

THE NEXT-GENERATION SC-7 RISC SPACEFLIGHT COMPUTER

D. J. Shirley and M. K. McLelland
Southwest Research Institute
San Antonio, Texas

Abstract

Southwest Research Institute (SwRI) recently undertook a search for the microprocessor for the next generation SwRI spaceflight processor. The driving selection criteria were versatility, performance, radiation characteristics, and availability of development tools. The microprocessor selected was the RH320C30 from Texas Instruments, a silicon-on-insulator implementation of the popular TMS320C30. Although the TMS320C30 is primarily classified as a Digital Signal Processor (DSP), the C30's open RISC design makes it ideal for use as a general-purpose computer core. A Harvard architecture, dual independent 32-bit data buses, instruction pipelining, internal RAM, and on-chip I/O peripherals make the C30 easy to use in embedded applications. The C30 will soon be available in the radiation hard SIMOX process. With its chip power requirement of approximately one watt, a total radiation dose capability of 1M Rad, single-event latchup-free operation and very low single event upset rate, the C30 incorporates all the requirements needed for a radiation hardened spacecraft computer design. Southwest Research Institute has used the C30 as the core CPU for its SC-7 Spacecraft Computer design. This new design is presented along with the power, speed, and other parametric data measured for the computer.

I. Introduction

In general, an embedded processor design intended to be used for spaceflight applications poses a number of system constraints that limit the number of useable processor types available. With the advent of NASA's "Faster, Cheaper, Better" philosophy, the demands on both spacecraft and payload processors have increased to the point that older designs are not applicable, performing at too low a processing speed, having too much mass and volume, requiring too much power, and being readily susceptible to various forms of radiation damage.

In order to remain competitive in the new spaceflight design environment, Southwest Research Institute (SwRI), having based its previous spacecraft computer designs on the Intel 80X86 chipset, embarked on a new development using a Reduced Instruction Set Computers (RISC) for spaceflight applications.

The use of RISC-based computers in space has been pioneered by such firms as IBM, TRW, Honeywell, TI, Harris and IDT, who have developed chipsets and systems under Ballistic Missile Defense Organization (BMDO) funding in order to meet the requirements of the space-based missile protection systems promoted in the past decade. Although these new RISC-based systems have the requisite processing power, less thought has been given to reduction of electrical power, mass and volume requirements. Another factor looming on the horizon is the cessation of the Cold War

and resultant decrease in emphasis on the space-based weapons systems. It is probable that existing RISC processor systems will not be taken further past existing design stages, and may not be supported by their developers.

The impetus for development of the SC-7 design was the requirement for a small, low-powered, low cost, RISC-based spaceflight computer suitable for use on small satellites in earth orbit or interplanetary flight.

II. Design Background

In order to develop the necessary design requirements for the SC-7, a set of criteria was developed for selection of the core processor chip or chip set. The criteria were versatility, performance, radiation characteristics, and availability of development tools. It was found early on that some of the criteria were somewhat mutually exclusive, and that others were mutually dependent, but it was felt that all were of equal importance and were thus kept intact.

A. Versatility

The design of the SC-7 is intended to fulfill the requirements of a large range of spaceflight applications ranging from spacecraft command and data handling (C&DH) computers, attitude control processors, experiment control processors and number crunching compute engines for specialized sensors. Although the versatility of a computer design is more a function of the peripheral and interface circuits, it was felt that the level of integration of the peripheral functions would affect the versatility of the final design, i.e., built-in serial ports, DMA controllers, timers and memory are more quickly and readily

accessible by the processor if they are inside the chip and do not require access through an external address/data bus. Versatility also requires that an appropriate set of instructions be available, but this is of less importance in RISC-based applications where higher order instructions are implemented by compiled RISC instructions, making versatility dependent upon the availability of an optimizing compiler. Another factor affecting versatility is the availability of a hardware floating-point processing capability. Floating point operations are quite often a requirement in spaceflight applications, and reliance on software implementation of floating point processing can have a major impact on processing speed.

B. Performance

The speed with which a computer can process instructions and data is of primary importance in any application. However, for spaceflight purposes, a ratio of processing speed to required electrical power is a more useful performance criteria. Some computer workstations have enormous computing power but pay a heavy price in the amount of power required to do the processing. Many of the RISC spaceflight computer systems now being developed are based on rad-hard implementation of a commercial workstation chipset and thus suffer from a low speed/power ratio. To produce a useful design, the power required for the core processing circuits must be kept below the three to four watt level to meet power requirements for a satellite with total available power on the order of 50W or less. In order to produce this power level, the processor chip must have a power drain below the one watt level. On the other hand, the speed of such a processor must be higher than the one to two million operations per

second (MOPS) of already existing rad hard MIL-STD-1750 computers.

It was recognized that in order to meet the speed/power goals for the SC-7, a 32-bit wide data bus would be required. Although 64-bit wide bus systems are currently being brought on the market, it was felt that the stage of development would preclude finding a currently available radiation-hardened system.

C. Radiation Characteristics

A computer that is designed for use in a spacecraft must be able to survive the radiation environment that is encountered during its lifetime whether it is in low earth orbit (LEO), geosynchronous orbit (GEO) or in interplanetary flight. Each of these environments presents different problems for spaceflight computer design.

A system in LEO will require that the components be able to withstand a high total dose of energetic particles because of trapping of these particles by the earth's magnetic field. The effect of this radiation dose is to degrade the operating parameters of electronic circuits and will eventually cause them to cease to operate. Quite often, simple shielding by the spacecraft and enclosures can help meet this requirement and has been used in the past to enable non-radiation hardened components to be used. However, this introduces an increase in mass of the system since shielding effectiveness is directly proportional to the density of the shielding material. In order to reduce the necessary shielding required for a LEO satellite with a one to two year life expectancy, a spaceflight computer should be able to withstand at least 50 KRad of total dose.

Highly energetic particles from the solar environment and from cosmic rays cause two other effects to take place in

electronic circuits. High energy particles are especially prevalent for GEO and interplanetary applications. The first effect is single-event upset (SEU) which is the occurrence of a bit flip in a digital circuit when the passage of a particle releases sufficient energy to overcome the flip-flop feedback voltage. This bit flip is non-destructive and can be recovered from in various ways, including error detection and correction (EDC) on memory and bus circuits and watchdog timers to reset the processor in case of a non-recoverable error. Radiation hardened integrated circuits use processes that reduce the susceptibility to SEU to a low level, although none can reduce it to zero.

A second effect from high-energy particles is single event latchup (SEL). SEL can be destructive since it involves the firing of an SCR structure inherent in bulk CMOS circuits. SEL can be completely eliminated by the use of special IC manufacturing techniques which place the active CMOS structures on an insulating substrate such as silicon-on-sapphire (SOS) or silicon-on-insulator (SOI). Some of the CMOS Epitaxial (CMOS/EPI) processes in which selected junctions are isolated with epitaxial structures can also eliminate SEL or at least create conditions such that only low-probability, extremely high energy particles can interact and cause SEL.

Since an SEL event creates an extreme hazard for a spacecraft computer, an acceptable design would reduce the possibility of an occurrence to the smallest degree possible.

D. Availability of Development Tools

One often overlooked factor in new computer designs is the ready availability of software development tools. The best computer designs are often reduced to

practical unusability by the lack of easy-to-use development tools. These tools include high-level language compilers, especially optimizing compilers for RISC development, debuggers, assemblers (not too useful for RISC) and in-circuit emulation systems. The ready availability of easy-to-use development tools was made a factor of prime importance in our processor selection.

II. TMS320C30 Processor

During the initial design phase, all the available rad hard RISC processor systems were examined for suitability, and the final selection was the Texas Instruments (TI) RH320C30 Digital Signal Processor chip.

Although not immediately recognized as a viable candidate for a general-purpose RISC processor, close examination of its design, operation and instruction set reveals that it is an open architecture system with all the requirements for a RISC processor.

A. Architecture

The RH320C30 is a radiation-hardened version of the popular

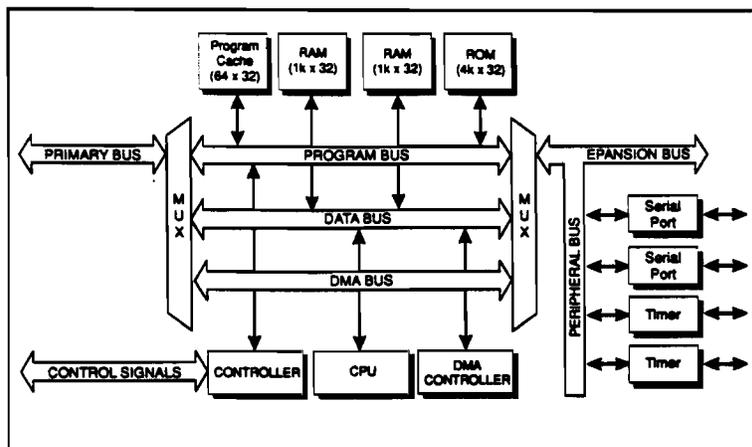
Harvard architecture computer system with on-board program cache, RAM, PROM, DMA Controller, CPU, FPU, and peripheral circuits consisting of timers and serial ports. The Harvard architecture is quasi in the sense that the internal program and data buses are separate inside the chip but are common external to the chip. The C30 uses register-to-register instruction operations at a rate of one instruction every two clock cycles. Floating point operations occur once every clock cycle. Thus, at the maximum 28 Mhz clock speed of the RH320C30, the processor can operate at 14 million instructions per second (MIPS) and 28 million floating point operations per second (MFLOPS).

Figure 1 shows a low-level block diagram of the C30 internal architecture. Three separate buses, the program bus, the data bus, and the DMA bus allow a high degree of parallelism to be realized. Program, data, and DMA accesses can be made during the same clock cycle. The busses connect all of the physical address spaces of the C30 including on-chip memory, off-chip memory and on-chip peripherals. The DMA controller can access the complete 24-bit address space (16 Mword) of the C30. All data accesses are done at a full 32-bit word level with no byte accesses.

The C30 contains two independent data buses, the Primary Bus with 32-bit data and 24-bit address, and the Expansion Bus with 32-bit data and 13-bit address. Accesses on the two buses are independent and may be made in parallel during the same clock cycle.

A 64 X 32 Program Cache section is contained in the C30 and allows pipelining of the processor instructions. The cache is configured in a 64-way, set-associative manner to provide a

Figure 1 RH320C30 Block Diagram



TMS320C30 DSP chip available from TI in their proprietary SIMOX[®] process. The C30 is a single-chip implementation of a quasi-

instructions. The cache is configured in a 64-way, set-associative manner to provide a

significant performance improvement. The cache operates autonomously with no operator intervention required.

Two blocks of 1K X 32 user RAM are also contained inside the C30. These RAM sections can be accessed by the processor with zero wait states to enable a small but fast data storage area.

On the TMS320C30, the 4K X 32 PROM is used to store masked programmed instructions and is accessible only in the Microcomputer Mode. The Microcomputer Mode and Microprocessor Modes are pin-selectable operating modes with the Computer Mode intended to operate as a single-chip computer. In the RH320C30, This has been changed to contain boot code for power-on initialization and the pin that programs the modes in the TMS version selects the internal boot code or external boot code.

A Controller Section connected to the Program Bus contains the Program Counter Register and Index Registers. The Controller provides interfacing to external signals such as clock, reset and interrupt signals.

The CPU section contains the register-based processing circuits and consists of a floating-point/integer multiplier, an arithmetic logic unit (ALU) for arithmetic and logical operations, a 32-bit barrel shifter, auxiliary register arithmetic units (ARAUs), internal buses, and the CPU register file.

The multiplier performs single-cycle multiplication on 24-bit integer and 32-bit floating point values. When performing floating-point multiplication, the inputs are 32-bit floating-point numbers and the result is a 40-bit floating-point number. For integer multiplication, the input values are 24-bit data and yield a 32-bit result.

The ALU performs single-cycle operations on 32-bit integer, 32-bit logical, and 40-bit floating-point data. This also

includes single-cycle integer and floating-point conversions. ALU results are either 32-bit integer or 40-bit floating-point values. The function of the barrel shifter is to shift 32 bits either left or right in one cycle. Two internal CPU and two internal Register buses are used to carry two operands from memory and two operands from the register file to allow parallel multiplies and adds/subtracts on four integer or floating-point operands in a single cycle.

To support addressing with displacements, two index registers, and circular and bit-reversed addressing, two ARAUs are used to generate two addresses in a single cycle and operate in parallel with the multiplier and ALU.

The CPU register file supplies 28 registers that are 32-bit wide. These registers are tightly coupled to the CPU and support a wide range of general and special purpose operations.

The C30 contains a single-channel DMA controller with its own internal bus. The DMA controller thus can read or write to any location in the memory map without interfering with CPU operations. The controller contains its own address generator, source and destination registers and transfer counter.

The peripheral circuits included on the C30 chip includes two programmable 32-bit serial ports and two programmable 32-bit timer/counters. The serial ports are totally independent and can be programmed to synchronously transfer 8, 16, 24, or 32 bits of data per word with either internal or external clocks. An RS-232 format asynchronous serial port can also be emulated with the proper software code. Three external lines are provided for each serial port for data, clock and enable input/output.

The timer/counters are general-purpose and operate in two signaling modes

with either internal or external clocking. Each timer has an external I/O pin for use as an external clock or as an output pin driven by the timer. All the serial and timer external pins can also be programmed as general-purpose I/O pins.

The serial ports and timers have their own internal address and data bus internally connected to the Expansion Bus.

The external buses of the C30 are totally independent and contain their own control signals. Both buses can be used to address either external program/data memory or I/O space. Each bus contains an independent READY signal for extra wait-state generation by external slow-access circuits. Wait states can also be generated under software control internally by the C30.

B. Speed

The commercial and military versions of the TMS320C30 can operate up to 40 MHz. The RH320C30, due to limitations of the radiation hardening process is limited to a maximum of 28 MHz. This clock speed will allow a CPU speed of 14 MIPS and 28 MFLOPS. A speed of 14 MIPS requires an external access to be performed in 71 ns which places some extreme speed requirements on external memory and address decoding components, especially when error-detect and correct (EDC) circuits are used to reduce random bit errors. Typically, one or two wait states must be generated for external memory accesses. In the case of the rather slow PROMS and EEPROMS that are available, a system that dumps operating code from PROM into RAM for running the code out of RAM is required to maximize system performance. Use of the internal instruction pipeline and temporary RAM storage areas can help the speed reductions due to external-access wait states.

C. Radiation

The RH320C30 is implemented on a 0.8 μm SIMOX[®] process. Preliminary design goals for this implementation were: (1) an SEU rate of less than $1\text{E}-10$ errors per bit per day (worst case), (2) a transient dose hardness without latchup up to $1\text{E}11$ rad (Si)/s, and (3) a total dose hardness of greater than 1 Mrad(Si). SEL immunity is guaranteed by the use of the SOI process.

Preliminary radiation testing of the RH320C30 parts show that the actual total incident dose (TID) will be between 1 and 2 Mrad(Si) with an linear energy transfer (LET) of 140 Mev/gm/cm² and a cross-section of $1\text{E}-9$ cm² for SEU.

D. Self-Test

Another feature that was implemented for the RH320C30 but is not available on the TMS320C30 are two self-test features. A set of Built-In-Test (BIT) code is included with the boot code to test the internal parts of the C30 during initialization. This test can complete in 58 ms. Another self-test is the inclusion of JTAG boundary scan. This feature is accessed through the existing C30 emulation/scan test port by pin selection using a previously unused pin.

E. Development Tools

An extensive base of development support tools exists for the C30, both from TI and from third party vendors. The software tools available from TI include a Macro Assembler/Linker, an Optimizing ANSI C Compiler and a Simulator. An ADA 9X compiler is also under development for the C30 by TI. A real-time operating system, SPOX, is available from third party

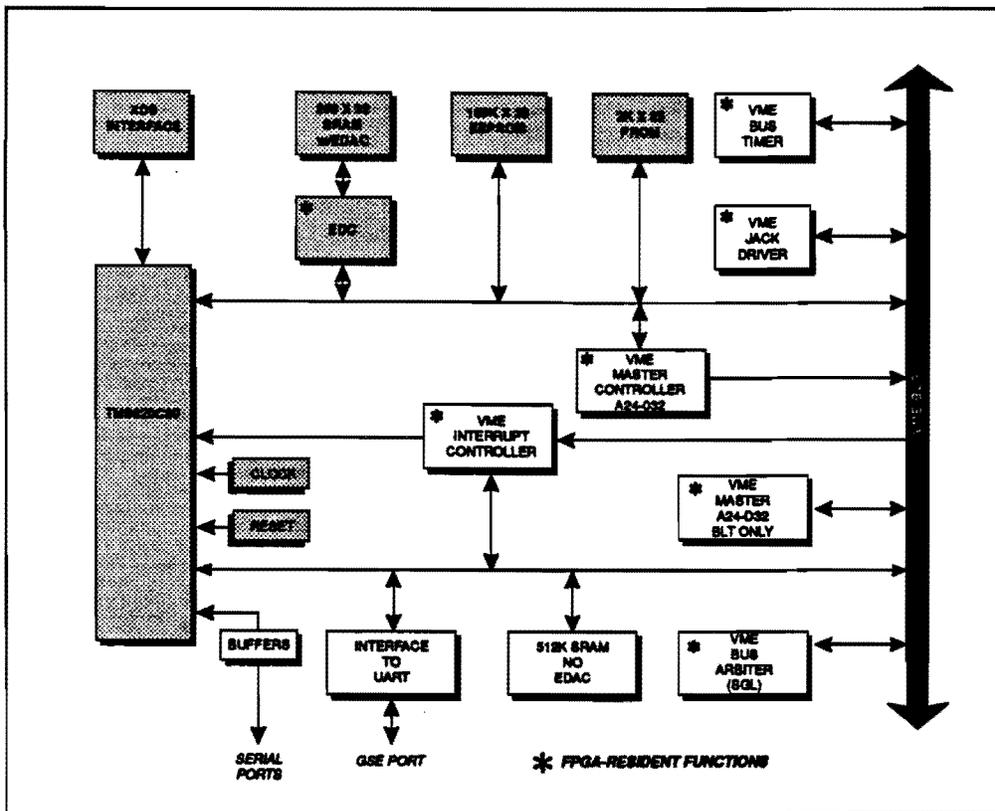
vendors as well as various packages of application software, although the latter are mostly aimed at DSP applications.

Hardware development tools available include an Evaluation Module (PC Plug-in), and a XDS510 Emulator, both from TI. An H-P 64776 Analysis Subsystem is available from Hewlet-Packard, and a multitude of PC plug-in cards for both PCs and other workstations are available from third party vendors.

IV. SC-7 Design

The following sections will describe

Figure 2 SC-7 Block Diagram



the design of the spaceflight computer based upon the RH320C30 chip. The SC-7 is the latest in the family of SC-Series of affordable and highly reliable spacecraft computers from Southwest Research Institute. The goal of the design effort was to develop a

core processor system (CPU, memory, interrupt, and control circuits) which would be available for inclusion in a complete design with I/O and interface circuits. A secondary goal was to implement a complete design in order to evaluate the speed and utility of the computer system operating in a real environment. The resultant design is shown in the block diagram of Figure 2. In the figure, the core processor blocks are shown with shading, while the I/O and interface circuits are shown without shading. As implemented in hardware, the SC-7 comprises a 6U VME card and functions as a VME Bus Master, Bus Timer, Bus Arbiter,

and IACK Driver. The Bus Master functions are divided between the primary bus which implements an A24:D32 Bus Master, and the expansion bus which implements an A24:D32 Block Transfer Mode (BLT) Only Bus Master. The expansion bus Master transfers VME Bus Data only into the expansion bus RAM.

Line drawings of the SC-7 card are shown in Figure 3. The assembly is fabricated with all surface-mounted components with the exception of connectors. Components are mounted on both sides of the board, with a

heavy copper plane in the center of the board to conduct heat from the components to the card edges.

A. Processor Core

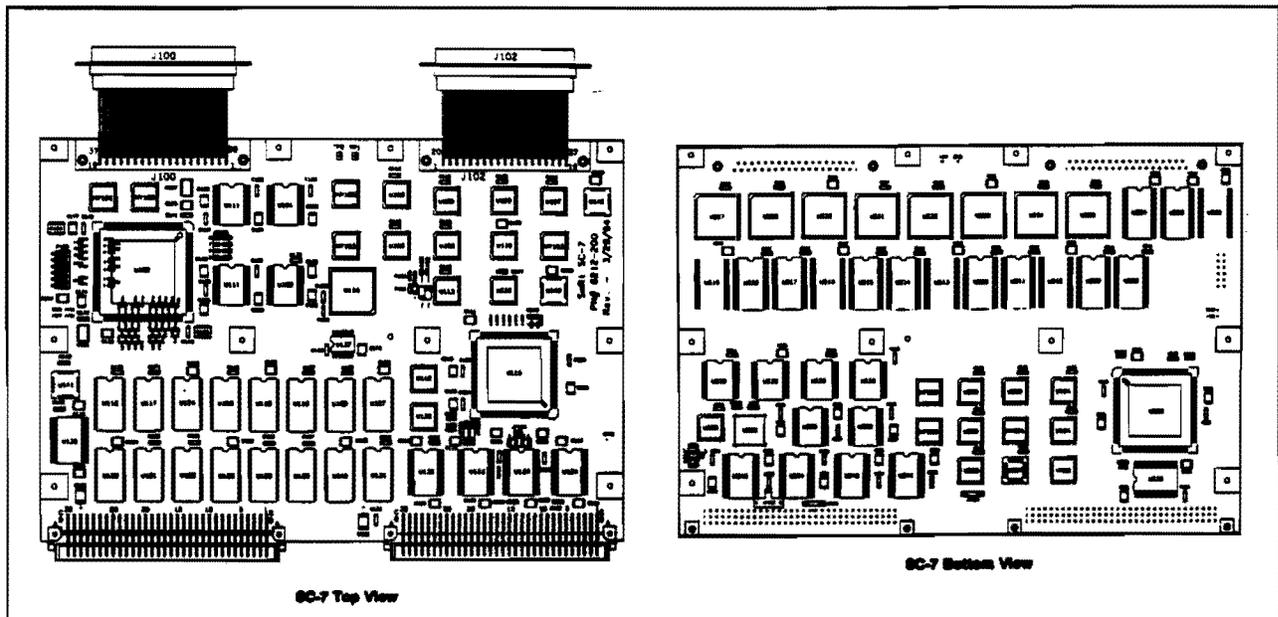
The Processor Core of the SC-7 computer consists of the C30 processor, the RAM, EEPROM, and PROM memory, an interrupt processor, and a memory management unit. These four blocks are capable of functioning alone to perform computations and calculations

in a desk-top computer and is used for in-circuit emulation and testing. The XDS port can also be used to download code into either the RAM or the EEPROM via the processor chip. The processor implements both the primary bus which contains the RAM, EEPROM and PROM memory, and the expansion bus which contains the UART interface, and an additional RAM area.

2. RAM

The primary bus carries the memory

Figure 3 SC-7 Board Layout



1. Processor

The processor for the SC-7 is the RH320C30 although hardware fabricated for testing in the program utilized the TMS320C30. Included with the processor are a clock operating at 24 MHz, a power-on reset circuit, and an Extended Development System (XDS) interface. The XDS interface is comprised of a 6-line connection to the C30 chip using a 12-pin header arrangement connected to an XDS emulator card resident

for the core processor. Since high-density rad hard RAMs are not easily obtainable at this time, the 256K X 32 SRAM area is implemented with Micron 1 MBit static RAMs. These SRAMs have good radiation characteristics as far as TID and SEL are concerned, although they are rather soft to SEU. For this reason, a flow-through EDC circuit was implemented to correct and flag single-bit errors and to flag double and some higher order bit errors. Bit errors are flagged by interrupts to the processor. Automatic error scrubbing is not

implemented with this EDC circuit. As indicated by an asterisk on the block diagram, the EDC is implemented with a Field-Programmable Gate Array (FPGA). Address decoding is also implemented in the FPGA. Although the SRAM has an access time of 20 ns, the inherent delays involved with using the FPGA to implement decoding and EDC requires the insertion of two wait states for SRAM accesses.

3. EEPROM

Since the SC-7 is intended for spacecraft use, having the capability to upload updated code on-orbit is an important factor. A 256K X 32 EEPROM memory was implemented in the SC-7. This design used the SEEQ 1 MBit EEPROMs which have an internal EDC function so an external EDC was not required. The EEPROMs can be accessed in either the read or write mode by the processor for code downloading. The access time for the EEPROM chips is 200 ns which is prohibitively slow for running code. The intended operation of the SC-7 requires the transfer of operating code from EEPROM to RAM during the boot sequence. The code may then be run out of the RAM. The write function to the EEPROM is protected with a software lock feature as well as a hardware feature which locks out writes when the power supply level decreases below 4.5V.

4. PROM

In order to have a block of code that has the least possibility of being corrupted and thus to have a capability to reboot the computer and recover a system crash, a 2K X 32 block of rad-hard PROM is included in the design. Harris HS-6617RH PROMs were used. This PROM contains the startup boot code which initializes the processor and

transfers the operational code from EEPROM to RAM.

5. Interrupts

The C30 processor has four external priority interrupt inputs as well as seven internal interrupts for the serial ports, the timers, and the DMA controller. The highest priority external interrupt is used to flag a correctable or uncorrectable error, the second highest is used to indicate a VME Bus interrupt, and the third highest a completion of a VME Block Transfer. The fourth interrupt is unused. An interrupt processor with registers is implemented in the FPGA to store data that can be accessed by the C30 to determine the exact interrupt and the action that must be taken.

B. VME Interface

The SC-7 includes the following VME interfaces:

1. An A(24):D(32) VME Master with capability for bus arbitration per IEEE Standard 1014;
2. An A(24):D(32) Block Transfer Only VME Master with capability for bus arbitration per IEEE Standard 1014;
3. An IH(1-7):D(16) VME Interrupt Handler with capability for bus arbitration per IEEE Standard 1014;
4. A VME Bus Arbiter;
5. A VME Bus Timer; and
6. A VME IACK Driver.

The inclusion of the Bus Arbiter, Timer and IACK Driver circuits on the computer card can reduce the number of cards required in a VME system.

C. GSE Port

Although the XDS interface, which is a standard feature of the C30 system allows access to the SC-7 by an external GSE computer system, a UART interface is also included on the SC-7 design for use in case a GSE system with only an RS-232 serial port. The GSE port only includes buffers and address decoding and not the GSE UART itself. An external board with the serial port circuits can be attached to one of the front connectors on the VME card to create the GSE interface.

V. Test and Evaluation

The prototype board has been tested with test software developed for the SC-7 program to evaluate design performance. The board was tested in two stages in order to evaluate the power required for the core processor (CPU, Memory and Memory Controller) as well as for the complete VME card which includes the core processor and the VME controllers and interface circuits. After completion of the board, it was evaluated for computational speed and performance using a standard Gibson Instruction Mix, and tested in a VME system for operational bus speed using standard read/write over the VME Bus and Block Transfer Mode.

A. Power

1. Core Processor

The core processor power was evaluated during fabrication of the prototype

board by populating only the processor chip, the RAM chips, the EEPROM chips, the PROM chips, and the memory controller FPGA. This configuration provides the core processor necessary for a design in which other peripheral circuits could be added to make a specialized computer for any application. This core design offers 256K X 32 of static RAM with Error Detection and Correction (EDC), 256K X 32 of EEPROM memory with no EDC, and 2K X 32 of rad hard PROM with no EDC. The core system was operated at 24 Mhz clock speed and required only 3.25W of power

2. Full VME Processor

After completion of fabrication of the full prototype card, the system was again evaluated for power usage. The card was installed in a standard commercial VME card cage with a Crislin Industries 16 Mbyte DRAM VME card (Type CIVME40). The power was measured for both cards operating in the block transfer mode at maximum speed (>16 Mbytes/second). The power was then measured for the memory card alone operating statically and this power was subtracted from the previous measurement value. The power for the SC-7 VME card operating in a backplane environment, fully operational at maximum speed, was found to be 6.85W.

B. Computational Speed

Once the prototype board fabrication was complete and ready to be evaluated, benchmark tests were run on the system to evaluate its performance under varying applications. A Gibson Mix benchmark test was done. This benchmark was selected in order to compare the processing speed of the SC-7 to that of the previous spacecraft computer designs at SwRI.

1. Gibson Mix Test

Operational tests done using a Gibson Mix of instructions are used to evaluate performance of a processor system where a mix of logical, fixed point and floating point instructions are required. The routine is designed to derive a performance rating by running a mix of 1000 instructions and deriving a running time. The instruction mix is presented below in Table 1. Cascaded timers are used to calculate the number of microseconds needed to run all of the instructions. The performance figure is calculated by dividing the result by 1000 to express the number of seconds per average instruction and inverting that number to the number of operations per second. This figure is usually given in millions of operations per second (MOPS) or Thousands of operations per second (KOPS).

Results of the Gibson Mix test show that the SC-7 performs at a rate of 3.08 MOPS when tested in a mode with two wait

A test in which zero wait states and no caching was used, the performance was 8.78 MOPS. When two wait states plus cache was used the performance was 5.85 MOPS. For a comparison, the SC-3 computer design that is based upon a 16 MHz 80386/80387 processor set with cached memory, is capable of performing at 2.4 MOPS Gibson Mix and requires 15W of power for the core processor circuits. The SC-7 design thus has a 1.8 MOPS/W speed/power (two wait states, cache used) ratio whereas the SC-3 has a 0.16 MOPS/W ratio.

2. VME Block Transfer Test

The speed at which the SC-7 board could perform in the VME Block Transfer mode was tested. The block transfer mode of the SC-7 operates to transfer 64 Words (32 bits per Word) without interruption and without requiring new addresses for each Word. The test setup used for power testing was used and the speed at which a block transfer read or write could be performed was measured. The SC-7 can transfer 18.29

Table 1 Gibson Instruction Mix

<u>INSTRUCTION</u>	<u>MIX FACTOR</u>
Load Register (No Indexing)	0.312
Load Register (Indexing)	0.180
Condition Test and Branch (Half of Branches Taken)	0.166
Compare	0.038
Shift (3 Bits)	0.044
Logical AND	0.016
Instruction with Minimum Execution Time	0.053
Fixed Point Addition	0.061
Fixed Point Multiplication	0.006
Fixed Point Division	0.002
Floating Point Addition	0.069
Floating Point Multiplication	0.038
Floating Point Division	0.015

states generated for a RAM access and no caching of data or instructions (i.e., no instruction loops were used in the test code).

Mbytes/second on the VME Backplane operating in this mode. Operating in a non-

block transfer mode, the speed is 4.0 Mbytes/second.

VI. Discussion and Conclusion

This paper has described the successful development of a RISC-based spacecraft computer design that has been designated the SC-7. This development has resulted in a design that meets the requirements of a spacecraft computer as far as electrical power requirements, processing speed, and radiation environment capability. The design represents a ten-fold increase in the speed/power ratio over that of the SC-3 that was previously the fastest processor design at SwRI using an 80386/387 chipset. Although presently applied as a VME processor, the core processor part of the design is adaptable to many other types of designs with changes to the mechanical configuration and modification of the interface circuits. The present design uses parts that are available either in fully radiation hard technology or in radiation tolerant technology and meets the criteria of >50K Rad Total Dose and is Single Event Latchup immune.