THE WIDE-BAND SIGNAL PROCESSOR

J. J. Stiffler
Consulting Scientist
Raytheon Company
528 Boston Post Road
Sudbury, Massachusetts 01776

ABSTRACT

The Wide-Band Signal Processor (WBSP) is a spaceborne communications processor designed to operate as a peripheral to the Fault-Tolerant Spaceborne Computer (FTSC) currently being developed for the U. S. Air Force. Its function is to demodulate and decode received FDM and TDM signals and to re-encode the recovered information and use it to modulate signals for retransmission. The major difference between the WBSP and other processors designed to perform similar functions is in the fact that the WBSP, like the FTSC itself, is designed to survive its own hardware malfunctions.

INTRODUCTION

The Fault-Tolerant Spaceborne Computer (1) is designed to operate in space for periods of five years or more without degradation in performance. It was sized to handle typical on board processing tasks: navigation, guidance, attitude control, thermal control, solar panel control, telemetry, command decoding, etc. It was recognized at the outset, however, that certain specialized, complex functions should be performed by peripheral hardware specifically designed for that purpose. Accordingly, the FTSC's input/output specifications were defined so as to accommodate the additional data load anticipated for these dedicated processors.

If the function to be implemented by a peripheral processor is critical to the success of the mission, such a processor must evidently achieve a level of reliability comparable to that of the FTSC itself. The fact that the FTSC is on board the spacecraft and capable of communicating with the peripheral processor considerably simplifies the task, however, since the properly functioning computer can be relied upon to diagnose the peripheral's malfunctions and to coordinate its subsequent recovery. Nevertheless, great care must be exercised in designing the processor if the desired reliability is to be achieved without incurring an excessive penalty in redundant hardware.

This work was supported by U. S. Air Force (SAMSO) contract F04701-77-C050.

One such high-reliability peripheral to the FTSC, a communications processor called the Wide-Band Signal Processor (WBSP), is the subject of this paper. The WBSP is requird to perform the following tasks: (1) Demultiplex and demodulate FDM and TDM signals received over five different frequency bands. (The total number of channels involved is of the order of 100 or fewer, with each channel operating at a data rate of either 75 bps or 2400 bps.) (2) Deinterleave the demodulated digital data stream. (3) Decode the deinterleaved data. (4) Encode and interleave information intended for retransmission and use it to modulate the five downlink (or crosslink) carriers. The received information also has to be formatted and routed to the proper channels for downlink transmission. This task is more readily performed in the FTSC itself, however, so it is not identified as a WBSP task.

The system configuration selected to implement these tasks and the throughput and reliability considerations leading to that configuration are described in the next section.

WBSP SYSTEM CONFIGURATION

One of the most important lessons learned during the early phases of the FTSC study concerned the dramatic improvement in reliability that could be achieved by pooling spares. A spare module capable of replacing any one of, say, five active modules in the event of a failure improves the system reliability nearly as much as if each of the five modules had a dedicated spare; yet the latter configuration obviously requires five times as much redundant hardware. Two conditions must be satisfied, however, in order to exploit this potential

- The switching device used to isolate the faulty module and to replace it with one of the pooled spares must be so designed that its unreliability does not dissipate much of the potential of this approach.

- The processor must be partitioned to the largest extent possible into interchangeable modules so that spares can in fact be pooled.

The first of these requirements was addressed extensively in the FTSC design and the results of this effort were adapted with only minor modifications for the WBSP. The second requirement was relatively easily satisfied in the WBSP due to the fact that it is inherently a multichannel processor; i.e., the processor must perform similar operations on a number of parallel channels. This fact alone suggests a multiprocessor configuration. As will be seen, the WBSP architecture takes full advantage of this inherent parallelism.

The initial candidate WBSP architecture consisted of a number of identical processors and a memory system all attached to a single bus structure. The active processors were

assigned various functions (e.g., demodulation, decoding) under FTSC control. Data were transferred from one processor or memory to another (e.g., from a demodulator to the deinterleaver memory) under the control of a dedicated bus controller. When a fault was detected, the FTSC intervened, isolating the faulty module and programming a spare module to take over its function.

The major disadvantage of the single-bus configuration was that the number of loads on the bus became quite large, causing excessive bus capacitance, and hence limited throughput, combined with high throughput requirements. Several modifications of the original configuration were investigated in an attempt to overcome this problem. The most promising configuration, and the one to be described here, involves four buses, all controlled by a single (dual-redundant) controller. The four-bus configuration obviously reduces both the number of loads and the throughput requirements on each bus. It does restrict the extent to which spares can be pooled, however, since a module can serve as a spare only for those modules interfacing with the same bus. This disadvantage is compensated for by the fact that the modules associated with each bus can be specifically tailored to a more restricted set of processing tasks, thus increasing their efficiency.

The WBSP configuration is shown in Figure 1. Each of the four buses consists of an eight-bit data byte, an eight-bit address byte, a data parity bit, an address parity bit and a spare byte to be used to replace any failed data or address byte or parity segment. Associated with each bus are two triple-modular-redundant (TMR) control lines. One is used by the controller to indicate that a valid address is on the address bus; the second is used to distinguish between hard and soft address modes (see next page).

In normal operation, the bus controller simply reads from its own memory a sequence of addresses for each of the four address buses, gates these addresses onto the buses and raises the valid-address control signal thereby initiating a new bus cycle. The most significant half of the address designates the address of the module that is to transmit onto the data bus during the following bus cycle; the least significant half designates the module that is to receive the data currently being transmitted.

The addresses used in normal operation are referred to as “soft” addresses. These addresses designate functions rather than modules. That is, a module responds to a given soft address only if it has been previously programmed to implement the function corresponding to that soft address. This initialization is done using “hard” addresses. Each module is permanently assigned a unique (for a given bus) hard, or physical, address to which it responds only in the hard-address mode. This mode is used only to test and to reconfigure the WBSP. When a fault is detected either by the FTSC or by the fault monitors (decoders) associated with each module interface, the FTSC is interrupted and, in effect, takes over control of the WBSP. By transmitting hard addresses over the WBSP

address buses, it 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 allows the WBSP to resume normal operation.

(The bus structure is critical to the fault tolerance of the system. Provisions must be made, for example, to ensure that failed modules can be deactivated, regardless of the nature of the failure; that once a failed module is deactivated, it cannot disable either of the buses with which it communicates, etc. All of these contingencies were successfully addressed in the FTSC design. Accordingly, except for the modifications needed to accommodate the narrower WBSP buses, the WBSP bus structure is identical to that used in the FTSC and is implemented exclusively with LSI devices being developed for the FTSC.)

In normal operation (cf. Figure 1), the bus controller routes information from the analog-to-digital converters to the demodulator processors; from there the demodulated (soft-decision) data are routed to the appropriate de-interleaver processor and then to one of the decoder processors. The decoded data is then stored into the FTSC's main memory where the integrity of each frame is tested (by checking the message "tails"). If a tail check indicates a malfunctioning WBSP module, the FTSC initiates the appropriate fault diagnostic routine and, if necessary, reconfigures the WBSP. Otherwise, the processed message is sent back to the WBSP to one of its deinterleaver processors where it is encoded and interleaved. From there, the bus controller routes it to one of the decoder processors where it is modified as appropriate (e.g., by being added to a pseudo-noise sequence). It is then sent to a modulator for down-link or cross-link transmission.

As indicated in Figure 1, the A/D converters and the modulators are configured in dual-redundant pairs, with each pair dedicated to a particular r.f. link. This configuration precludes the pooling of spares and hence results in a less efficient utilization of redundancy than would otherwise be possible. Pooling spare A/D converters or modulators would require analog multiplexers between these devices and their analog interfaces. Although this option need not be ruled out, it should be noted that both the A/D converters and the modulators are relatively simple devices; thus, the fact that redundancy is added somewhat inefficiently is not nearly as significant as it would be, for example, for the considerably more complex demodulator or decoder processors.

The communication link between the WBSP's bus controller and the FTSC could utilize either one of the FTSC's direct memory access (DMA) ports or one of its serial bus device interface units (DIUs). If the DMA port is not preempted for some other purpose and if the FTSC and WBSP are to be deployed in reasonably close proximity, the DMA interface is preferable since it enables more rapid communication between the two systems and hence shortens the time needed for diagnosis and recovery. Since the DIU port is fast enough for

normal operation, however, the only penalty in using it is in somewhat longer recovery period following a WBSP failure.

The only fault monitoring currently envisioned for the WBSP are the bus decoders provided at each bus interface and the “tail checking” already mentioned in the FTSC. Since the data are highly encoded, it is felt that the tail checks by themselves will provide an adequate means of monitoring the health of the system. The main function of the bus decoders is to help the FTSC isolate faults involving inter-module communication. If additional fault monitoring is found to be desirable, it would be relatively simple to add one processor on each bus and to program it to monitor, on a time shared basis, the performance of each of the other processors on the same bus. (It should be noted that the monitoring processor can in general execute considerably simpler algorithms than the processor being monitored and still thoroughly check the latter’s performance.) ,

The bulk of the processing capability, and hence of the system complexity, in the WBSP resides in the three sets of processors shown in Figure 1. These processors are the subject of the following section.

THE CORE PROCESSOR

The “core processor” forms the heart of the WBSP. The three major elements comprising the WBSP, the demodulator, deinterleaver, and decoder processors shown in Figure 1, are all variations on this basic processor design.

A block diagram of the core processor is shown in Figure 2. As can be seen, the processor architecture is that of a fairly conventional general-purpose computer consisting of a control section, a RALU section (register array and arithmetic logic unit), an input/output section, and a 256-word by 8-bit random-access, data storage (scratchpad) memory. It does have some unusual features, however, which merit discussion. Perhaps most important is the fact that it is to be implemented using LSI devices currently being developed for the Fault-Tolerant Spaceborne Computer (FTSC). Circuitry needed to support a fault-tolerant configuration is already designed into these devices and hence does not have to be re-invented for the present application.

The need for special fault-tolerant circuitry in the bus interface section (cf. Figure 1) has already been noted. It is of little benefit to have spare processors to replace defective processors if the latter cannot be logically isolated from the rest of the system. The safeguards needed to be certain that a processor can be successfully isolated, regardless of its failure mode, has been meticulously designed into the bus interface logic (cf. Ref. 1 for details).

The architecture shown in Figure 2 is also unconventional in three other minor but, for the applications of concern here, important ways: (1) The control memory, a RAM rather than a ROM, can be reloaded under external control, thereby allowing the processor to be assigned new tasks dynamically as the system requirements change due to processor failures or for other reasons. (2) Certain of the ALU output bits can be individually specified as control RAM branch bits. (Usually only sign, carry-out, overflow, and the all-zeros conditions are used to effect branches.) (3) The microcode can specify that certain of the bits of the data RAM address generated by the processor are to be complemented, thereby allowing the microcode to modify processor-generated base addresses. The utility of the first of these features is immediately apparent when it is recognized that the demodulation or decoding algorithms to be implemented need not be the same for all communication channels. The latter two features proved to be quite useful in defining the microroutines needed to implement certain decoding and demodulation algorithms.

The decoder processor shown in Figure 1 is identical to the core processor. The deinterleave processor consists of a core processor with a considerably expanded buffer memory. The demodulator processor is basically two cor processors operating in parallel, each augmented with a multiplier chip and, again, an expanded memory for storing sampled input data and intermediate results.

CONCLUSIONS

The WBSP communications processor configuration shown in Figure 1 satisfies one of the fundamental constriants essential to efficient high-reliability design: The most complex processing tasks have been partitioned so that they can be performed by a number of identical, interchangeable modules. A detailed comparison of this approach to a more conventional approach, for example, one using special-purpose processors with system-level redundancy, is beyond the scope of the present paper. Obviously the results of such a comparison depend upon the details of the alternative implementation, the number of channels to be processed, the channel data rates, the decoding and demodulation algorithms to be implemented and the overall reliability objectives. When long-term high reliability is required (e.g., a 7 year survival probability of 95% or greater), comparisons have shown that the WBSP approach offers significant advantages relative to conventional methods of implementing the same processing tasks. The potential disadvantages of having to implement a function with a general-purpose processor (the core processor and its variants) rather than a special-purpose processor is more than overcome by the efficiencies encountered in partitioning the processing task among a number of interchangeable modules and providing these modules with pooled spares.

REFERENCES

(1) Stiffler, Jack, "Fault-Tolerant Spaceborne computer (FTSC) Study," Report No. ER78-4173, Raytheon Company, May, 1978.

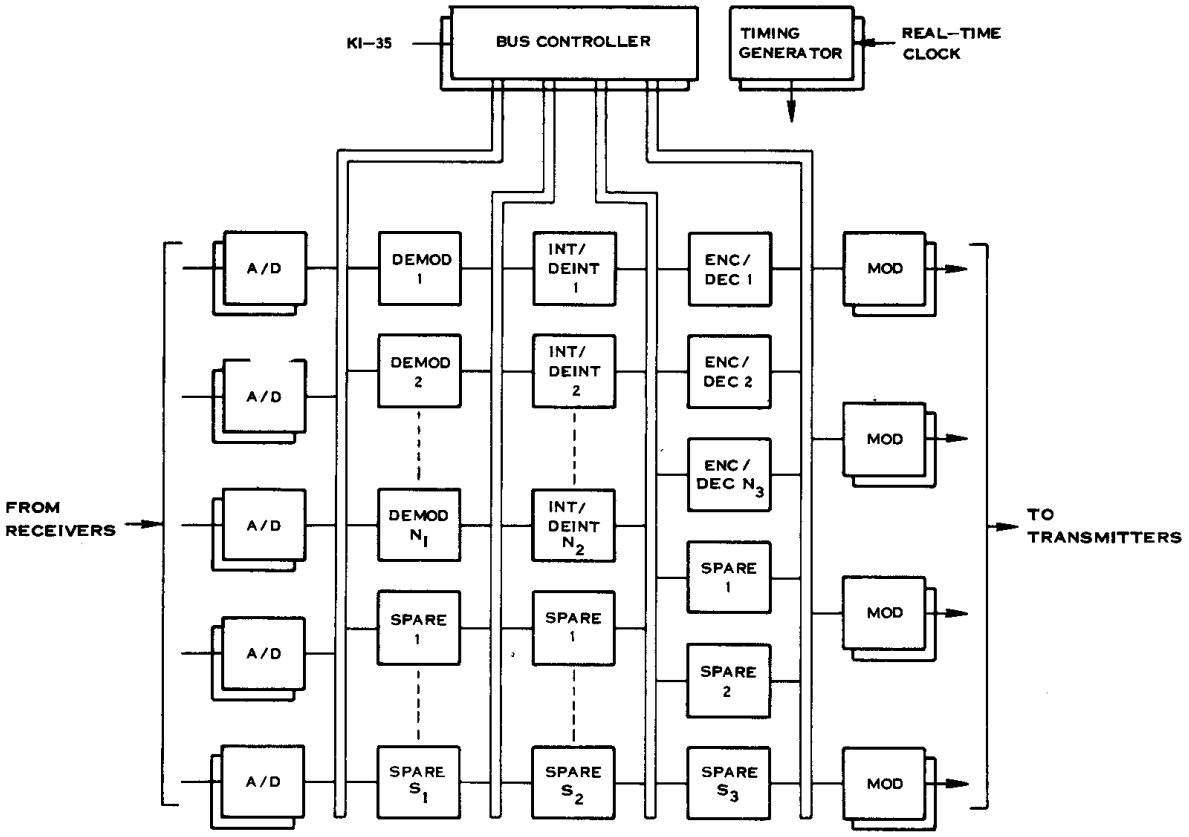

Figure 1 WIDE-BAND SIGNAL PROCESSOR

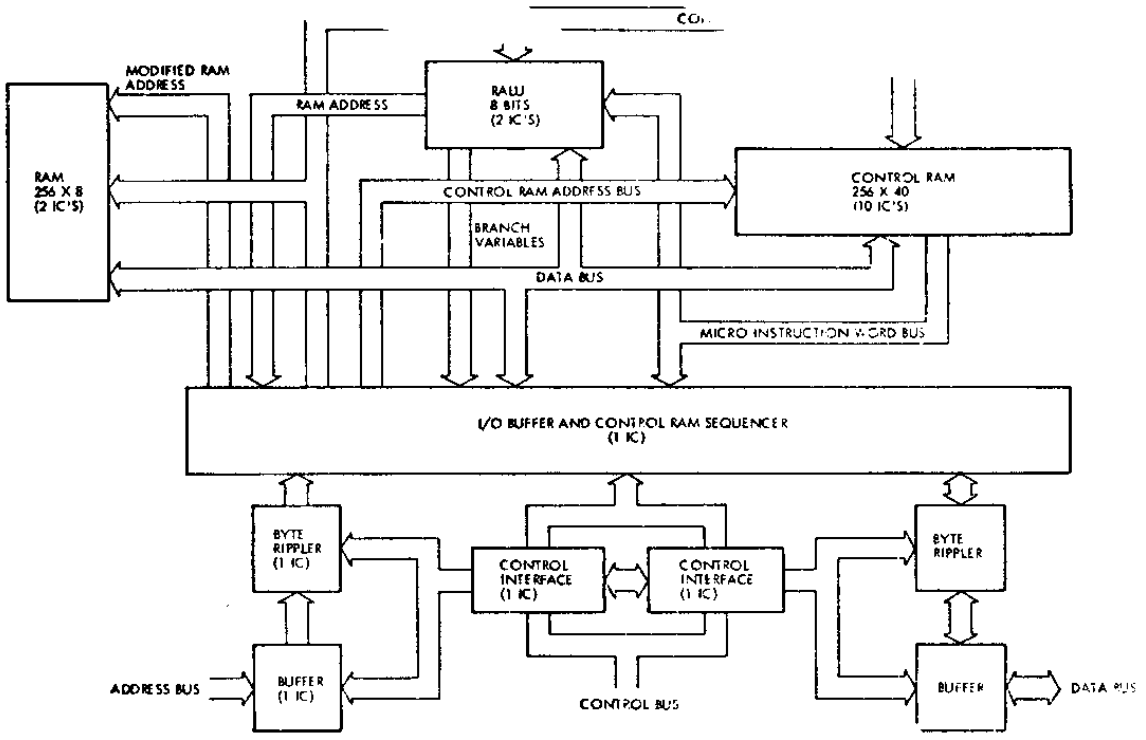

Figure 2 CORE PROCESSOR