and the WBSP. During normal operation, data are typically routed from the analog-to-digital converters to the stage-1 demodulators where the asynchronous frequency-multiplexed channels are demultiplexed and the synchronous channels demodulated. The soft-decision data produced by the demodulators are routed to the deinterleaver processors, the deinterleaved data are passed on to the decoder processor, and the decoded data are transferred through the bus controller to

Presented as Paper 79-1948 at the AIAA/IEEE/ACM/NASA Computers in Aerospace II Conference, Los Angeles, Calif., Oct. 22-24, 1979; submitted Dec. 18, 1979; revision received Feb. 17, 1981.

*Consulting Engineer.

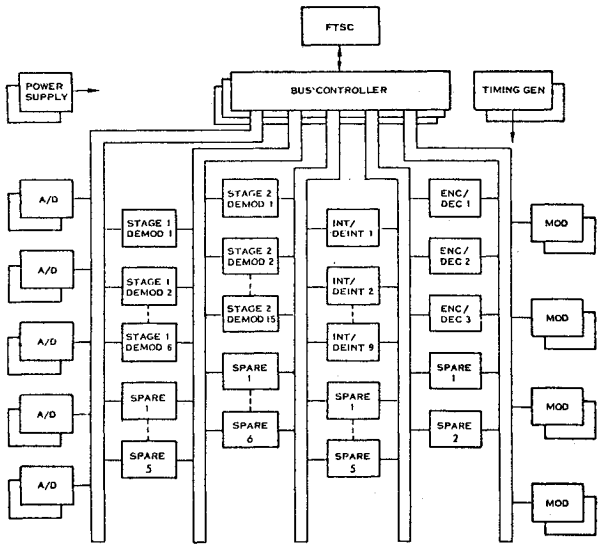

Fig. 1 Wide-band signal processor (asynchronous case).

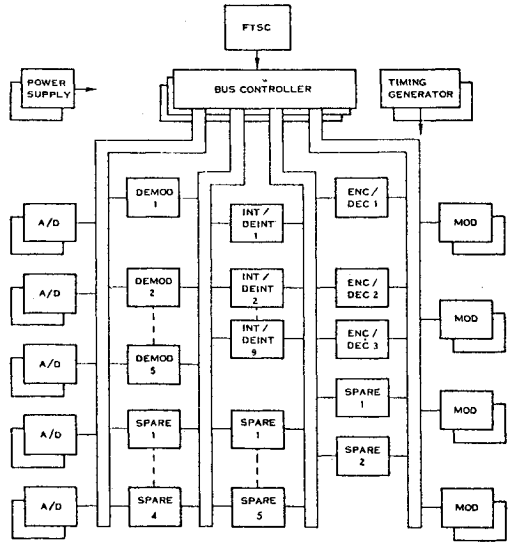

Fig. 2 Wide-band signal processor (synchronous case).

the FTSC's main memory. The FTSC then checks the integrity of the data, eliminates duplicate messages, and routes the resulting data back through the bus controller to an interleaver processor for encoding and interleaving. The interleaver processor output is routed through an encoder (where a cover sequence may be added) and from there to a modulator for retransmission.

The synchronous WBSP configuration shown in Fig. 2 has only one demodulator stage and four data buses rather than the five required in the asynchronous case. In all other respects, it is similar to the asynchronous WBSP just described.

Fault Monitoring

As noted earlier, one of the functions of the FTSC is to monitor the performance of the WBSP in order to detect failures in any of its processors or buses, and to effect the necessary reconfigurations (deactivation of the faulty processor or bus segment and activation of a spare) following any such detected failure. The first step in this procedure, of course, is to detect any anomalous behavior in the WBSP and to isolate the source of the problem to a replaceable unit.

A number of fault detection methods are compatible with the WBSP architecture. The most appropriate method or combination of methods depends on several factors such as the details of the message format, the amount of time allowed for fault detection, and the particular algorithms being executed. For purposes of this discussion, the different monitoring techniques will be separated into two categories: passive monitoring in which the FTSC monitors the WBSP's performance by testing the decoded information buffered in the FTSC's main memory; and active monitoring in which additional WBSP processors are activated solely for monitoring purposes.

Passive monitoring is effective only if the data buffered in the FTSC's main memory contain some redundancy. If the data were totally irredundant, all data patterns would be equally likely. A malfunction would be detectable only if it resulted in patterns significantly different statistically from those anticipated, such as an unlikely imbalance between "ones" and "zeros." In many applications, each frame of data received over each channel is terminated with a "tail" consisting, for example, of the modulo-n sum of all data words comprising that frame. If this is the case, a simple "tail check" in the FTSC can reveal a great deal about the WBSP's performance, particularly when the details of the signal flow through the WBSP are taken into account. The purpose of interleaving and encoding the data is to make transmission errors extremely unlikely. To the extent that this procedure is successful, a tail-check failure (i.e., a situation in which the received tail does not satisfy the requisite conditions) is highly unlikely. Therefore, tail-check failures occurring more frequently than some threshold rate are indicative of some hardware malfunction; the pattern of these failures can in many cases be used to locate the sources of that malfunction.

In particular, consider the two WBSP configurations described previously. Each asynchronous first-stage demodulator processes signals representing some subset of the total set of received channels. Since a given processor operates on a sizable fraction of the chips (waveform segments) comprising a frame of data (in some cases, every chip; in other cases, every other or every third chip), a malfunctioning processor will most likely cause tail-check failures in each of its associated channels. In contrast, the second-stage asynchronous demodulators are each dedicated to a single channel. Consequently, failures in these processors will affect only that channel. Deinterleaver processors, like the first-stage asynchronous demodulator and the synchronous demodulator processors, operate on subsets of more than one channel. It is not difficult to make these two sets of subsets disjoint, however. In most cases, in fact, throughput constraints of the different processors force these sets to be disjoint. If this is done, no single deinterleaver failure can affect the same subset of channels affected by a demodulator failure, and the converse is also true. Similarly, if the subset of channels processed by any single decoder is distinct from the subset of channels operated on by any other single processor, a malfunctioning decoder also produces a unique pattern of tail-check failures. The FTSC need only compare the observed failure pattern to a stored set of failure syndromes to determine not only the fact of a hardware failure but also its location.

It must be acknowledged that the discussion in the preceding paragraph has been overly simplified. Further diagnosis may be needed, for example, to distinguish between a faulty A/D converter, an unused channel, and a malfunctioning second-stage demodulator. Moreover, it is entirely possible for a processor to fail in such a way that its performance deteriorates only slightly (e.g., one of the less significant adders in an arithmetic unit fails). In this event, the tail-check failure rate may increase but not to the point at which the FTSC would attempt to identify and replace the faulty element. As a result, the level of performance over some subset of channels might be needlessly reduced. And, although the interleavers and encoders constitute only a small

fraction of the entire WBSP, they too can fail. Obviously, failures in these elements cannot be detected through tail checks.

All of these problems can be mitigated by programming each processor to transmit periodically to the FTSC certain easily calculated performance measures such as the bit-error rate as perceived by a decoder, or the average difference between the demodulated symbol metric and that of its nearest competitor as determined by a demodulator.

The major objection to the passive-monitoring technique concerns the amount of time needed to detect and isolate a faulty element. Indeed, from 10 to 100 frames of data might well be required to isolate some of the more subtle malfunctions. This delay may in some cases be unacceptable. (It should be noted, however, that any malfunction resulting in serious performance degradation can generally be isolated relatively quickly.)

More rapid fault isolation is possible if the WBSP is provided with active monitors; that is, if additional processors are activated on each bus and programmed to monitor the performance of the other processors on the same bus. The cost of doing this is one additional active processor per bus (plus a possible additional spare to preserve the overall system reliability). The monitor processor need only duplicate the functions of the various active processors on a time-shared basis and generate an interrupt to the FTSC whenever it detects an output from the monitored processor with which it disagrees. In response to the interrupt, the FTSC queries the monitor concerning the identity of the indicated processor, and, after further diagnostics of the sort described in the passive-monitor case, determines whether the failure resides in that processor or in the monitor itself. The monitor processor, incidentally, can often simultaneously monitor all active processors by, for example, monitoring, at any given time, only some of the channels being processed by each processor, or by executing different, less-demanding algorithms by, for example, encoding a decoder's output and comparing the result with its input.

As the preceding discussion suggests, very little circuitry has been added to the WBSP solely for fault-monitoring purposes. The main reason for providing fault-monitoring hardware is to decrease the time delay between the occurrence of a fault and its detection. This time reduction can be crucial in systems like the FTSC in which a malfunctioning computer can jeopardize an entire mission (e.g., by erroneously activating a control jet). It is not so important, however, in the WBSP where the effect of a malfunctioning processor is a garbled communication that, if necessary, can be repeated. The mean-time between component failures in the WBSP should be on the order of a few months. Thus, even if several seconds elapse between the occurrence of a fault and its detection and repair, the resulting several seconds' outage every few months does not represent a serious deterioration in performance. The tail-checking procedure described above is generally satisfactory and when supplemented with cross monitoring is certainly satisfactory under these conditions, so the cost of more extensive fault-monitoring hardware is not justified. (The modest amount of fault-monitoring circuitry that has been designed into the WBSP has been included primarily to aid in the fault-isolation process. The previously mentioned byte parity used over the buses, for example, makes it easier to distinguish between failures in the originator of a message, typically a processing element, and the bus over which this message is transmitted.)

Fault Recovery

Once a fault is detected and isolated, the recovery process is relatively straightforward. No attempt is made to preserve communication capability during the recovery process. The FTSC reconfigures the WBSP through the bus controller. To initiate this operation, the FTSC places the WBSP in a "hard address" mode; i.e., it sets the appropriate hard/soft control lines to the hard states. To reach a given processor, the FTSC transmits that processor's hard address, along with an appropriate two-bit command, over each of the two buses with which the processor of interest communicates. Any processor can be deactivated if it can accept information over either of its associated buses; it can be activated and armed for subsequent data initialization only if it can accept information over both buses.

When a fault is detected either by the FTSC or by the fault monitors, the FTSC takes direct control over the WBSP. It then can test any specified module, deactivate it if necessary, activate a spare module, and program it (by loading its control memory with the appropriate microinstructions) to take over the function vacated by the failed module. The FTSC then reinitializes the WBSP and normal operation resumes.

The most crucial constraint on the WBSP hardware, so far as fault isolation and recovery are concerned, is the requirement that the effects of a failure be properly contained. If, for example, a processor could fail in such a way that it constantly transmitted data over one of its associated buses, the identity of that processor might be difficult to establish. If there were a failure mode in which it could constantly transmit over both buses, it might not be possible to deactive the faulty processor even if it could be identified.

To avoid such difficulties, extensive use was made of the fault-containment techniques developed for the FTSC. In particular, the processor-bus interfaces are implemented with the same LSI devices used to implement the FTSC's data and address-bus interfaces. These devices include circuitry to generate and monitor bus parity, to send interrupts to the FTSC in the event of detected parity violations, and, on the basis of control signals relayed through the bus controller, to replace a defective byte with a spare byte should the FTSC determine that the failure was in fact in a bus rather than in a processor. In addition, the circuitry is so designed that no single failure can allow a processor to transmit data at other than during its assigned time slots, thus virtually eliminating failure modes of the sort mentioned in the previous paragraph.

System Reliability

The predicted 10-yr reliability of the asynchronous and synchronous versions of the WBSP are summarized in Tables 1 and 2, respectively. The hazard rates of the individual devices used to implement the WBSP were obtained from MIL-HDBK-217C (Ref. 2); class S parts and a 25°C ambient temperature were assumed throughout.

The hazard rates (column 2 in Tables 1 and 2) associated with the bus controller, decoders, A/D converters, modulators, power supplies, and timing units are just the sum of the hazard rates of their respective components.

The bus hazard rates represent the rate of occurrence of failures that disable an entire byte-width section of a bus. As already noted, the WBSP bus interfaces are implemented with FTSC LSI devices. These devices are segregated into 8-bit-wide sections; the sections are physically isolated from each other in order to keep the failures in the separate sections uncorrelated. A failure in any one section may or may not prevent other modules attached to that bus from using the corresponding segment of that bus. If a tri-state driver fails in the "on" state, for example, the entire bus segment could be disabled; if it fails in the high-impedance state, the segment can still be used by the other modules on that bus. The hazard rate for bus-disabling failures was estimated as $n\gamma a\lambda_b$, with n the number of interfaces attached to the bus in question, γ the fraction of bus interface failures that fall into the bus-disabling category, a the fraction of the LSI device area used to implement all the drivers and receivers associated with each section, and λ_b the hazard rate of the entire device. The

Table 1 WBSP 10-yr reliability asynchronous demodulation

	WBSP			Irredundant equivalent	
	Hazard rate ($\times 10^{-7}$/h)	Number available/ number required	Reliability (10 yr)	Hazard rate ($\times 10^{-7}$/h)	Reliability (10 yr)
Bus controller	36.39	3/1	0.9797	20.97	0.8322
Bus 1	1.16	4/3	0.9994	0.58 ($\times 2$)	0.9899
Bus 2	1.69	4/3	0.9987	1.06 ($\times 2$)	0.9816
Bus 3	1.83	3/2	0.9992	1.20	0.9895
Bus 4	1.06	3/2	0.9997	0.63	0.9941
Bus 5	0.79	3/2	0.9999	0.39	0.9966
Bus 6	0.86	3/2	0.9998	–	–
Demodulator	20.03	11/6	0.9967	18.25 ($\times 6$)	0.3832
2nd stage	14.79	21/15	0.9994	21.93 ($\times 15$)	0.0560
Deinterleaver	14.98	14/9	0.9957	20.80 ($\times 9$)	0.1940
Decoder	11.95	5/3	0.9916	8.33 ($\times 3$)	0.8034
A/D converter	7.36	$(2/1)^5$	0.9807	3.12 ($\times 5$)	0.8723
Modulator	7.46	$(2/1)^4$	0.9840	3.22 ($\times 4$)	0.8933
Power supply	13.80	2/1	0.9870	12.50	0.8963
Timing unit	6.30	2/1	0.9971	5.30	0.9546
Total reliability			0.9039		0.0018

Table 2 WBSP 10-yr reliability synchronous demodulation

	WBSP			Irredundant equivalent	
	Hazard rate ($\times 10^{-7}$/h)	Number available/ number required	Reliability (10 yr)	Hazard rate ($\times 10^{-7}$/h)	Reliability (10 yr)
Bus controller	32.29	3/1	0.9850	20.97	0.8322
Bus 1	0.91	4/3	0.9996	0.48 ($\times 2$)	0.9916
Bus 2	1.11	3/2	0.9997	0.67	0.9941
Bus 3	0.91	3/2	0.9998	0.58	0.9949
Bus 4	0.71	3/2	0.9999	0.38	0.9967
Bus 5	0.86	3/2	0.9998	–	–
Demodulator	20.03	9/5	0.9923	18.25 ($\times 5$)	0.4496
Deinterleaver	14.98	14/9	0.9957	20.80 ($\times 9$)	0.1940
Decoder	11.95	5/3	0.9916	8.33 ($\times 3$)	0.8034
A/D converter	7.36	$(2/1)^5$	0.9807	3.12 ($\times 5$)	0.8723
Modulator	7.46	$(2/1)^4$	0.9840	3.22 ($\times 4$)	0.8933
Power supply	13.80	2/1	0.9870	12.50	0.8963
Timing unit	6.30	2/1	0.9971	5.30	0.9546
Total reliability			0.9153		0.0380

fraction γ was somewhat arbitrarily but conservatively equated to ½; this parameter will be adjusted in future reliability estimations as better data become available.

The hazard rates for the demodulator and the interleaver/deinterleaver processing elements are less than the sum of the hazard rates of their various components because these elements contain internally redundant memories. Specifically, a spare bit line is appended to these memories along with a switching device by which this spare can be switched in to replace any other line found to be defective. One or two parity-bit lines, depending on the width of the specific memory, are also added to each memory in order to detect bit-line oriented failures. Once a failure is detected, the processing element can, upon FTSC command, diagnose the failure to determine whether it was transient or permanent. In the latter case, it can also, in conjunction with the switching device, isolate the faulty bit line and switch in the spare. (This technique and the hardware to implement it were developed in the FTSC program.) For small memories, the increase in memory reliability obtained in this way is largely offset by the unreliability of the switching hardware, so the added complexity is not justified. For the larger sample-point and deinterleaver data memories, however, the reliability gain is substantial. As a result, the effective hazard rates listed in Tables 1 and 2 for the demodulator and interleaver/deinterleaver processors are considerably smaller than they would have been had their memories not been made internally redundant.

The third column in Tables 1 and 2 indicate the number of functioning elements of a given kind that are initially available in the WBSP and the number that are needed for nondegraded performance. Entries of the form $(2/1)^l$ indicate that l pairs of dual-redundant elements are available, with at least one of each pair required for nondegraded performance. It will be noted that three bus segments are needed to implement a 16-bit-wide bus and its associated control lines and two bus segments are needed for a 8-bit-wide bus. This is because the WBSP is implemented with FTSC parts that are segmented into 8-bit-wide sections. An additional segment is therefore required to accommodate the parity bits associated with each bus byte. It will also be noted that both the asynchronous and synchronous WBSPs contain a bus not explicitly shown on their respective block diagrams. This bus is not used for normal data transfer. It is needed, however, to provide a second configuration-control input to the A/D converters and the modulators and must be accounted for in the reliability assessment.

The fourth column in Tables 1 and 2 show the probabilities that the minimum number of elements in each row survives for a period of at least 10 yr. The products of these probabilities, the 10-yr reliabilities of the two systems, are listed at the bottom of these columns.

The last two columns in Tables 1 and 2 list the hazard rates and reliabilities of the various elements of irredundant versions of the asynchronous and synchronous WBSPs. All fault-tolerant hardware (parity generators and checkers, spare

bit-lines and switches, interface protection circuitry, spare bus segments, etc.) has been removed from each of these WBSP elements, and the respective systems are implemented using only the minimum number of elements needed for non-degraded operation. The hazard rates for the individual elements have in most cases decreased, often considerably, due to the elimination of this extra hardware. The first-stage demodulator hazard rate decreases only slightly, however, and the second-stage demodulator and the interleaver/deinterleaver hazard rates actually increased. This is due to the fact that the increased reliability of the internally redundant memories in the fault-tolerant elements compensate for the increased complexity of their fault-monitoring hardware and their fault-tolerant interfaces.

As is apparent in Tables 1 and 2, the redundancy techniques used in the WBSPs dramatically increases the reliability of these systems relative to that of their irredundant equivalents. Both the asynchronous and the synchronous WBSPs have a better than 90% chance of surviving 10 yr without degradation; their irredundant equivalents have a less than 4% (less than 0.2% in the asynchronous case) chance of surviving that long.

A measure of the cost of achieving the increased reliability can be obtained by comparing the number of logic and memory devices needed to implement the various WBSP systems. In particular, the total number of those devices needed for the processing elements and the bus controller are: asynchronous, fault-tolerant WBSP, 5062; asynchronous, irredundant WBSP, 2675; synchronous, fault-tolerant WBSP, 2390; synchronous, irredundant WBSP, 1123. Thus, the complexity of this portion of the fault-tolerant WBSP is roughly twice that of its irredundant equivalent in both the asynchronous and the synchronous cases. Since the remainder of the fault-tolerant WBSPs consist entirely of dual-redundant components, this 2 to 1 complexity ratio in fact applies for the entire system.

Concluding Remarks

The wide-band signal processor architecture appears to provide a versatile structure for a wide class of communication-signal processing functions. The WBSP's tolerance to faults is achieved through the provision of a pool of spares, each of which is capable of taking over the function of any module communicating over the same pair of buses. The task of monitoring the performance of each of the active WBSP processors, in order to determine whether or not it needs to be replaced, is not as formidable as it might at first appear. Either, or a combination, of two techniques—passive monitoring of the decoded data buffered in the FTSC's main memory and active monitoring through the use of monitor processors—can be used effectively for this purpose. That such fault-tolerant structures can be effective in increasing the reliability of complex systems is readily demonstrated by the fact that the 10-yr survival probability of the WBSP can be increased by a factor of more than 20 at the cost of a roughly twofold increase in complexity.

Acknowledgment

This work was supported in part by U.S. Air Force under Contract number F04701-77-C050.

References

[1] Stiffler, J.J., "The Wide-Band Signal Processor," *Proceedings of the International Telemetering Conference,* Los Angeles, Nov. 1978, pp. 73-77.

[2] MIL-HDBK-217C, Rome Air Development Center, Griffiss Air Force Base, N.Y., May 1, 1980.