



Fermi National Accelerator Laboratory

FERMILAB-Conf-89/52
[SSC-207]

**Microprocessors and Other Processors
for Triggering and Filtering at the SSC***

Irwin Gaines
Fermi National Accelerator Laboratory
P.O. Box 500, Batavia, Illinois 60510

March 1989

* Presented at the Workshop on Triggering and Data Acquisition for Experiments at the Superconducting Super Collider, Toronto, Canada, January 16-19, 1989.



Operated by Universities Research Association, Inc., under contract with the United States Department of Energy

Microprocessors and Other Processors for Triggering and Filtering at the SSC

Irwin Gaines

Fermilab Advanced Computer Program, Batavia, IL 60510

INTRODUCTION

The rapid increase in processing power available in commercial integrated circuits presents the high energy physics community with important opportunities for SSC era data acquisition systems. The processors that will be available in the late 90's will allow enormous amounts of computing power to be utilized on-line and will permit commercial high level language programmable devices to be used for tasks previously performed by home-brew, hard-wired, or microcoded devices.

I will describe this processor revolution, in particular with respect to the new RISC (reduced instruction set computer) microprocessors now becoming available. These processors are already commercially available with processing power of 20 VAX 11/780 equivalents per chip, and a number of different manufacturers expect chips of 100 VAX power by the early 1990's. I will also discuss the plans the Fermilab ACP group has to exploit one such RISC chip for both off-line and on-line use. Finally, I will mention digital signal processors (DSPs) and other more specialized chips that offer even greater amounts of processing power with only slightly less convenience.

RISC PROCESSORS

Reduced Instruction Set (RISC) microprocessors have gone from an academic research project to practically a computing industry standard in a short period of time. Every leading semiconductor manufacturer and computer vendor have RISC projects underway. The current generation of RISC processors has already surpassed the more common CISC (Complex Instruction Set) architecture (typified by the Motorola 68032 and the Intel 80386) in performance, and shows signs of surpassing mainframe performance as well. Figure 1, showing one vendor's projections of the computing power available in different architectures, illustrates these ideas.

What is RISC, and why do these processors have such high performance? The principle of RISC is to keep the instruction set of the processor as simple as possible, so that all instructions can be executed in a single clock cycle. This is in contrast to the prevailing design philosophy of the 60's and 70's, where instruction sets were filled with an enormous variety of instruction types and addressing modes, supposedly to make it easier to write

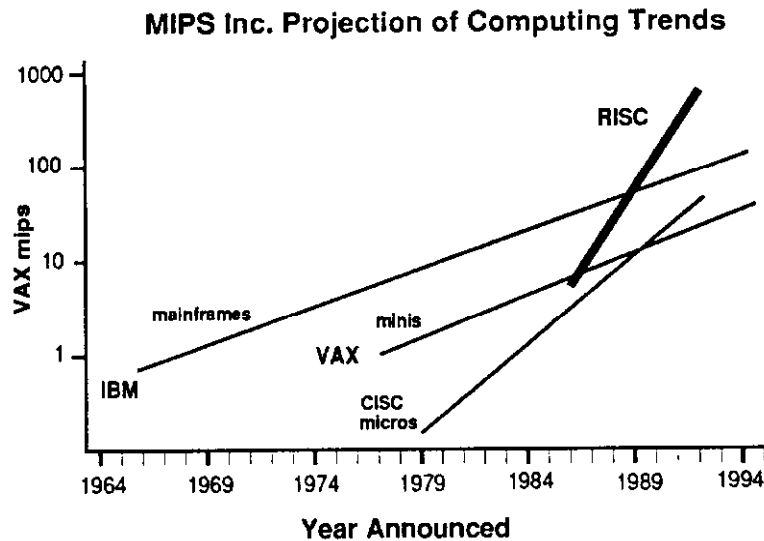


Figure 1. Processing power available in different architectures.

compilers for high level languages. However, study of the instructions used by compilers indicated that only a small fraction of the complex instruction sets were frequently used, and that the complexity led to a large loss in performance, even when performing the simplest instructions. Eliminating many of these instructions results in an architectural simplicity allowing the remainder of instructions to be executed in one processor clock cycle. Infrequently performed complex operations are done in software rather than in hardware, so that there will be no performance penalty for the vast majority of simpler operations.

More specifically, the time to perform any given computing task depends on the product of three factors: the number of instruction needed to do the task; the number of clock cycles needed for each instruction; and the amount of real time required for each clock cycle. RISC processors win by giving an enormous reduction in the 2nd component (cycles/instruction)—while a VAX 11/780 averages 10.6 cycles per instruction and the Motorola 68020 averages 6.3 cycles, a modern RISC processor like the MIPS R3000 requires only 1.25 cycles per instruction (and future RISC processors expect to lower this average to 1.0 or below).¹ Moreover, the simplicity of the architecture allows the RISC processors to run at higher clock speeds than CISCs (as well as allowing implementation in faster technologies like ECL or GaAs that are not suitable for CISC architectures). This gain is accompanied by an increase in the required number of instructions, but only by 20-50%, leaving the RISCs with a large overall performance gain.

A variety of architectural techniques allow RISC processors to achieve their goal of single cycle instruction execution. Typically there is a relatively small number of instructions and addressing modes, and a fixed instruction format. The RISC designs are often a load/store architecture, with large register sets and no memory-to-memory instructions. Control logic is usually hard-wired, with none of the microcoded control typical of

minicomputers and CISC processors. Much more of a burden is put on the compilers, with sophisticated optimizations required to achieve full performance in high level languages. (This is why the currently leading RISC implementations are those where significant effort was put into compilers right from the very beginning). In a sense, the RISC philosophy shares the complexity of processor design between the chip architects and the software writers, rather than putting all the burden on the hardware design as in a CISC processor.

The advantages of RISC architecture are by now widely recognized, and there are a large number of RISC processors now commercially available. These include the MIPS R2000 and R3000, the SUN SPARC, the Motorola 88000, the AMD 29000, the Intergraph (formerly Fairchild) Clipper, and the Intel 80960. Moreover there are proprietary RISC chips in use in systems from IBM (the RT personal computer), Apollo (the PRISM family of workstations), and Hewlett Packard. Table 1 summarizes some of the features and measured performance for the current generation of RISC processors.

Furthermore, all of the leading RISC manufacturers have announced plans for higher speed versions of their chips that will meet the performance projections shown in figure 1. To cite some examples:

- 1) The Intergraph Clipper is available in 50 MHz (14 MIPS) versions now, with 20 MIPS expected in March 1989 and 60 MIPS ECL versions in 1990;
- 2) Data General is designing a 100 MIPS ECL 5 chip version of the Motorola 88000 (instruction processor, memory management unit, cache controller, system controller, and system bus interface) expected by 1991;
- 3) Sun has licensed the SPARC technology to LSI Logic, Fujitsu, Bipolar Integrated Technology, and Cypress Semiconductor. All are working on higher speed implementations, with Fujitsu having 25 MHz (15 MIPS) parts available now, 33 MHz in 1989 and 40 MIPS in 1990, LSI with 40-50 MIPS in biCMOS in 1990, and BIT with 40 MIPS ECL in 1989;
- 4) MIPS has licensed the R3000 technology to LSI Logic, Integrated Device Technology, Performance Semiconductor, Siemens, and NEC. 33 MHz (25 MIPS) parts are available now, with IDT, for example, projecting 40 MHz in 1990, 60 MHz in 91 and 100 MHz by 92.

Clearly it is not overly optimistic to expect individual processors with between 50 and 100 VAX equivalents in performance well before the turn-on of the SSC.

Finally, the simplicity of the RISC designs has led to successful implementations in Gallium Arsenide, where limitations on the number of gates has prevented CISC implementations.² Two examples (each of which has the goal of building a 200 MHz or 150 MIPS processor by 1992) are:

- 1) a TI/CDC collaboration (CDC did the chip architecture while TI did the GaAs implementation). This chip has 12,895 gates, 6 pipeline stages, has already run at 68 MHz, and is expected to draw 1W at 200 MHz; and

Table 1. RISC Chips Features Comparison

	R3000	SPARC	CLIPPER	29000	88000
VLSI Chip Count CPU + FPU + MMU + Cache Control	2 CPU/MMU/ Cache + FPU	6 CPU + MMU + Cache + FPU(3 chips)	3 CPU/FPU + ICACHE/MMU + DCACHE/MMU	3 CPU/MMU + FPU + Cache Control	3 CPU/FPU + 2 Cache/ MMU/RAM
Bus Structure	"Harvard" 2 Address/Data (MUXed)	von Neuman 1 Address/Data	"Harvard" 2 Address/Data (MUXed)	"Harvard" 2 Address/Data (MUXed)	"Harvard" 2 Address/Data
Cache Support	64K ICACHE 64K DCACHE	No Direct Cache Support	8K ICACHE 8K DCACHE	No Direct Cache Support	Special Cache RAM w/MMU
MMU Support	Onchip, 64 TLB CAM	No Direct MMU Support	Onchip, 2 x 64 TLB	Onchip 64 TLB	On Cache Chips
Registers	32 General Purpose 16 x 64-bit Floating Point	128 8 x 32 Windows	64 General Purpose 8 Floating Point	192 General Purpose 3 Floating Point	32 General Purpose/ Floating Point
Clock Rate	25 Mhz	16 Mhz	33 Mhz	25 Mhz	20 Mhz
Measured Performance	20 VaxMips	10 VaxMips	5 VaxMips	14 - 17 VaxMips	14 - 18 VaxMips

2) an MDAC (McDonnell Douglas Astronautics Corp.) chip, with 23,178 transistors, a 5 stage pipeline, working 60 MHz versions, and 4-6W expected at 200 MHz.

Both of these chips will require GaAs cache memory to allow 1 memory access every 5 nanosecond cycle, and will require extremely sophisticated compilers to keep pipelines full even in the presence of branches. Most likely, the cost of the chips will prevent them from being in widespread use, but they will be available, for example, for on-line applications requiring the utmost in processing power from a single chip.

RISC AND THE ACP

As an example of what this new RISC technology makes possible for both off-line and on-line applications, I will briefly describe the Fermilab Advanced Computer Program (ACP) group's Second Generation Multiprocessor Project.³ Their original parallel processing system was based on the Motorola 68020, with several systems of more than 100 processors in use. The Second Generation System uses the MIPS R3000 RISC chip set to provide an increase in processing power per board of more than a factor of 20.




Figure 2. The ACP MIPS Processor Block Diagram.

The MIPS chip set was chosen after benchmarking several potential RISC processors. The R3000 was selected primarily because of its superb FORTRAN compiler, which now offers full VMS FORTRAN compatibility and highly sophisticated optimization. The MIPS compiler is an in-house product, and the compiler writers were involved even in the architectural design of the chip, which has given MIPS a lead over other RISC vendors who have up till now relied on outside third parties to provide compilers. The R3000 performs at 15 VAX 11/780 equivalents on a variety of high energy physics reconstruction programs (or roughly 20 times the original 68020's). Note that DEC has also selected the MIPS chips for use in their new line of high performance UNIX workstations (the DECstation 3100).

A VME processor board using the R3000 set has been designed (see figure 2 for a block diagram of the processor board). The board features a 25 MHz R3000 CPU and 25 MHz R3010 Floating Point Unit, 32 KB each of both instruction and data cache, 8 MB of on-board memory with parity, 256 KB of EPROM, a serial port, a full VME master/slave interface, and a memory expansion interface allowing expansion up to 32 MB of memory. Prototypes of this board will be available this spring, and commercial availability should follow shortly.

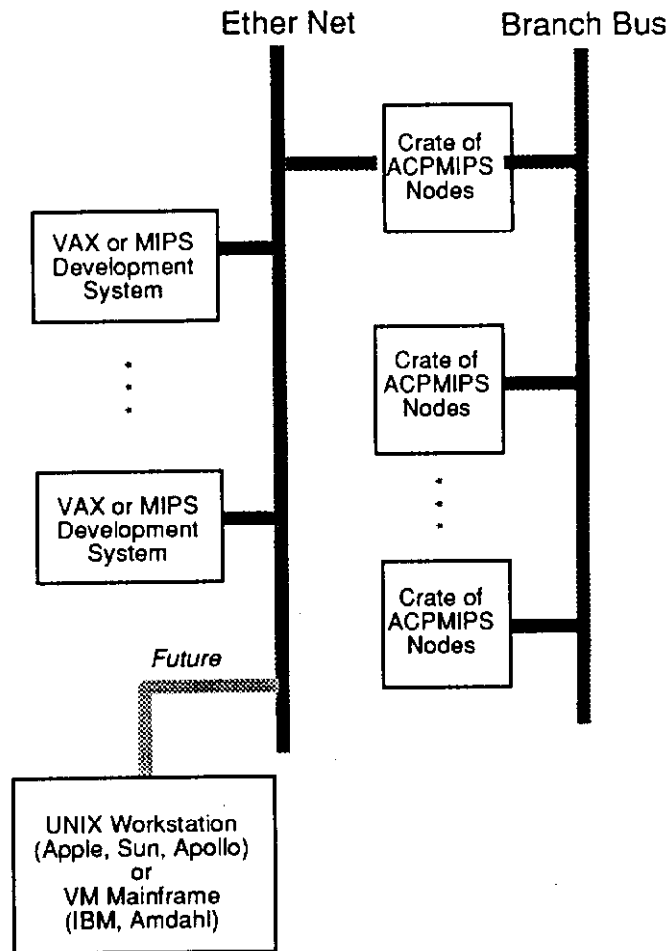


Figure 3. Second generation ACP multiprocessor systems allow any VMS or UNIX processor or workstation to take part.

The board runs the full UNIX operating system, including support for virtual memory, disk and tape I/O, interprocess communication and multi-tasking. A full suite of program development tools (compilers, linkers, librarian, and debugger) also runs on the board. Besides this vastly improved programming environment, the existence of a full-blown operating system also allows more complicated forms of multiprocessing than was allowed in first generation ACP systems, where the processors were all slaves to a dedicated VAX host.

The software support is thus provided in a layered fashion.⁴ Underneath is UNIX, a commercial standard operating system. Next are networking protocols, again using the commercial standards of NFS and TCP/IP. Finally, a set of higher level service routines is provided to:

- 1) transfer blocks of data between processes;
- 2) call a remote subroutine on another process;
- 3) send a datagram message to another process; or
- 4) place and remove processes to and from queues.

Most users will only use these service routines, remaining ignorant of any of the operating system details, but more sophisticated users will also have the full resources of UNIX and the networking protocols available. This reliance on standards also allows any other processors or workstations running UNIX or VMS to be included as full members of a second generation multiprocessing system, providing graphics or other capabilities not present in the VME processors, as shown in figure 3.

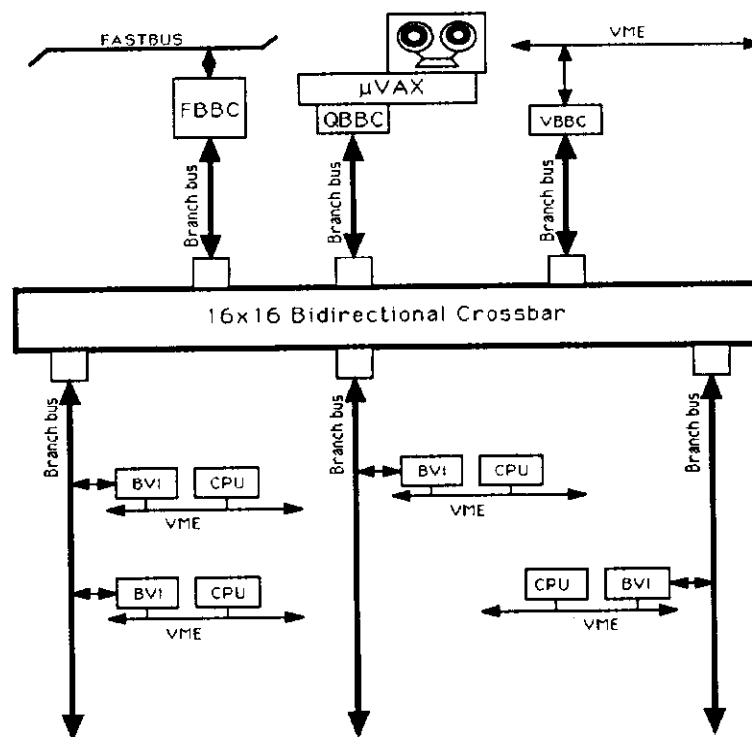


Figure 4. The bus switch used to interconnect data acquisition buses with VME processor farms.

On-line use of these powerful processors is aided by the existence of several interface modules. The ACP Branch Bus (a high-speed parallel data bus allowing transfers at 20 MB/sec) permits VME and other data acquisition buses to be interconnected. FASTBUS (FBBC), VME (VBBC) and Q-Bus (QBBC) interfaces to the Branch Bus exist, as does the Bus Switch, a full 16x16 crossbar switch allowing arbitrary interconnection of Branch Buses. These modules can be combined to provide high performance connections between DA systems and farms of processors, as shown in figure 4. In addition, the memory expansion (XBus) bus on the MIPS processor boards can be used to provide direct access from a DA bus into processor memory (see figure 5), much as the D0 data acquisition system does with microVAXes.

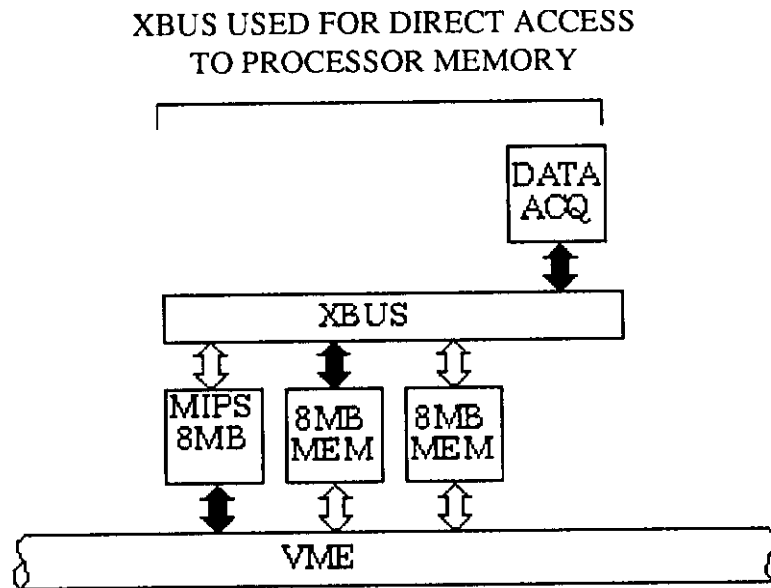


Figure 5. The memory extension bus (XBus) allows direct access into the memory of the ACP MIPS processor board.

DSPS AND OTHER PROCESSORS

Finally, mention should be made of digital signal processors (DSPs). These chips have traditionally only been used for special purpose applications such as dedicated trigger processors, and have not been suitable for more general purpose use due to limited memory space, no floating point, and lack of high level languages and good program development tools.

The current generation of DSPs remedies most of these deficiencies. These new chips typically have full IEEE floating point, large memory address spaces, and speeds of up to 100 MFLOPs. They are supported by operating systems offering high level languages and good program development tools. Examples of such chips' are:

- 1) the Texas Instruments TMS320C30, TI's third generation DSP, now available in a 16 MHz version offering 16 MIPs and 33 MFLOPs of performance. Despite the more than 700,000 transistors on the chip, it sells for under \$100 in quantity. It also supports the SPOX operating system, which provides math libraries, memory management and I/O services together with a small real-time kernel;
- 2) the Motorola DSP96002, available in 27 MHz offering 13.5 MIPs and 40.5 MFLOPs, with hardware hooks to support parallel processing applications;
- 3) the United Technologies UT69532 IQMAC (In-phase Quadrature Multiplier Accumulator), a highly pipelined chip with 3 ALUs and 2 multipliers, offering 100 MFLOPs of performance at 20 MHz.

While not yet appropriate for general purpose processing farms (the RISC processors have the important advantage of allowing code development on workstations or minicomputers using the same chip sets and software tools as the processor farms), these DSPs are highly suited for front end processing and triggering applications. The high level software tools should make these "special purpose" processors much more accessible to the average physicist than ever before, when their coding was normally restricted to a small group of experts.

As a last indication of things to come, I would like to mention Intel's N10 chip, details of which were revealed at the recent ISSCC (International Solid State Circuits Conference). This million transistor chip, originally planned as a co-processor for the forthcoming Intel 80486, instead combines a general purpose RISC core, a high-speed double precision floating point unit, and a 3-D graphics processor. A 50 MHz version of the chip allows 150 MIPs and 100 MFLOPs of performance. No pricing or availability is yet known for this superchip, but it indicates the kind of technology we can expect from industry to be used at the SSC. (On February 27, 1989, Intel formally announced availability of 33 MHz versions of this chip, now known as the Intel 80860, for under \$1000.)

CONCLUSIONS

Industry will provide us with extremely powerful high level language processors for both filtering (level 3) and triggering (level 1-2) applications. We should resist as much as possible the temptation to build hard-wired, non-programmable or microcoded devices. Processor farms of 10^5 VAX equivalents are likely for both off-line and on-line applications: 100-1000 processor boards with 100-1000 VAX power on each board (possibly using multiple processors on each board).

The real challenge will not be in providing the processing power, but rather in insuring that the extraordinarily powerful arrays of processors are actually doing what we want them to do. It is not too early to start developing tools for program specification and

verification. Without these tools, we will be unable to enjoy the full benefits that processor technology can supply.

REFERENCES

- 1 R. Wilson, *Computer Design*, "RISC Architectures Take on Heavyweight Applications", p. 59-79, May 15, 1988.
- 2 B. Cushman, *VLSI Systems Design*, "GaAs Technology meets RISC Architectures", p. 68-77, September 1988.
- 3 T. Nash, H. Areti, R. Atac, J. Biel, A. Cook, J. Deppe, M. Edel, M. Fischler, I. Gaines, R. Hance, D. Husby, M. Isely, M. Miranda, E. Paiva, T. Pham, and T. Zmuda. "High Performance Parallel Computers for Science: New Developments at the Fermilab Advanced Computer Program", to be published in proc. of the Workshop on Computational Atomic and Nuclear Physics at One Gigaflop, Oak Ridge, TN, April 14-16, 1988.
- 4 J. Biel, "Second Generation ACP Multiprocessor System: System Specification Document", July 25, 1988 and revisions; Fermilab Advanced Computer Program internal document; unpublished. Availability of ACP publications may be determined by reference to the file at the HEPNET location: fnacp::acpdoc_root:[docs]doclist.doc
- 5 IEEE Computer Society, *Micro, Digital Signal Processors*, vol. 8, no. 6, December 1988.