

Today's Supercomputers

K. G. Stevens, Jr.

NASA Ames Research Center
Moffett Field, California

Abstract

This paper will discuss the extant supercomputers of 1986. For the purposes of this paper, a supercomputer is a general purpose computer designed to do scientific and engineering calculations at speeds on the order of hundreds of floating point calculations per second. This class of computers includes products from Cray Research, Inc. (CRI), ETA Systems, Hitachi, Fujitsu/Amdahl, and NEC. The basic hardware architectures will be discussed and compared. This will include a discussion of the register structure, instruction set, memory hierarchy, and I/O structure. Performance figures will be given as will cautions about using such figures. The paper will conclude with some comments about the next generation of supercomputers.

What Is A Supercomputer?

There is no generally accepted definition of a supercomputer. Moreover, the concept of a supercomputer must be related to a specific period of time because the supercomputers of the 1970s, 1980s, and 1990s are or will be quite different. Supercomputers are normally considered to be the fastest computers of a particular generation of computers. They are intended to solve a wide range of scientific and engineering applications. Historically they have been used by government laboratories, e.g., Department of Energy, NASA, etc., to solve problems in nuclear energy, fluid dynamics, and meteorology. Today, the class of users and problems has expanded to include oil companies--exploration and reservoir modeling, automobile companies--structural modeling, semiconductor manufacturers--circuit simulation, and movie companies--animation. In the 1970s, the market for supercomputers was measured in tens of systems but has grown to hundreds of systems in the 1980s. This market growth is expected to continue as more industries find that computing is an inexpensive alternative to traditional methods.

For the purposes of this paper we will be considering the extant supercomputers of today, i.e. at least one has been delivered to a customer prior to March 1986. These machines are all expensive general purpose scientific and engineering computers. They are the fastest computers made by the respective companies. Their speeds are all in excess of 100 MFLOPS (millions of floating point operations per second), they have cycle

times of less than 20 nanoseconds, and use architectural features such as pipelining to achieve their speed. The systems which fall into this class are the CRI Cray X-MP and CRAY 2, ETA Systems (CDC) Cyber 205, Hitachi S-810, Fujitsu VP400, and NEC SX-2. Each of these systems will be described in turn as will their performance.

CRAY X-MP

The CRAY X-MP series of computers is the evolutionary successor of the CDC 6600, CDC 7600, CRAY 1, and CRAY 1S. The CRAY X-MP utilizes a register-to-register architecture, i.e. all operands for logical or arithmetic operations must be in a register, and the result is stored in a register. Like its predecessors, the X-MP uses multiple special purpose functional units as shown in figure 1. These functional units are pipelined and can operate concurrently. Unlike the CDC 7600 which only had scalar registers, the Cray has eight vector registers, each 64 words long. The Cray has instructions for manipulating the vectors as units. The registers receive data from and send data to memory. A starting location and an increment are specified to get data from memory, and the increment is not necessarily one. On newer X-MPs there are also instructions for doing scatters and gathers under the control of index lists. Each CRAY X-MP Central Processing Unit (CPU) has two read paths and one write path to memory which is an increase in bandwidth over previous Cray architectures. The main memory of the X-MP can range from 1 million to 16 million 64-bit words and is divided into from 16 to 64 single word banks.

The Cray X-MP has a feature called chaining which permits the linking of one vector operation to another. What this does is create a pipeline which performs multiple operations in a single clock cycle. The cycle time on the X-MP line is 9.5 nanoseconds. Thus, if a single X-MP CPU has chained two operations, it is capable of producing 210 million floating point operations per second (MFLOPS). A four processor X-MP is thus capable of 840 MFLOPS.

In addition to the main memory, the CRAY X-MP has an optional Solid-state Storage Device (SSD). The SSD can provide from 16 to 128 million words of semiconductor memory which is accessed as if it were a disk drive. Transfer between the SSD and main memory is at 1250 Mbyte/sec. The X-MP line also comes standard with an I/O Subsystem which contains two to four I/O Processors and from one to eight

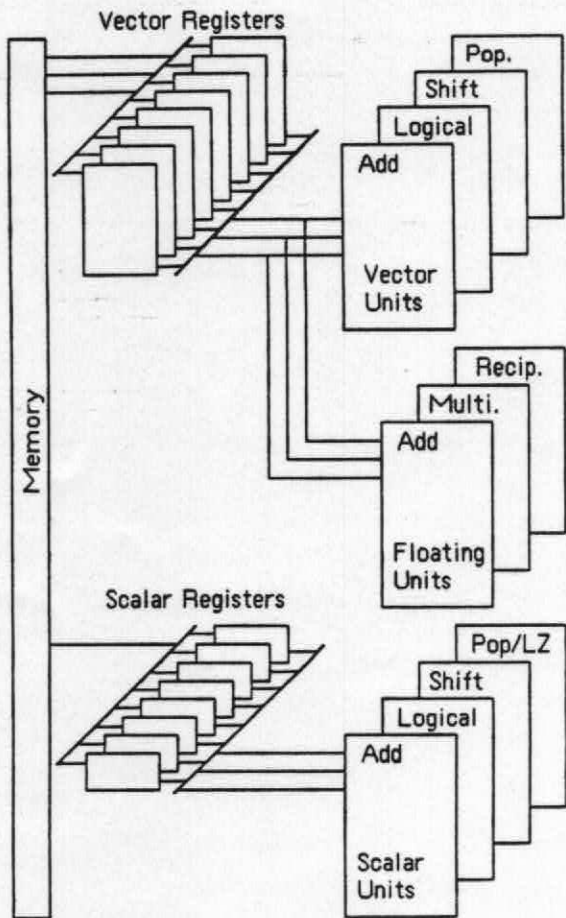


Figure 1. CRAY X-MP CPU

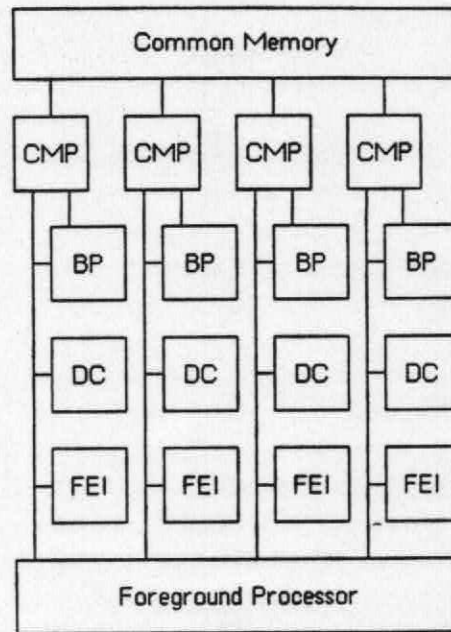
million words of buffer memory. This subsystem interfaces the X-MP with disks, tapes, and front-end processors.

CRAY 2

The CRAY 2 preserves much of the architecture of the X-MP. The clock rate is 4.1 nanoseconds. At this rate the pipelines cannot be chained. Like the X-MP/48, the CRAY 2 has four processors (see Figure 2). It has a main memory of up to 256 million words and no provision for any solid-state secondary memory. The main memory is divided into 128 single word banks. Only 32 banks are accessible to a given processor on any one cycle. The CRAY 2 is capable of performing close to 1 billion floating point operations per second.

CDC/ETA Systems Cyber 205

The Cyber 205 is another example of a vector processor and is the current successor the the CDC Star 100 which was first introduced in the early 1970s. The Cyber 205 (see Figure 3) differs from the CRAY X-MP in that it processes all vector instructions to and from main memory. The vector unit has general purpose pipelines. The Cyber 205 may be configured with 2 or 4 such pipelines. The five segment pipelines are feed data by the stream unit. The vector unit also has a string



- CMP Common Memory Port
- BP Background Processor
- DC Disk Controller
- FEI Front-end Interface

Figure 2. CRAY 2 System Overview

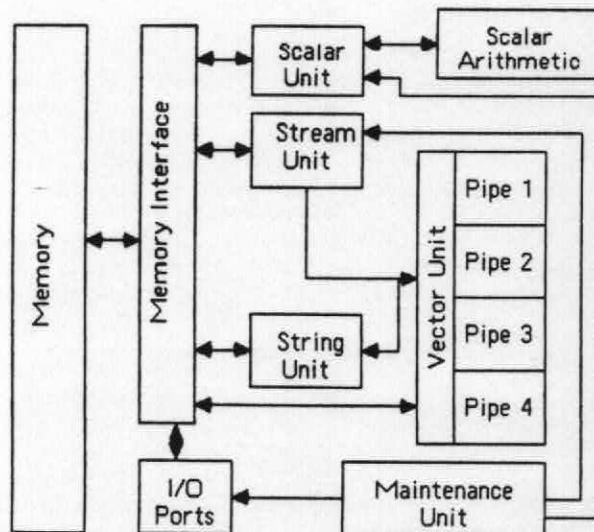


Figure 3. CYBER 205

unit which processes the control vector during streaming operations and is capable of BCD and binary arithmetic, address arithmetic, and logical operations. In addition to the vector unit, there is a scalar unit with 256 scalar registers, functional units for arithmetic, a load and store unit to move data between registers and memory, and the instruction decode unit. There is a memory interface unit and the memory which can be configured with up to 32 million words. The Cyber 205 does not use single-word banks, but superword banks which are eight 64-bit words wide. This memory configuration is necessary because all vector operations are memory-to-memory. The vector startup on a Cyber 205 is significantly longer than that of the Cray because of this memory-to-memory architecture. However, the Cyber 205 can operate on vectors as long as 65,635 consecutive memory locations. Each vector pipeline can perform 100 megaflops for 32-bit results (two 32-bit results every 20 nanoseconds) or 50 megaflops for 64-bit results. The Cyber 205 can perform linked vector multiply/adds at a rate of 800 or 400 megaflops for 32- or 64-bit results, respectively, on a four-pipe configuration.

The memory is a virtual memory with virtual to physical address translation. There are two page sizes, small (1K, 2K, or 8K) and large (64K). The Cyber 205 has 16 I/O channels, each 16 bits wide. The Cyber is typically interconnected to its peripherals and front-end processor via Lossely Coupled Network (LCN), a 50 megabit coax local area network.

Hitachi S-810

The Hitachi S-810 Array Processor System (S-810 for short) is the oldest of the Japanese supercomputers having been announced in 1982. It is available in two models, S-810/10 and S-810/20. The S-810/20 has twice the vector computing rate of the S-810/10 and a large maximum memory capacity. Figure 4 is a block diagram of the S-810/10. The S-810/20 has twice as many vector arithmetic units and paths to memory. This architecture is similar to that of the CRAY CPU. However the S-810/20 has 32 256-word vector registers, 32 scalar registers (only 16 on the S-810/10) and 8 256-word vector mask registers. Each of the vector units is independent. The maximum performance of the S-810/20 is 630 MFLOPS. The main memory capacity is 256 megabytes and the extended memory capacity is 1024 megabytes with a maximum transfer rate between the two memories of 1000 megabytes per second. This system has a number of 3 megabyte per second disk channels with a maximum I/O processor capacity of 48 megabytes per second. It is possible for the S-810 to perform as a stand-alone system or be front-ended by a Hitachi M-Series computer. In the latter case there is a channel to channel connection between the two CPUs and they share direct access storage devices (DASD).

Fujitsu VP Series

Fujitsu announced the FACOM vector processors VP-100 and VP-200 in July, 1982, and the VP-50 and VP-400 in 1985. These computers are marketed in the United States by Amdahl Corporation. Figure 5 is a block diagram of the VP-200. Like the CRAY X-MP,

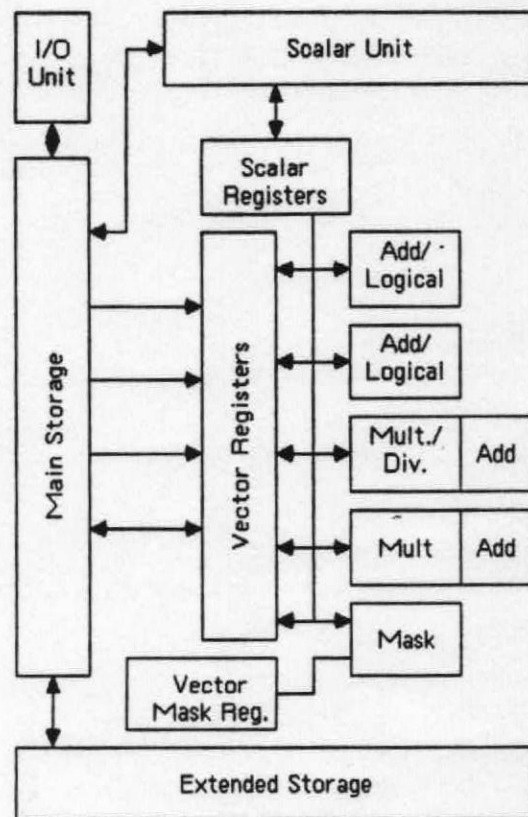


Figure 4. HITACHI S-810 (Model 10)

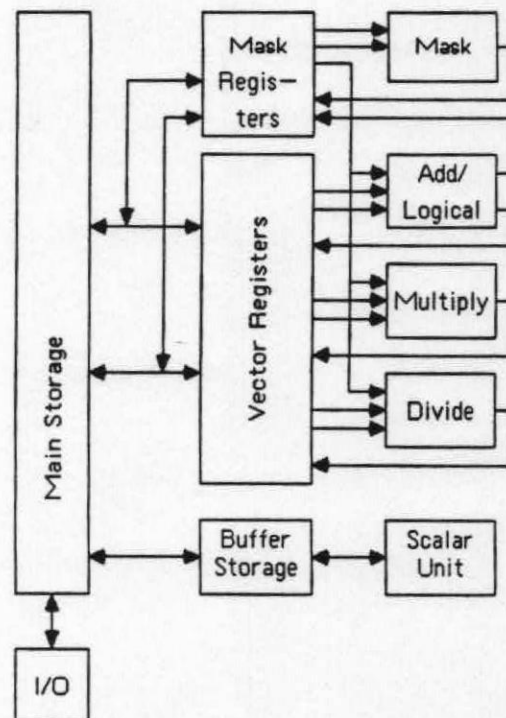


Figure 5. VP-200

the system has a scalar processor and a vector processor which can operate concurrently. The scalar unit has 16 general purpose registers, 8 floating point registers, and 64K bytes of cache memory. The vector unit consists of two paths to memory, a set of vector registers, a mask unit, and 3 arithmetic pipelines. Each path to memory has a data bandwidth of 32 bytes every 15 nanoseconds in either direction. The vector registers total 64K bytes on the VP-200 and can be used as 256 32-word registers, 128 64-word registers, 64 128-word registers, ..., 8 1024-word registers depending on the contents of a programmable register. For the VP-200, the throughputs of add/logical pipe and multiply pipe are 267 MFLOPS each, whereas the divide is 38 MFLOPS. Thus, 533 MFLOPS is the maximum speed when the add and multiply are active concurrently. In order to control conditional vector operations and vector editing operations, bit-string masks are stored in the 256 Mask Registers and the operations are performed by the Mask Pipe. The maximum capacity of the main memory is 256 megabytes interleaved into 256 banks for the VP-200 and half that capacity and number of banks for the VP-100. The more recently introduced VP-400 has twice the vector capacity of the VP-200 while the VP-50 has one half the capacity of the VP-100.

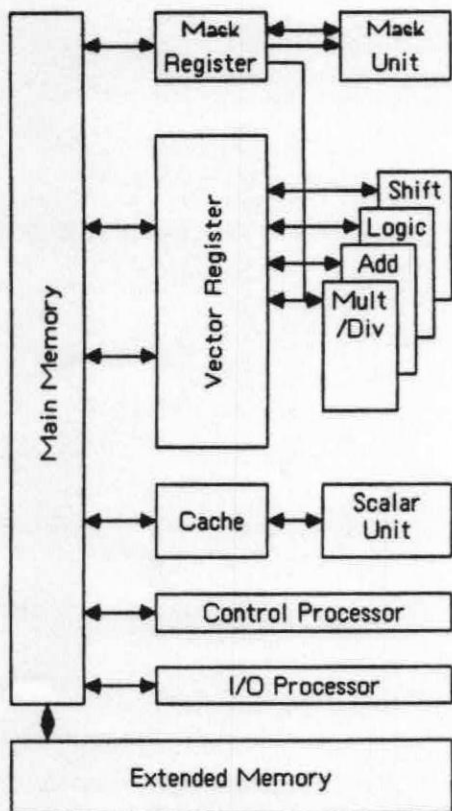


Figure 6. SX-2

NEC SX-1/SX-2

The NEC SX-1/SX-2 series of machines is the most recently announced of the Japanese supercomputers. As shown in Figure 6 the NEC system has 4 different vector arithmetic pipelines; add, multiply and divide, logical, and shift. In the case of the SX-2, there are four sets of these pipe which operate in parallel. There are only two sets in the SX-1. When performing a multiply add combination on the SX-2, four adds and four multiplies are produced every six nanosecond clock cycle for a maximum speed of 1300 MFLOPS. Vector operations are register to register from a set of 40 256-word vector registers. As with the other register-to-register architectures there are a set of mask registers and a mask unit to perform vector editing functions. The scalar unit has 128 registers as well as a cache. The main memory of the SX-2 is composed of 512 banks with a capacity of 256 Megabytes. NEC also can provide an extended memory of up to 2048 megabytes with an intermemory transfer rate of 1300 megabytes per second. The input/output processor has a maximum capacity of 50 megabytes per second through up to 32 channels.

Performance Comparison

No single performance comparison is adequate to compare these machines because their relative performance is very dependent on the benchmark being performed. Figure 7 plots the maximum speed of three generations of super computers. When actually running applications programs the speeds will be much slower and vary widely depending on the application. Table 1 gives the speeds of these supercomputers in MFLOPS for the Livermore Loops which were derived from several codes in use at Lawrence Livermore National Laboratory in the 1970s. Again, caution should be taken when using these numbers because they may or may not reflect how other applications would perform on these supercomputers. In particular, the numbers for the




Figure 7. Plot of Maximum Speeds vs Year

Loop	CRAY X-MP*	CRAY 2*	Cyber 205	S-810	VP-200	SX-2
Hydro excerpt	115	167	122	298	331	745
MLR. inner product	138	69	13	329	180	418
Inner product	136	91	79	332	338	553
Banded linear equation	48	33	3	113	88	123
Tri-diagonal elimination (below)	11	9	8	36	10	14
Tri-diagonal elimination (above)	9	8	5	36	10	15
Equation of state excerpt	165	213	146	349	331	795
P. D. E. integration	91	80	22	155	90	153
Integer predictors	123	170	81	332	261	638
Difference predictors	64	56	30	62	86	122
First summation	8	4	8	35	5	24
First difference	76	53	73	111	115	251
2-D particle pusher	5	5	3	10	6	8
1-D particle pusher	7	9	7	11	14	24
Average	74	69	43	158	133	277

* Single CPU only

Table 1. Livermore Loops in MFLOPS

CRAY X-MP and CRAY 2 reflect the speed of only a single CPU not the multiple CPUs which could be used on an application.

FUTURE TRENDS

The primary trend in supercomputers as defined in this paper will be to the use of multiprocessor architectures. This trend has been followed by CRAY in its current offerings and there are indications that the number of processors will increase in subsequent generations of supercomputers from Cray Research, Inc. ETA Systems, the manufacturer of follow ons to the Cyber 205, has also indicated that its future products will be multiprocessors. This trend is necessary because the gains in speed which can be achieved with new technology and architectural variations on a single vector processor are not sufficient for a new generation of supercomputer. The companies which have been discussed previously in this paper will tend to start with a small number of processors, e.g. 4, and work themselves up to larger numbers of processors. There are quite a number of startup companies which are jumping in to the multiprocessor business with larger, e.g. 256 to 1024, numbers of processors. These companies like Ncube, Ametek, etc. may or may not become important contenders in the supercomputer market place. Another trend is that supercomputers will have more and more memory with the next generation being measured in gigabytes and gigawords. And finally, the market for supercomputers will continue to increase as more and more industries find them to be a cost effective alternative to the way that they are currently doing business.

Conclusions

Today's super computers can be grouped into one of three classes: single CPU register-to-register vector architectures (Japanese machines), multiprocessor register-to-register vector architectures (CRAY), and single CPU memory-to-memory vector architectures (CDC/ETA). The next class which will be

added will be the multiprocessor memory-to-memory vector architectures (ETA). Another class which may enter the supercomputer category is the large number of scalar processor multiprocessor architecture. This class of machines currently lacks mature software and considering that it took at least 10 years for vector software to fully mature it will be a while before this software is available. Within a class there are only subtle difference whose effect on performance are difficult to predict. Thus, with today's supercomputers it is a clear case of *caveat emptor* with price being an important factor in a highly competitive market place.

References

1. Arnold, C. N., "Performance Evaluation of Three Automatic Vectorizer Packages", 1982 International Conference on Parallel Processing, IEEE Computer Society Press.
2. Furukatsu, T., Watanabe, T., and Kondo, R., "NEC Supercomputer SX System", *Nikkei Electronics*, 1984, 11, 19 pp 237-272.
3. Hockney, R. W. and Jesshope, C. R., *Parallel Computers*, Adam Hilger Ltd, Bristol, 1981.
4. Hwang, K. and Briggs, F. A., *Computer Architecture and Parallel Processing*, McGraw Hill, New York, 1984.
5. Kawabe, S., Kanagawa Works, Hitachi, Ltd., personal communications.
6. Kowalik, J. S., *High-Speed Computation*, Springer-Verlag, Berlin, 1984.
7. Nelson, H and McMahon, F., Lawrence Livermore National Laboratory, personal communication, 1985.
8. Schulbach, C. H., *Evaluation of Existing and Proposed Computer Architectures for Future Ground-Based Systems*, AIAA-85-3018.