
ANL/MCS-TM-57, Rev. 2

ANL/MCS-TM--57-Rev. 2
DE90 001045

ARGONNE NATIONAL LABORATORY
9700 South Cass Avenue
Argonne, Illinois 60439

Advanced Architecture Computers

by

*Jack J. Dongarra** and *Iain S. Duff***

Mathematics and Computer Science Division

Technical Memorandum No. 57 (Revision 2)

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

September 1989

* Electronic mail address: dongarra@mcs.anl.gov.

** Permanent address: Computer Science and Systems Division, Building 8.9, Harwell Laboratory, Oxfordshire OX11 ORA, England. Electronic mail address: na.duff@na-net.stanford.edu.

This work was supported by the Applied Mathematical Sciences subprogram of the Office of Energy Research, U.S. Department of Energy, under Contract W-31-109-Eng-38.

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

DISCLAIMER

Portions of this document may be illegible in electronic image products. Images are produced from the best available original document.

Contents

Section	Page
Abstract	1
1. Introduction	1
2. Summary and Classification of Machines Considered	2
3. Supercomputers	3
4. Minisupercomputers	5
5. Enhanced Mainframes	5
6. Parallel Machines	6
7. High-Performance Graphics Workstations	7
8. Template for Machine Description	7
Alliant FX Series	9
Amdahl Vector Processors	12
AMETEK Series 2010	15
Active Memory Technology DAP	17
Ardent Titan	20
BBN Butterfly GP1000 and TC2000 Parallel Processor	23
CDC CYBER 180 990E/995E	26
CDC CYBER 205	28
CONVEX C-120/130; C-210/220/230/240	30
CRAY-1	33
CRAY-2	34
CRAY-3	36
CRAY X-MP	37
CRAY Y-MP	38
CULLER 8	40
CDC CYBERPLUS	41
Cydrome CYDRA 5	43
ELXSI 6400	44
Encore Multimax 310/320	47
ETA-10	48
FLEX/32 Multicomputer	50
Floating Point Systems MP32 Series	52
Floating Point Systems FPS-5000 Series	54
Floating Point Systems M64/40, 140, 50, 60	56
Floating Point Systems M64/20, 30	58
Floating Point Systems FPS T Series	60

Section	Page
Galaxy YH-1	62
Loral MPP (Goodyear)	63
Gould NP1	64
Hitachi S-810, 820	66
IBM 3090/VF	68
Intel iPSC/2	71
IP-1	74
Isis	76
Kendall Square Research	78
Loral System 500	79
Meiko	80
Multiflow TRACE 7/200, 14/200, 28/200	83
Myrias	86
National Advanced Systems AS/91X0	88
NCUBE	89
NEC SX-JA, 1EA, 1A, 2A	91
Prevec	94
PS 2000	95
Saxpy MATRIX 1	96
SCS-30/XM, 40/XM	98
Sequent Balance 8000, 21000	101
Sequent Symmetry S27, S81	103
Silicon Graphics IRIS 4D/70	106
Star Technologies VP Series	108
Stellar GS1000	110
Supertek S-1	113
Thinking Machines Connection Machine CM-2	114
Unisys ISP 1100/90	115

Advanced Architecture Computers

by

Jack J. Dongarra and Iain S. Duff

Abstract: We describe the characteristics of several recent computers that employ vectorization or parallelism to achieve high performance in floating-point calculations. We consider both top-of-the-range supercomputers and computers based on readily available and inexpensive basic units. In each case we discuss the architectural base, novel features, performance, and cost. We intend to update this report regularly, and to this end we welcome comments.

Keywords

vector processors, array processors, parallel architectures, supercomputers, high-performance computers

1 Introduction

In the past few years several machines have been announced that use some form of parallelism to achieve performance in excess of that attainable directly from the underlying technology of the constituent chips. To a large degree the availability of low-cost chips as building blocks has given rise to many of these new machines.

After listening to numerous technical and sales presentations on these new computers, we became overwhelmed and confused with the characteristics of each product and its relative strengths and weaknesses. In an effort to clarify these issues - both for ourselves and for other computational scientists - we have written this report summarizing the range of machines available, the architectures employed, and the principal features of each machine.

In Section 2, we list the computers considered and discuss the criteria we have used to select them. We present a rough classification based on architectural features and their niche in the marketplace. This classification divides the machines into five categories: supercomputers, minisupercomputers, vector add-ons or vector-assisted mainframes, parallel processors, and high-performance graphics workstations. Each category is discussed in turn in Sections 3 through 7.

The guidelines used in preparing the detailed descriptions are given in Section 8. In some cases, our data are incomplete and nonuniform. This situation reflects the technical level of the presentations, the documentation available to us, the stage of development of the product being described, and the comments received from vendors on draft copies of our document. We welcome comments and criticisms that might help to remedy any deficiencies. This report is a second edition. We intend to continue updating this report to reflect both the changing marketplace and further information on currently listed machines.

2 Summary and Classification of Machines Considered

In the past few years there has been an unprecedented explosion in the number of different computers in the marketplace. This explosion has been fueled partly by the availability of powerful and cheap building blocks and by the availability of venture capital. There have been two main directions to this explosion. One has been the personal computer and workstation market, and the other the development and marketing of computers using advanced architectural concepts. In this report we restrict our study to the latter group, with particular interest in architectures that use some form of parallelism to increase performance over that of the basic chip.

We also restrict our attention to machines that are available commercially, and thus exclude research projects in universities and government laboratories and products still at the design stage. We would, however, welcome being alerted to ongoing activities.

We have necessarily had to exclude information obtained under non-disclosure agreements. We will update this report as such information is released through product announcements.

A much-referenced and useful taxonomy of computer architectures was given by Flynn (1966). He divided machines into four categories:

- (i) SISD - single instruction stream, single data stream
- (ii) SIMD - single instruction stream, multiple data stream
- (iii) MISD - multiple instruction stream, single data stream
- (iv) MIMD - multiple instruction stream, multiple data stream

Although these categories give a helpful coarse division, we find immediately that the current situation is more complicated, with some architectures exhibiting aspects of more than one category.

Many of today's machines are really a hybrid design. For example, the CRAY X-MP has up to four processors (MIMD), but each processor uses pipelining (SIMD) for vectorization. Moreover, where there are multiple processors, the memory can be local, global, or a combination of these. There may or may not be caches and virtual memory systems, and the interconnections can be by crossbar switches, multiple bus-connected systems, time-shared bus systems, etc.

We thus choose a method of subdividing and classifying the machines different from that used in our original report (Dongarra and Duff 1987). As before, we identify the supercomputers separately and discuss these in Section 3. However, we split the other machines according to their niche in the marketplace rather than their connectivity or mode of data access or data transfer. Minisupercomputers can be defined as junior versions of supercomputers that offer a similar interface to the larger machines but with lower performance and reduced costs. We consider machines in this class in Section 4. Some powerful vector computers do not fall into either of the previous classes but are based on an enhancement to a mainframe computer through the addition of an array processor or an integrated vector facility. We discuss both types of computer in Section 5. In Section 6, we consider machines that rely primarily on parallelism rather than pipelined vector processing, and divide these into two categories depending on whether we regard them as good experimental vehicles for studying parallelism and parallel algorithms or whether we

consider them as potential supercomputers of the future. In Section 7 we summarize the high-performance graphics workstations that do not themselves qualify for the previous categories but that are clearly in a different class from regular top-of-the-line workstations.

3 Supercomputers

Supercomputers are by definition the fastest and most powerful general-purpose scientific computing systems available at any given time. They offer speed and capacity significantly greater than mainframe computers, defined as top-of-the-range widely available machines built primarily for commercial use. The term supercomputer became prevalent in the early 1960s, with the development of the CDC 6600. That machine, first marketed in 1963, boasted a performance of 1 Megaflops (millions of floating-point operations per second).

During the next 15 years, the peak performance of supercomputers grew at an rapid rate; and since 1980, that trend has accelerated. The 1995 machine is expected to have a maximum speed of 200 Gigaflops, more than 200,000 times that of the CDC 6600 (see Table 1).

Table 1. Performance Trends in Scientific Supercomputing

Year	Machine	Speed	Speed Increase	
			10 years	20 years
1963	CDC 6600	1 MFLOPS	-	-
1969	CDC 7600	4 MFLOPS	4	-
1979	CRAY-1	160 MFLOPS	100	-
1983	CYBER 205	400 MFLOPS	100	400
1986	CRAY-2	2 GFLOPS	500	2000
1990-1995	-	200 - 1000 GFLOPS	1000	250,000

Many companies have devoted their resources to producing the fastest and most powerful machines on the market. Their strategy has been to develop a few state-of-the-art machines that enable scientists and engineers to tackle problems previously considered computationally infeasible. From these commercial ventures we have seen the development of vector and, more recently, parallel computers capable of solving complex numerical and nonnumerical problems. The second generation, with higher speed and more parallelism, is already under development. In Table 2, we summarize the currently available supercomputers.

Table 2. Supercomputers

Machine	Maximum Rate, in MFLOPS	Memory, in Mbytes	OS	Number of Processors
CRAY-1	160	32	Own	1
CRAY X-MP	941	512	Own/UNIX	4
CRAY Y-MP	2667	256	Own/UNIX	8
CRAY-2	1951	4096	UNIX	4
CYBER 205	400	128	Own	1
ETA-10G	5714(a)	2048(b)	UNIX/VSOS	8
ETA-10E	3810(a)	2048(b)	UNIX/VSOS	8
ETA-10Q	526(a)	512(b)	UNIX/VSOS	2
Fujitsu VP-400E	1714	1024	Own	1
Fujitsu VP-200E	857	1024	Own	1
Fujitsu VP-100E	429	1024	Own	1
Fujitsu VP-50E	286	1024	Own	1
Fujitsu VP-30E	133	1024	Own	1
Hitachi S-820/80	2000	512(c)	Own	1
Hitachi S-810/20	857	512(c)	Own	1
NEC SX-2A	1300	1024(d)	Own	1
NEC SX-1A	650	1024(d)	Own	1
NEC SX-1E	324	1024(d)	Own	1

(a) for 64-bit processing on 2 pipelines with linked triad and overlapped scalar processing

(b) Also 16 MWord (128 Mbyte) local memory for each processor

(c) Also a 12-Gbyte extended memory

(d) Also a 8-Gbyte extended memory

The actual price of the systems in Table 2 depends on the configuration, with most manufacturers offering systems in the \$5 million to \$20 million range. All use ECL logic with LSI, except the CRAY X-MP, the CRAY-1 in SSI, and the ETA-10 in CMOS ALSI (Advanced Large Scale Integration), and all use pipelining and/or multiple functional units to achieve vectorization/parallelization within each processor. Cray is the only supercomputer manufacturers to offer multiple-processors machines, although other vendors have announced multiprocessor machines for delivery in the near future. The form of synchronization on the Cray machines is essentially event handling. Both Fujitsu and Hitachi systems are IBM System 370 compatible. We have included the CRAY-1 computer in the above table largely as a benchmark since it could not now be considered a supercomputer in terms of performance and is no longer manufactured by Cray. The Fujitsu machines are marketed in Europe and North America by Amdahl (the 500E to 1400E range) and by Siemens (the VP-50 to 400 range).

4 Minisupercomputers

Below the supercomputer market, a new class of near-supercomputers or minisupercomputers has emerged. These systems typically feature strong vector or advanced scalar capabilities and have been utilized for traditional high-performance technical computing applications. Priced well under supercomputers, \$100,000 to generally no more than \$1 million, minisupercomputers are frequently sold when budgets are limited to this price range or when stand-alone capabilities are required. Early leaders in the field of minisupercomputing were Alliant, Convex, and Scientific Computer Systems. More recently, this market has experienced high growth, and many new products and companies have emerged, including Multiflow, and Gould (see Table 3).

Table 3. Minisupercomputers

Vendor	Theoretical Peak Performance Mflops (64 bits)	LINPACK Performance Mflops	First Shipment
Alliant FX/8	94	7.6	1985
Alliant FX/80	188	8.5	1987
Astronautics	90	7.1	1988
Convex C1	20	7.3	1984
Convex C2	200	16	1987
FPS 500	—	—	1988
Multiflow Trace 28/200	60	10	1987

5 Enhanced Mainframes

An alternative in the near-supercomputer category is the add-on array processor. Companies such as Floating Point Systems, and Star Technology are actively marketing these add-on products in an effort to attract current supercomputer users.

In a related vein, vector-processing enhancements are now being marketed for commercial mainframes. These vector enhancements allow machines produced for general-purpose applications to offer users increased numerical capability. In some cases, the ability to apply vectors is extended to more than one processor in multiprocessing mode. Companies currently offering such vector-processing capabilities include Control Data, Hitachi (marketed in the West by NAS and COMPAREX), Honeywell, IBM, and UNISYS.

We summarize some of the machines in this category in Table 4.

Table 4. Power-assisted mainframes

Machine	Maximum Rate, Mflops	Memory, Mbytes	OS	Number of Processors
CDC 180 990	125	256	NOS/VE	1-2
FPS M64/140	187	128	Own	1
IBM 3090S/VF	696	256 (a)	Own	1 - 6
NAS AS/91X0	?	64	Own	1 or 2
Unisys 1190/ISP	266	128	Own	1,2,4 (c)

(a) Also a 2-Gbyte extended memory

(b) In 32-bit arithmetic

(c) Only 1 or 2 ISPs can be attached

6 Parallel Machines

While most of the supercomputers and minisupercomputers utilize vector processing to provide performance, a number of new companies are developing parallel processing systems. Such systems range from smaller (8- to 30-processor) machines like the Sequent or Encore to massively parallel (16,384-processor) systems like the Thinking Machines CM-2. Others in this area include Floating Point Systems, Myrias, BBN Advanced Computing, and DEC; and they may be joined soon by IBM, which has indicated that it will offer a product in this category by 1989.

While it certainly true that the parallel architectures fall into two camps depending on whether or not they are potential supercomputers, it is less easy to assign a particular machine to one of these classes. We have, however, made a partly subjective judgment and compare the parallel architectures in two tables. Table 5 summarizes those parallel architectures that are designed for experimentation with parallel constructs, and Table 6 lists machines with potential for future elevation to the status of a supercomputer.

Table 5. Experimental parallel machines

Machine	Chip	Max. Parallelism	Connection
Elxsi 6400	ECL	12	bus
Encore Multimax	32332/32081	20	bus
Flex/32	32032/32081	20	bus
Sequent Symmetry S81	80386/80387	30	bus

Table 6. Potential supercomputers

Machine	Chip	Parallelism	Connection
Active Memory (DAP)	CMOS	4096 (SIMD)	near-neighbor
BBN Butterfly TC 2000	88000	256	Banyan network
CYBERPLUS	Own	256	ring
Intel iPSC/2	80386/80387	128	hypercube
IP-1	Own	33	cross-bar
Meiko	Transputer	No limit (a)	user-configurable
Myrias SPS-2	68020/68882	512 minimum	hierarchical bus
NCUBE	VLSI	1024	hypercube
TMC CM-2	VLSI	65536 (SIMD)	hypercube

(a) Maximum system delivered to date has 1024 processors

Because of the widely differing architectures of the machines in Tables 5 and 6, it is not really advisable to give one or even two values for the memory. In some instances there is an identifiable global memory; in others there is a fixed amount of memory per processor. Additionally, it may be possible to configure memory as either local or global. A value for the maximum speed is even less meaningful than in the previous tables, since a high Megaflop rate is not necessarily the objective of those machines and the actual speed will depend on the algorithm and application.

7 High-Performance Graphics Workstations

Finally, the supercomputer market has been expanded by the introduction of supercomputing workstations, such as those from Apollo, and single-user high-performance graphics systems such as those from Apollo, Ardent, Stellar, and Silicon Graphics. We summarize these machines in Table 7.

Table 7. High-performance graphics workstations

Machine	Chip	Peak performance, Mflops	Memory, Mbytes
Apollo DN10000	Own	?	?
Ardent TITAN	MIPS/Weitek	64	128
Silicon Graphics IRIS GT	MIPS/Weitek	100	16
Stellar GS2000	Own/Weitek	80	128

8 Template for Machine Description

As we mentioned in the introduction, the level of technical information on each machine varied significantly. We have, however, attempted to organize the available information in

a consistent manner. In Table 8, we give the template used in presenting the data in the appendixes.

Table 8. Template for Description of Machines

Name of machine, manufacturer, backers, etc.
Architecture
Basic chip used
Local, shared memory, or both
Connectivity (for example, grid, hypercube)
Range of memory sizes available; virtual memory
Floating point unit (IEEE standard?)
Configuration
Stand-alone or range of front-ends
Peripherals
Software
UNIX or other?
Languages available
Fortran characteristics
F77
Extensions
Debugging facilities
Vectorizing/parallelizing capabilities
Applications
Run on prototype
Software available
Performance
Peak
Benchmarks on codes and kernels
Status
Date of delivery of first machine, beta sites, etc.
Expected cost (cost range)
Proposed market (numbers and class of users)
Contact: technical and sales

Alliant FX Series

Vector Register, Parallel, Shared-Memory Architecture

Formerly, the company was called Dataflow.

Architecture: Computational elements (CEs) execute applications code using vector instructions. The CEs transparently execute the code of an application in parallel. CE: Weitek 1064/1065 plus ten different gate array types with 2600 to 8000 gates. First-generation computational elements (FX1, FX4, FX8) may be added in the field, increasing performance without recompilation or relinking. Advanced Computational Elements (ACEs) for second generation (FX40, FX80, VFX) are based on the BIT floating point chips. Each CE has 8 vector registers, each with 32 64-bit elements, and 8 64-bit scalar floating point, 8 32-bit integer, and 8 32-bit address registers.

Interactive Processors (IPs) execute operating system, interactive code, and I/O operations. An FX/1 has 1-2 IPs. An FX/4 and FX/40 have 1-6 IPs. An FX/80 has 1-12 IPs.

Basic chips used: IP. Motorola 68020. 4 Mbyte local memory in each IP. ACE 64-bit processor 20,000 gate CMOS VLSI gate array, with BIT floating-point processors. 64 Kbyte instruction cache.

The cycle time is 170 nsec. Only six different PC boards are used.

CEs are cross-bar connected on the backplane to a 512 Kbyte write-back computational processor (CP) cache (FX/80). Bandwidth is 376 Mbyte/sec.

Each 32-Kbyte IP cache is connected to 1-3 IPs (FX/80) or 1-2 IPs and a CE (FX/1). The FX/80 has 1-4 IP caches; the FX/4 and FX/40 have 2 IP caches; the FX/1 has one IP cache.

The CP and IP caches are attached by two 72-bit buses to the main memory. Memory bus bandwidth is 188 Mbyte/sec, and memory cycle time is 85 nsec.

Connectivity: crossbar (CE to cache), bus (cache to memory, cache to cache)

Range of memory sizes available: 32-64 Mbytes (FX/1), 32-160 Mbytes (FX/4 and FX/40), and 32-256 Mbytes (FX/80), using 1 Mbit chips with ECC.

Virtual memory: 2 Gbytes per process.

Floating-point unit: IEEE 32- and 64-bit formats including hardware divide and square root and microcoded elementary functions.

Configuration: Stand-alone. TCP/IP network support.

Size: FX/1 system - 28" x 13" x 25" (the FX/1 I/O expansion cabinet is the same size); FX/4, FX/40, and FX/80 systems - 43.5" x 29.5" x 33.8" (the I/O expansion cabinet is 24.5" and same height and depth, while the tape cabinet is 61" in height).

Cooling: All systems are air-cooled.

FX/1 : 1155 Watts (max. configuration), 725 Watts (I/O Expansion)

FX/4 : 4500 Watts, 2100 Watts (I/O Expansion)
FX/40 : 4200 Watts, 2100 Watts (I/O Expansion)
FX/80 : 5100 Watts, 2100 Watts (I/O Expansion)

Peripherals:

- 800/1600/6250 BPI start-stop tape drive
- 550 Mbyte (formatted) Winchester disk drives
- 45 Mbyte cartridge tape drive
- Floppy disk drive
- 8/16 line multichannel communications controllers
- 600 lpm printer
- Ethernet controller

Software: Concentrix, Alliant's enhancement of Berkeley 4.2 UNIX with multiprocessor support.

Languages: Fortran, C, Pascal, Ada, Lisp, STSC APL, 68020 Assembler

Fortran characteristics:

- F77 - Conforms to 1978 ANSI standard.
- Extensions - Most of VAX/VMS extensions and Fortran 8x array extensions. Debugging facilities.
- Vectorizing/parallelizing capabilities - Automatic detection of vectors and feedback to user via diagnostic messages.
- Can employ COVI (concurrent outer, vector inner) on nested loops.
- User control of transformations via directives in the form of Fortran comments
- Interprocedural dependency analysis for automatic determination of parallel subroutine calls in loops.

Performance: Advanced CE's (ACE's).

Scalar 32-bit : 7.2 mips / CE (14700 Kwhetstones).

Scalar 64-bit : 6.2 mips / CE (13700 Kwhetstones).

Vector 32-bit and 64-bit : 23.5 Mflops / CE.

FX/80 on 1,000 x 1,000 LINPACK benchmark: 69.3 Mflops. Peak performance 188.8 Mflops.

Applications: Engineering and scientific end-user and OEM applications, stand-alone or as a computational server to a network of engineering workstations.

Status: First beta delivery May 1985; first production shipment September 1985. Alliant's customers include Asahi Chemical Corp, AT&T, Boeing Airplane Co., Ford Motor Co., Hughes Aircraft Corp., Motorola Inc., Siemens, The Whittle Laboratory at the University of Cambridge, CERFACS at Toulouse, and the Jodrell Bank Observatory at the University of Manchester.

Entry level package prices:

FX/1 : discontinued; FX/4 : \$99,900; FX/40 : \$149,000; FX/80 : \$299,000
ACEs are priced at \$59,000 each

Contact:

Alliant Computer Systems Corp.
1 Monarch Drive
Littleton, MA 01460
508-486-4950

President: Ron Gruner
Technical: Craig J. Mundie, Vice President Business Development
Sales: Roger Parsons, Vice President Worldwide Sales

Alliant Computer Systems UK Ltd
10 Heatherley Road
Camberley
Surrey
GU15 3LW
England
0276-682765
FAX 0276-65235

John Harte, President of European Operations

Chic McGregor, Sales

Jane Doorly, Systems & Applications

Amdahl Vector Processors (Fujitsu VP) Vector Register Architecture

Architecture: The Amdahl 500, 1100, 1200, and 1400 Vector Processors are marketed by Amdahl Corp. in the U.S., Canada, Europe, and the Pacific Basin. These products are manufactured by Fujitsu, and similar models are marketed in Japan as the VP-50, VP-100, VP-200, and VP-400. The VP-100 and 200 is also marketed by Siemens in mainland Europe. In 1987, the range was upgraded by the addition of E models. The main change was to the functional pipes.

These are all register-to-register machines. All models have one scalar and one vector unit which can execute computations independently. The scalar unit fetches all instructions and passes each instruction to the appropriate unit for execution. The scalar processor is based on the Fujitsu M380/382 series mainframes and runs the IBM S/370 extended architecture instruction set. A recent Amdahl proprietary software program product, called VP/XA, allows Amdahl vector processors to run current MVS/XA releases, and permits Amdahl supercomputers to use standard operating environments.

Configuration: The vector unit consists of 5 or 6 pipelines, a vector register memory, and a mask memory. The 5 or 6 pipelines comprise 1 or 2 load/store pipelines, plus 1 mask pipeline, 1 add/logical pipeline, 1 multiply pipeline, and 1 divide pipeline. In the E models, the multiply pipe is replaced by a multifunctional pipe for floating-point addition, multiplication, or concurrent multiplication/addition. The number of concurrent pipelines, vector register size, and mask register size differ for each model, as shown below. Main memory capacity ranges from 32 Mbytes to 1024 Mbytes (4 to 128 M 64-bit words).

Configuration	Model			
	500	1100	1200	1400
# pipes total	5	6	6	5
# concurrent load/store pipes	1	2	2	1
# 64 bit words/vect cyc/pipe	1	1	2	4
Scalar cycle time (nsec)	14	14	14	14
Vector cycle time (nsec)	7	7	7	7
# concurrent arith pipes	2	3	3	3
# 64-bit results/vect cyc/pipe	1	1	2	4
Vect. reg. size (Kbytes)	32	32	64	128
Mask reg. size (Bytes)	512	512	1024	2048
Max. main memory (Mbytes)	512	512	1024	1024
Min. main memory (Mbytes)	32	32	64	64
Max. interleaving (ways)	128	128	256	256

The total vector register capacity is 32-128 Kbytes. The registers can be reconfigured dynamically to 6 different combinations with varying vector register lengths, as shown below:

Configuration of Vector Registers

# registers	Register Length by Model (# of 64-bit word elements)			
	500	1100	1200	1400
8	512	512	1024	2048
16	256	256	512	1024
32	128	128	256	512
64	64	64	128	256
128	32	32	64	128
256	16	16	32	64

Other features:

400 and 1300 gate ECL, 350-picosecond delay
 main memory - 256 KB, 55 nsec, MOS static RAM
 380-470 square feet
 36-62 KVA power consumption
 air cooled

Performance: The vector performance varies according to model as follows:

Model	Peak Mflops	Model	Peak Mflops
500	143	500E	286
1100	286	1100E	429
1200	571	1200E	857
1400	1143	1400E	1714
2000	1600+		

The scalar processor cycle time is 14 nsec, compared to the CRAY X-MP's 8.5 nsec, but a sampling of scalar instructions indicates that the VP operations may be slightly faster than the X-MP's. All scalar work can overlap vector operations.

Software:

VP/XA operating system offering IBM MVS/XA system support
 Automatic vectorizing Fortran compiler (Fortran 77/VP)
 Scalar Fortran compiler
 Interactive debugger
 Performance measurement tools
 Interactive vectorizer
 STREAM77 Language Converter
 SIMUL38 IBM 3838 array processor simulator
 Scientific subroutine library (223 routines)

Contact:

Phil Howell
Amdahl Corp.
1250 East Arques Ave.
P.O. Box 3470
Sunnyvale, CA 94088
408-746-6880

Dr. Horst-Peter Rother
Amdahl International Management Services Ltd.
Dogmersfield Park
Hartley Wintney
Hampshire RG27 8TE
England
0252-24555 Telex 858486 G

AMETEK Series 2010

The company is no longer marketing this product.

MIMD Reconfigurable Local-Memory Architecture

AMETEK is a Fortune 500 company with 26 domestic divisions, four European manufacturing sites, and 6300 employees. The Computer Research Division was formed in 1983, and the first generation machine, the AMETEK 14 hypercube, was announced in 1985. The AMETEK Series 2010 is the second generation of AMETEK Concurrent Processing Systems and was announced on January 18, 1988.

Configuration: Each node is based on a 25 MHz, 4 mips 68020 processor with a 68881 420 Kflop arithmetic coprocessor or an optional 68882 630 Kflop coprocessor. Standard local memory for each node is 1 Mbyte, which can be upgraded in 1 Mbyte increments to 8 Mbytes per node.

A VMEbus interface on each node allows up to three VME devices to be attached to every node. One option that can be interfaced through a VMEbus is the vector floating-point accelerator (VFPA). The VFPA is based on Weitek chips rated at 20 Mflops peak and has from 2 to 10 Mbytes of on-board memory. The LINPACK benchmark for each VFPA is 7 Mflops, and the execution rate for scalar operations is 1.2 Mflops.

Message routing is organized through the "GigaLink" network consisting of interlinked modules called Automatic Message Routing Devices (AMRDs), which are full custom VLSI semiconductor devices in CMOS technology. Each AMRD has five bidirectional parallel channels, four for communication with other AMRDs, and one for access to its local node through a special AMRD interface on the node board. Peak traffic volume transmission on each link can exceed 20 Mbytes/sec in each direction for a throughput of more than 80 Mbytes/sec over the network. The routing of messages is automatic and asynchronous of work at the nodes so that computation is not interrupted for message forwarding. Since the topology is defined by the linking of the AMRDs, there is no restriction to a hypercube architecture. The actual configuration can have any number of nodes and is hardware reconfigurable determined by the topology of the GigaLink interconnect. The maximum configuration offered has 1024 nodes, with 8 Gbytes of local memory, 10 Gbytes of VFPA memory, and a peak performance of over 4000 mips and 20 Gflops.

The user interface to the Series 2010 is a SUN-3 workstation. Programs are compiled on the SUN, then downloaded to the Series 2010. The system is a "space share" multi-user system.

Because of the VME interface, a whole range of devices can be directly coupled to each node. These include SMD disk drives, optical disk drives, A/D and D/A converters for signal acquisition and processing, high-speed line printers, external communications controllers, high capacity (up to 700 Mbyte unformatted) SCSI disk storage devices with cartridge or half-inch tape backup, and user-designed VME interfaces. The configuration is totally heterogeneous so different nodes can be fitted with different devices. Graphical output can be transferred at 80 Mbytes/sec through the GigaLink to a series 2010 graphics processor installed on the SUN-3 host workstation.

Software: The operating system on each node is called the Reactive Kernel, and nodes with local disk storage also have a resident UNIX-compatible file server, as well as an interface to SUN's NFS. A Reactive Kernel simulator is available on the SUN for program development and debugging. Fortran 77 with VMS extensions and C are supported with optimizing and vectorizing compilers. Concurrent LISP is also available and Ada is scheduled for mid-1989. A node-level dbx-type debugger is also available.

Applications: AMETEK provides a library of parallel mathematical routines, including matrix and signal-processing subroutines. Additionally, compatibility libraries allow the Series 2010 to execute applications developed for the earlier AMETEK 14 systems or for the Intel iPSC or JPL Mark III hypercubes. Applications software includes FLO 57, an Euler transonic fluid flow program developed by Tony Jameson.

Status: Pricing for a 4-node system starts at \$45,000, with an 8-node system less than \$100,000 and a 64-node system \$495,000. A fully configured 32-node Series 2010 with VFPA's and a peak performance of 640 Mflops is priced at under \$1M. The price quoted for the VFPA (each rated at 20 Mflop) is \$1,000 per Mflop.

The system is boxed in standard 19" RETMA racks. Up to 32 nodes with VFPA's or VME controllers can fit into a single system cabinet measuring 24" w x 48" d x 60" h. A larger cabinet capable of holding 128 nodes is also available. The system is air cooled.

Deliveries of Series 2010 machines to about 6 beta sites is scheduled for 3Q 1988. Production deliveries are expected to begin in October.

Contact:

AMETEK Computer Research
610 North Santa Anita Avenue
Arcadia, CA 91006
818-445-6811

Technical Contact: Dr. Jeff Fier
Sales: John C. Wyckoff III

Active Memory Technology DAP

Bit Parallel Architecture

Architecture: The AMT DAP is an SIMD lockstep machine which operates on multiple data one bit at a time. It can support, via software, variable-length arithmetic. Configuration is a grid of processing elements with nearest neighbour connections and row/column data highways. The row/column data highways allow efficient global fetches and broadcasts giving the system the properties of associative processors.

The major differences over the ILLIAC IV are:

- bit processors
- row/column highway
- much larger memory per processor
- high input/output capability

AMT offers two models of the DAP. The DAP 510 is a 32 x 32 array of processors, and the DAP 610 is a 64 x 64 array of processors. The DAP array is constructed using custom CMOS VLSI chips which contain 64 processor elements. Both models of the DAP currently operate with a 100 nsec cycle time. A real-time graphic display interface is available for the DAP systems. The following table summarizes the characteristics of the two DAP models.

The development environment (cross-compilers and run time debugging aids) are supplied running under UNIX. The DAP is linked as a peripheral via a 1.5 Mbyte/sec parallel interface.

Model	Memory Bandwidth	I/O Data Rate	Processing Elements	Memory Configurations
DAP 510	1.2 Gbytes/sec	50 Mbytes/sec	1024	4, 8, 16 Mbytes
DAP 610	4.8 Gbytes/sec	100 Mbytes/sec	4096	16, 32, 64 Mbytes

Configuration: The DAP 510 is small enough to fit under a desk, while the DAP 610 is housed in a standard EIA rack cabinet. Both DAP models can be hosted by Sun or DEC VAX computers and workstations. The DAP can be connected to a Sun host via the SCSI interface. Connection to DEC VAX systems is via DR11W or DRB32 interfaces. Connection to the Aptecl IOC is supported as well as direct connection to VME bus.

	DAP 510	DAP 610
Array size	32 x 32	64 x 64
Array memory	8 Mbytes	16 Mbytes
Code store	512 Kbytes	512 Kbytes
Instruction rate	10 MHz	10 MHz
host	Sun or VAX	Sun or VAX
Size	17 x 13 x 20 in.	45 x 25 x 38 in.
Price	\$155,000	\$320,000

Software and Languages Available: The principal programming language used is Fortran plus, an augmented Fortan that includes most of the array features proposed for Fortran 8X. APAL, an assembler language, is also available.

Applications: The variable length arithmetic capabilities of the DAP make it particularly well adapted to large scale signal and image processing applications.

AMT provides libraries of algorithms in subroutine form to support image and signal processing application development. A general-purpose algorithm library is also available. Major application areas include scientific and engineering computing, image processing, signal processing, defense applications, and database applications.

The present DAP systems are third-generation machines which started with a 64 x 64 array originally installed at QMC (Queen Mary College, University of London). The QMC machine, which had an effective cycle time of 250 nsec, proved highly adaptable to a wide range of numerical problems based on partial differential equations. The performance on large-scale Monte Carlo simulations in lattice gauge theory and molecular dynamics was found to be exceptional and, in some specialized applications such as the Ising model, the DAP outperformed a CRAY-1 by a factor of 10.

Status: Initial shipments of the DAP 510 began in February 1988. Shipments of the DAP 610 began in November 1988. At year end of 1988, 60 DAP 510 and 5 DAP 610 machines were installed in the United States and Europe.

Contact:

Dr. Geoff Manning
 Active Memory Technology Limited
 65 Suttons Park Avenue
 Reading RG6 1AZ
 0734-661111

Bill Terry
Active Memory Technology Inc.
16802 Aston Street Suite 103
Irvine, CA 92714
714-261-8901

Ardent Titan

Vector Register, Shared-Memory Parallel Architecture - Graphics Supercomputer

The company was founded by Allen Michels (from Convergent Tech) in November 1985, and was originally called Dana Computer Incorporated. Financed by venture capital and Kubota Ltd. of Japan, a \$2.5 billion worldwide industrial equipment manufacturer.

Heavy emphasis on interactive graphics for large computational problems. Graphics boards are integrated with system and many graphics calculations can be done in vector units.

Configuration: One to four processors are connected to a shared memory of between 8 and 128 Mbytes through a 256 Mbyte/sec bus. There are ten slots on the bus, of which six are available for memory boards or CPU boards.

Each CPU has a MIPS chip scalar unit rated at 16 mips with a 16 Kbyte instruction cache and a 16 Kbyte data cache. The vector unit uses a custom designed chip with divide, a pipelined multiplier, and a pipelined adder/subtractor as independent arithmetic function units. Data is streamed from shared memory direct to the vector registers via 1 store and 2 load pipes. The vector registers are 8192 words long and can be configured in any mode between 8192 registers of one word each and 32 registers each of 256 words. Each word is 64-bits long. The clock cycle time is 62.5 nsec, and each vector processor is rated at 16 Mflops, giving a maximum potential of 64 Mflops. Gather/scatter is supported by the hardware.

Memory uses 1 Mbit chips, and each memory board contains 8, 16, or 32 Mbytes. Interleaving is 8-way on odd boards and 16-way on paired boards of the same size. A maximum of four boards can be used. Access rate is 256 Mbyte/sec.

A major feature of the Titan is its integrated graphics support. Up to two graphics boards can be attached to the bus and are powerful processors in their own right. All pixel manipulation is done on the graphics boards minimizing traffic between them and the vector processors, which can be employed on related or independent computation to the graphics processing. Graphics is supported by PHIGS+ and CGI as well as Ardent's own software package called Doré (Dynamic Object Rendering Environment), which handles image representations from wire-frame through flat and smooth-shading to global ray tracing.

Standard interfaces to LANs and I/O devices are supported through an I/O board connected to a single bus slot. The I/O board supports two 4 Mbyte/sec SCSI channels, a keyboard, a mouse, Ethernet, 4 RS-232 ports, and 1 parallel port and can be fitted with a 15 Mbyte/sec VME bus adapter for SMD and other devices, such as knob-boxes, tablets, and stereo viewers.

Software: The operating system is fully compatible with the standard AT&T System V.3 UNIX operating system and Berkeley 4.3 Unix with enhancements for communications, high I/O bandwidth, and large applications. Asynchronous reads and a fast file transfer of 1000 Kbyte/sec using disk striping are also supported.

Vectorizing and parallelizing compilers are available for both Fortran and C and generate

a common intermediate form for subsequent code generation. Standard Fortran 77 is supported along with extensions compatible with the VAX/VMS extensions. CRAY vectorizing directives are also recognized. The parallelism which is automatically detected by the compiler is fine-grained microtasking and several multitasking primitives are also supported. A symbolic debugger, much extended from the basic UNIX debugger, is available. Both 32 and 64-bit floating-point arithmetic is supported, and the arithmetic conforms to the IEEE standard 754.

Performance: The performance on the LINPACK (100 x 100) benchmark on one processor is over 6 Mflops, with 24 Mflops performance on two processors for the 1000 x 1000 LINPACK test. The maximum possible computational rate is 64 Mflops.

Applications: Application agreements are being reached with companies supplying a wide range of applications software. Areas included are Mechanical CAE (NISA, DYN A3D), CFD (PLOT3D, ARC2D, KIVA, NEKTON, VSAERO, FLO87, PHOENICS), Computer-aided molecular design (BIOGRAF, GAUSSIAN 86, MOPAC, AMPAC, AMBER, CHARMM, BIOSYM, and COSMIC), Seismic (LANDMARK), Animation (WAVEFRONT), and Mathematical software (IMSL, NAG, MATHADVANTAGE, MATLAB, LINPACK, EISPACK).

Status: Beta test sites in 1987. General delivery of one- and two-processor machines both in the United States and worldwide by May/June 1988, and four-processor machines in 4Q 1988.

All packaged systems include a 380 Mbyte hard disk, 1/4" cartridge tape, operating system, TCP/IP, C compiler, linker/debugger, Doré, X-windows V.11 and UNIX Navigator, a user interface based on visual agents in the same style as Apple's Macintosh and DRI's GEM, together with a tutorial version of Smalltalk 80, 32 plane graphics subsystem, 19" 1024 x 1024 color monitor, 50 ft video cable, keyboard, mouse, Ethernet connectors, and four RS 232 connectors.

Cost: The cost of the minimum packaged system which has one processor and 8 Mbytes of memory is \$80K. A two-processor, 32-Mbyte system is priced at \$120K. All systems are field upgradable with a separate processor priced at \$21K and a 32 Mbyte memory board at \$16K. The Fortran compiler is \$5K. NFS is \$350. Doré source code licences are available for non-Titan hardware at \$200 (unsupported) for non-commercial customers and \$10,500 (supported) for commercial customers. CRAY- and SUN-specific versions are available if required.

Markets:

CAD/CAM/CAE

Molecular Modelling

Image Processing

Scientific/Engineering Research and Development

Computational Fluid Dynamics

Contact:

Ardent Computer Corporation
880 West Maude Avenue
Sunnyvale, CA 94086
408-732-0400; FAX 408-732-2806

President and CEO: Allen Michels
VP Research and Development: Gordon Bell
Mathematical Software: Cleve Moler

European Office
David G. Howes
Ardent Computer Limited
Brooke House Market Square
Aylesbury
Bucks HP20 1SN
England
(0296) 89911; FAX (0296) 87123; Telex 838811 BROOKE

BBN Butterfly GP1000 and TC2000 Parallel Processor **Parallel Butterfly Network Architecture**

Architecture: The Butterfly GP1000 is a tightly coupled, shared memory multiprocessor housing up to 256 processor boards, each with an MC68020 microprocessor and an IEEE-complying MC68882 floating-point coprocessor. Every processor board includes 4 Mbytes of globally shared memory. Any processor can access any memory location through the Butterfly switch, a fast, modular, multi-stage interconnect. Processors also have direct access to their own 4 Mbyte share of the global memory pool. Providing true parallel access to memory, the Butterfly performs up to 256 simultaneous reads or writes and automatically resolves contention for memory.

Other architectural features include:

- Multiple instruction, multiple data (MIMD) architecture.
- Up to 600 mips of processing power in 2.5 mip increments.
- All processors have access to as much as 1024 Mbytes
(one Gbyte) of main memory.
- Memory bandwidth up to 1024 Mbytes/sec (one Gbyte/sec).
- Memory access time is typically less than 1 microsecond,
4 microseconds worst case (without contention).
- Distributed I/O system supports RS-232, RS-449, Ethernet, Multibus,
and VME bus.
- Field expandable in single processor increments.

Configuration: The GP1000 is a standalone system supporting a full range of peripherals including 500 Mbyte and 850 Mbyte disk drives; 1/4" cartridge and 1/2" reel-to-reel tape drives; a flexible terminal control system; and an Ethernet interface.

Software: Mach 1000, the GP1000 operating system, is based on Berkeley 4.3bsd UNIX, with extensions for parallel processing. The GP1000 supports C, Fortran 77, Common Lisp, and Scheme (a Lisp dialect). Ada is being developed. All languages are extended naturally to support parallel structures. A rich, graphically-oriented debugging environment is provided.

Status: Prices vary from \$95,000 to \$3,500,000 depending on size and peripherals.

The BBN TC2000 incorporates Motorola's 88000 microprocessor. The multiprocessing architecture allows field-expansion from eight to 504 processors, with corresponding increases in memory, memory-access bandwidth, and I/O capabilities.

MAXIMUM SYSTEM PERFORMANCE (504 processors)

Integer	9,576 Dhystone MIPS
Whetstone	6,552 Whetstones
Floating Point	10,080 MFLOPS
Memory	16,096 MBytes
I/O Bandwidth	2,560 MBytes/sec

TC2000 system supports two operating systems concurrently. At the same time as some processors are running the pSOS+m real-time executive for time-critical applications, others can be using the nX operating system (based on UNIX 4.3 BSD) for either analysis or time-shared program development.

A major feature unique to the TC2000 system is its software-controlled clustering capability. Processors can be assigned to groups or clusters, which are then designated for either nX or pSOS+m operation. Different sections of an application can be run concurrently on each one. In addition, data can be shared within and between clusters, so a TC2000 system can integrate various segments of a complex application traditionally dispersed among a number of loosely-coupled computers. Processor allocation is dynamic, meaning that resources can be reallocated during an actual run.

To reduce the time and cost of applications development, the TC2000 system includes the only graphical development tools specifically designed for a multiprocessor environment. Based on the X Window System standard, the Xtra (X Tools for Runtime Analysis) environment makes it easier for programmers to handle the complexities inherent in multiprocessor programming. Included within the Xtra environment is the TotalView source-level, multiprocessing debugger and the Gist graphics-oriented performance analyzer. Optimized Ada, Fortran-77 and C compilers are also available for the TC2000 system.

Pricing for the TC2000 system begins at \$350,000 for a base model with 152 Dhystone MIPS and 160 MFLOPS. BBN Advanced Computers will sell the TC2000 system in technical markets through its direct sales force.

Contact:

Gary Schmidt
Bolt, Beranek and Newman Advanced Computers Inc.
10 Fawcett Street
Cambridge, MA 02238
617-873-2756

Gerry O'Neill
BBN Inc.
Heriot-Watt Research Park
Riccarton
Edinburgh EH14 4AP
Scotland
031-449-5488

CDC CYBER 180 990E/995E

It is possible to do a field upgrade from the one-processor 992-31 to the two-processor 992-32. Each processor has a 16 nsec clock, a 32 Kbyte data cache, and an instruction cache of 64 instruction words (equivalent to a maximum of 256 instructions). There are no vector registers. There are segmented functional units for addition/subtraction, multiplication, scalar multiplication, division (4 units), shift, integer addition/subtraction, compare, Boolean, increment, and character handling. There is shortstopping (that is, effectively chaining) but no scalar/vector overlap.

The memory size is from 16 to 256 Mbytes in 32 banks of 256 Kbit static MOS. There are four ports to memory, 2 for the CPUs and 2 for I/O. The bank busy time is 96 nsec. The virtual address space is 8.8 trillion words/user.

Each system has from 8 to 256 MWords (64-bit words) of MOS semiconductor shared memory using 256K DRAMS. Memory sizes are 8, 16, 32, 64, 128, or 256 million words. There is a 1 MWord fast communication buffer for interprocessor communication and synchronization. The system has virtual memory addressing using a 48-bit address. SECDEC is on each 32-bit half word. The maximum transfer rate is 1 word per clock cycle between each CPU and shared memory.

Up to 18 440 Mbit/sec I/O units are available for accessing disks, tapes, other mainframes, and networks. There is a dual-ported interface to 1.2 Gbytes (formatted) capacity disks with a data transfer rate of 12 Mbytes per second, an average seek time of 16 milliseconds, and an average latency of 8.3 milliseconds. Each IOU can support up to 16 such drives. Each IOU also can interface to as many as 20 1.1 Gbyte capacity, 3 Mbytes per second SMD disks. There are connections to 10 Mbit/sec Ethernet using TCP/IP, CDC's 50 Mbit/sec.

Loosely Coupled Network (LCN) using RHF and a 50 Mbit/sec hyperchannel using TCP/IP.

Configuration: The ETA-10 can have a variety of front-ends including CDC, IBM, and DEC. It can also operate stand-alone, with access from terminals and workstations via Ethernet.

Software:

ETA System V:

- SVID compliant version of AT&T's System V, Release 3.0
- TCP/IP/Telnet/FTP
- BSD 4.3 sockets and "r" commands
- Sun Microsystems' Network File System
- Network Queuing System batch support

EOS:

- VSOS user environment

- TCP/IP/Telenet/FTP
- CDC Loosely Coupled Network
- UNIX utilities

Utilities:

- Interactive symbolic debugger
- Symbolic postmortem dump
- Performance analyzer
- Source and object code maintenance

Many matrix algebra routines are available, including the BLAS in the LIB99 vectorized subroutine library. Also, an object code utility called Afterburner is available that provides user-selected in-lining of system and user subroutines for reduction of call/return overhead.

Languages available and Fortran characteristics: The operating system is NOS/VE and languages supported include Fortran, Cobol, Lisp, Prolog, Pascal, C, Cybil, Basic, and APL. There is a hot spot analyzer available for Fortran codes. The Fortran compiler allows many extensions including some anticipated 8x constructs and generates vector code automatically. Compiler directives are also available.

Performance: The peak performance of the 992 is 125 Mflops.

Contact:

Control Data Corporation
P.O. Box O
HQS09B
Minneapolis, MN 55440
800-828-8001 ext 88

CDC CYBER 205

The company is no longer marketing this machine.

Vector Architecture

Architecture: ECL/LSI logic (168 gates/chip)

Sequential and parallel processing on single bits, 8-bit bytes and
32- or 64-bit floating-point operands

20-nsec cycle time

Scalar Unit

Segmented functional units
64-word instruction stack
256 word high-speed register file

Vector Unit

1, 2, or 4 segmented vector pipelines
memory-to-memory data streaming
maximum vector length of 65,536 words
gather/scatter instructions
up to 800 million 32-bit floating-point operations/second

Memory

MOS semiconductor memory
Memory size: 1, 2, 4, 8 or 16 million 64-bit words
Virtual memory accessing mechanism with multiple, concurrently
usable page sizes
SECDED on each 32-bit half word
48-bit address (address space of 4 trillion words per user)
80 nsec memory bank cycle time
Memory bandwidth: 25.6 or 51.2 Gigabits/second

I/O

Eight I/O ports, 32-bits in width, expandable to 16
200 Mbits/second for each port
Maximum I/O port bandwidth of 3200 Mbits/sec

Miscellaneous

Cooling: freon
Dimensions: floor area (four pipe model) 23 ft x 19 ft
footprint (with I/O system) 105 sq ft

Software:

Virtual operating system
Batch and interactive access
FORTRAN compiler
ANSI 77 with vector extensions
32-bit half-precision data type
Special calls to machine instructions
Automatic vectorization
Scalar optimization utilizing large register file

Utilities

Interactive symbolic debugger

Source code maintenance

Object code maintenance

Performance:

Linked triad performance on long vectors approaches asymptotic speed of machine.

Performance can be severely degraded at short vector lengths (that is, the typical $n_{1/2}$ is around 100), and if vector is not held contiguously. For this reason most tuned software employs long, contiguously held vectors.

Contact:

ETA Systems, Incorporated

1450 Energy Park Drive

St. Paul, MN 55108

612-642-3400

Charles D. Swanson - Account Support

B. Lawrence

Control Data Limited

3 Roundwood Avenue

Stockley Park

Uxbridge

Middlesex UB11 1AG

England

01-848-1919

CONVEX C-120/130 and C-210/220/230/240

Vector Register, Parallel Processor, Bus-Based Architecture

Architecture: The C-2, which was available from January 1988, is a multiple-processor bus-connected, shared-memory computer. Each CPU is similar to (but a new design) the single CPUs of the C-1 computers.

The CPUs consist of a scalar and address unit (based on ECL 7K and 10K density chips) and a vector processor (using CMOS VLSI 20K gates/chip). The vector architecture is register-to-register with three asynchronous pipelined functions (load, store, and edit; add, subtract; multiply, add, divide, and square root). Each CPU has 8 vector registers, each with 128 elements (64-bit elements). VL and VS registers are also present. The scalar unit performs integer arithmetic and floating-point multiplies, adds, divisions, and square roots in hardware. There is a 64 Kbyte cache for the scalar unit with cache bypass for the vector unit. The cycle time is 40 nsec for the C-2 (100 nsec for the C-1). Scalar and vector units (fixed and float) can operate concurrently.

The C-2 has new microcode instructions for vector square root, mask operations, type conversions, intrinsic functions, and random memory access.

Real memory is up to 4 Gbytes (1 Gbyte for the C-1) of DRAM. The early C-1 memories were in 256 Kbit DRAMs, but the later memories and those of the C-2 use 1 Mbit DRAM. Virtual address space is 4 Gbytes (page size 4 Kbytes) with 2 Gbytes available per user. Memory is 64-way interleaved (32 bit) or 32 way (64 bit).

Transfer rates between memory and CPU on the C-1 are rated at 80 Mbytes/sec. There is a single memory pipe between memory and registers.

On the C-2, the access between each CPU and the memory is via a non-contentious, non-blocking 5-bus crossbar using ECL chips, with each bus rated at 200 Mbytes/sec.

The arithmetic is in floating-point IEEE standard format. Byte-addressable with integer*1, integer*2, integer*4, integer*8, complex*8, and complex*16 supported.

There is a 1/2 Mbyte IOP buffer. The IOP is 68000 based with event-driven monitor and I/O transfer rates of 80 Mbytes/sec on custom application boards, or standard Multibus at 8 Mbytes/sec, or VME bus at 16 Mbytes/sec.

Configuration: All machines are stand-alone, multi-user, interactive machines. They can be interfaced to most standard communication channels including Ethernet (TCP/IP), DECnet, and Hyperchannel. Pink book and color books over LAN and NFS are also available. X25 color book will be available shortly. Batch job submission from VAX to C-2 possible with output files and results returned to VAX.

Software: UNIX 4.2 BSD and COVUE shell offer emulation of most common VMS commands.

Languages available: Parallel Fortran, C, vectorized Ada, common Lisp, Prolog

Fortran characteristics: Fortran 77 with VAX extensions and excellent Fortran vector-

izing compiler. C compiler (VC) automatically vectorizes scalar code. HCR/PASCAL and HCR/UX-BASIC are available as third party compilers.

Applications: There is a very extensive range of application software. General-purpose packages include NAG, IMSL, ABAQUS, MSC, NASTRAN, ANSYS, DI-3000, DISPLA, GKSGral, UNIRAS, TELEGRAF, Q-Calc, Sir, and Oracle.

Performance: Peak performance for the C-120 is 20 Mflops in double precision (64-bit arithmetic) and 40 Mflops in single precision (32-bit arithmetic). LINPACK timings are 3.7 Mflops (100 x 100 matrix with unmodified code).

Peak scalar performance of C-210 is 22 Whetstone mips at 32 bit and 14 Whetstone mips at 64 bit (with in-line subroutine expansion). Peak vector performance is 50 Mflops. LINPACK benchmark runs at 10.0 Mflops (again for unmodified code on 100 x 100 case).

The following two tables compare the C-210 performance in Mflops of a single processor with the C-120. The first table compares the performances for the algorithm $a_i = b_i * k$

	64 bit	32 bit
C-120	6.6	13.3
C-210	16.3	25.0

The second table shows a comparison for an indirect vector addressing algorithm of the form $a_{x_i} = a_{x_i} * b_{x_i} * k$

	64 bit	32 bit
C-120	3.6	3.5
C-210	12.5	16.7

The C-210, used in these benchmarks, is the single processor version of the C-2 computer. The multiple processor C-220, C-230, and C-240 versions are available and are all field upgradable from the C-210.

Basic C-120 system: two 19-in. racks and 32 Mbytes memory, 1 I/O processor, service processor, 434 Mbyte Winchester, 6250 bpi tape drive.

Size: 25 x 62 x 40 inches for each cabinet. Base system requires two cabinets, each about 500 lb. Forced air cooling. Power consumption 3200-4500 Watts

A 2 CPU C-220 system consumes 12 KW.

Status: C-120 base system lists at about \$250K but generous academic discounts available. The C-210 will be about twice as much as a C-120 but trade-in possible. Trade-in value calculated by $\text{CPU} * (1 - M/36)$ where CPU is initial cost of C-120 processor and M is its age in months. The F77 compiler costs \$25K, as does the C compiler with GPROFF, PROF, and BPROF run-time profilers. A source level debugger and a range of editors, including VAX EDT emulation, are available.

CONVEX has sold 380 systems (280 C-1, 100 C-2) worldwide since 1985.

Contact:

CONVEX Computer Corporation
701 N. Plano Road
Richardson, Texas 75081
214-952-0200
FAX 214-952-0550 uucp convex!wallach

Technical: Steve Wallach
Sales: Adrian Wise

A.S. Nutt
CONVEX Computer Limited
Hays Wharf
Millmead
Guildford GU2 5BE
England
0483-69000 Telex 858136 Fax 0483-36775

CRAY-1

The company no longer markets this machine.

Vector Register Architecture

This machine is no longer being produced, although when first introduced in 1976 (Los Alamos), it was without doubt the fastest processor in the world and is still used as a benchmark for high-speed computing. Since many CRAY customers are currently upgrading their systems to an X-MP, there are opportunities to buy a second-hand CRAY-1S at knockdown prices.

Architecture:

A uniprocessor.

Vector processor, uses pipelining and chaining to gain speed.

12.5-nsec clock. Fast scalar.

Uses only four chip types with 2 gates per chip.

64-bit word size up to 4 Mwords of storage.

The CRAY 1-S has bipolar (in units of 4K RAM), and the newer (1982)

CRAY 1-M has MOS memory (in units of 16K RAM).

Logic chips - ECL with a gate delay of .7 nsec.

Main memory banked up to 16 ways. The bank busy time is 50 nsec (70 nsec on the 1-M) and the memory access time (latency) is 12 clocks (150 nsec).

Register-to-register machine

8 registers of length 64 (64-bit) words each

Word addressable (64-bits).

No half precision.

Double precision (128 bits) is through software and is extremely slow (factors of about fifty times single precision (64 bits) are common).

There is only one pipe from memory-to-vector registers, resulting in a major bottleneck with loads and stores to memory from registers. Loads can be chained with arithmetic operations; stores cannot.

Software: An extensive range of software exists for this machine. Since the instruction set is compatible with the X-MP range, this software will also run on that range.

Performance: Low vector start-up times and fast scalar performance make this a very general-purpose machine. Max. performance 160 Mflops; 64-bit arithmetic; max. attainable sustained performance 150 Mflops. There are codes for matrix multiplication and the solution of equations which get close to this. Maximum scalar rate is 80 mips. It is easy to attain over 100 Mflops for certain problems, even using Fortran.

Contact:

Cray Research Inc.

1440 Northland Drive

Mendota Heights, MN 55120

612-452-6650

CRAY-2

Vector Register, Parallel, Shared-Memory Architecture

Architecture: This is a 4-processor (quadrant) vector machine with pipelining and overlapping but no chaining. There are more segments in the pipes than in the other CRAYs. Multitasking primitives have same syntax as the X-MP.

The system has a 4.1-nsec clock cycle time.

Memory is 256 Mwords of 256 K DRAM in 128 banks. The bank busy time is 57 clocks, and the scalar memory access time is 59 clocks. Local memory is 16 Kwords, 4 clocks from local memory to vector registers. Vector references from local memory must be with unit stride. There are 8 vector registers each with 64 elements.

Overheads for vector operations are large:

- 63 cycles for vector load
- 22 cycles for vector multiply
- 22 cycles for vector add
- 63 cycles for vector store

Recent enhancements to the CRAY-2 include a 512 Mword memory and models with 128 Mword static RAM. Other improvements include implementing functional units in VLSI (and cutting latency time by half), a larger instruction buffer, reduced branch time, and faster issue rates for certain sequences of instructions.

Configurations: Cray has an ongoing commitment to high-speed peripherals and fast network links. HSX is a 100 Mbytes/sec link for connecting CRAYs together. CRAYs can be linked to Ultra Corporations 1.6Gbit Ultra bus in addition to standard connections with Ethernet (TCP/IP), and VME buses. The DD-40 disks each hold 5 Gbytes and have a transfer rate of 10 Mbytes/sec.

The machine is liquid cooled using inert fluorocarbon.

Software:

- UNIX-based OS (called UNICOS)
- C compiler
- CFT2 (Fortran compiler)
- CFT77

Performance: Peak performance is 488 Mflops per processor. A matrix multiply code has run at 1.7Gflops on 4 processors.

Status: Cost: \$15M - \$20M

Delivered: NMFECC, NASA Ames, University of Minnesota, Harwell Laboratory, Stuttgart, and Ecole Polytechnique (Paris).

Contact:

Cray Research Inc.
1440 Northland Drive
Mendota Heights, MN 55120
612-452-6650

CRAY-3

Vector Register Parallel Architecture

The machine is essentially a GaAs version of the CRAY-2 being developed by a new company under Seymour Cray at Colorado Springs. GaAs components are being developed in cooperation with Gigabit Logic.

Architecture:

- About 300 gates/chip with gate delay of 200 picosecs.
- 2 nsec cycle time
- 4 logical functions/clock period.
- Instruction issue every clock.
- 16 processors.
- 512 Mwords static RAM with 256 or 512 way interleaving.
- Bank busy time of 25 nsec and memory cycle time of 50 nsec.
- 2 ports to memory per CPU.
- Total memory bandwidth will be 16 times CRAY-2.
- VR tailgating can yield similar effect to chaining.
- Peak rate of about 16 Gflops.
- 1000 Mbyte/sec channels.
- Liquid coolant immersion.

CRAY-2 imbalance removed by increasing scalar speed to four times that of a CRAY-2 on each processor giving 12 times scalar speed. Aim is 100 times a CRAY-1.

Configuration: Boards reduced from the 4x8x1" of the CRAY-2 to 1x1x0.1". There are four modules to each processor, each module containing sixteen 1" cubes. Overall dimensions of 28" diameter and 4" to 6" high, with power dissipation of 180 KW as in CRAY-2. Power supplies take 10 cu ft and liquid coolant 100 cu ft.

Software: Operating system will be UNICOS, CRI's UNIX system.

Status: The price is likely to be around \$25M to \$30M for a full configuration, naturally dependent on market forces at time of launch. Full production by 1991.

Contact:

Cray Computer Company
P.O. Box 17500
Colorado Springs, CO 80935
719-579-6464

CRAY X-MP

The company is no longer marketing this product.

Vector Register, Parallel, Shared Memory Architecture

Architecture: This multiprocessor pipelined vector machine has the same architecture as the CRAY-1. The major difference is that there are three paths from memory to the vector registers, and the clock cycle time is 8.5 nsec on all machines shipped after August 1986 (machines built before August have a cycle time of 9.5 nsec.)

The current machines come with 1, 2, or 4 processors. Gather/scatter hardware is available on the 2- or 4-processor version of the machine. The gather/scatter can be chained to a load/store operation. Users can control all processors through calls in Fortran. The processors share memory.

Other features:

Memory up to 16 M (64-bit) words

X-MP-2 MOS. (Bank busy time is 68 nsec and memory access time is 17 clocks.)

X-MP-4 ECL. (Bank busy time is 34 nsec and memory access time is 14 clocks.)

ECL logic with .35-.5 nsec gate delay and 16 gates/chip.

Main memory - ECL 4K RAMs with 25-nsec access time. (Interleaving to 64 banks is possible.)

High-speed connection at 1024 Mbytes/sec per channel (max. 2) to a CRAY SSD. The SSD comes in various sizes up to 512 Mwords of secondary MOS memory. Data transfer to high speed (1200 Mbyte) DD-49 disk takes 10 Mbytes/sec. Recent peripheral enhancements as reported under the CRAY-2.

Configuration: There are many possible front ends including IBM, CDC, VAX, and Apollo.

Performance: Peak of 235 Mflops per processor.

Status: Cost: \$25M to \$30M for a fully configured system at present-day prices.

Delivery: Announced in August 1982, first system delivered in June 1983.

Contact:

Cray Research Inc.
1440 Northland Drive
Mendota Heights, MN 55120
612-452-6650

CRAY Y-MP

Vector Register, Parallel, Shared-Memory Architecture

Architecture: This is a multiprocessor pipelined vector machine. It has a similar architecture to the CRAY X-MP. A major difference is the availability of 32- as well as 24-bit addressing. The cycle time is 6 nsec, and it is an 8 processor machine. As in the X-MP there are three paths from memory to the vector registers.

There are only three module types, for the CPU (8 modules), the memory (32 modules, with 1 Mword/module), and the clock (1 module), making 41 modules in all compared with 144 in the X-MP. Each module is on an 11" by 21.2" board. 2.5μ ECL in 2500 gate arrays with gate delay of 350 picoseconds. There are 312 arrays per processor on four PCBs with a power dissipation of 9 Watts per array.

The processors share a common memory of 32 Mword in bipolar SRAM with a 15 nsec access time and a bank busy time of 102 nsec. Memory is interleaved in 256 banks. Total bandwidth is 340 Gbytes/sec (32 words/CP per processor). The Y-MP comes with a 128 Mword SSD as a standard feature. 2 IOSs can be fitted, each with a 4 Mword buffer memory.

Configuration: There are many possible front ends including IBM, CDC, VAX, and Apollo. There are four VHISP channels each rated at 1250 Mbytes/sec, eight 100 Mbytes/sec HISP channels, and eight LOSP 6 Mbytes/sec channels. A full range of disks, tapes, terminals, workstations, and networks (including TCP/IP) are supported.

Cooled using inert fluorocarbon but not with liquid immersion technology (as on CRAY-2).

The dimensions of the machine are 77" x 30" x 75" with a total footprint of 98 sq ft and weighing 5,000 lb.

Software: COS and UNICOS are both supported. In addition to CFT77 a new Fortran compiler will be available. Performance tools include dynamic and static analysis, tuning, and debugging aids. Can run in X or Y mode, and software can also run (in X mode) on the X-MP.

Performance: Peak performance of 4 Gflops. Overall performance of about 30 times a CRAY-1. Each processor should outperform a single X-MP processor by a factor of 1.4 in vector mode (1.2 in scalar).

Status: 1 CPU running in 1987; first deliveries in 1988; nine deliveries in 1989; full production and possible enhancements in 1990.

Cost: Around \$25M for a fully configured machine.

Contact:

Cray Research Inc.
1440 Northland Drive
Mendota Heights, MN 55120
612-452-6650

**Les Davis
1100 Lowater Rd.
Cray Research Inc.
Chippewa Falls, Wisconsin 54701
715-726-1211**

CULLER 8

The company is no longer marketing this product.

Bus-connected parallel architecture

Architecture: Proprietary 64-bit dual interleaved floating-point processors. Each processor has a local high-speed static memory of 32 Mbytes. A global dynamic memory is available in 24 Mbyte increments.

Configuration: The CULLER 8 is front-ended by a SUN 3 with three board set performing at 37.5 Mflops. An HSC link at 100 Mbytes/sec interconnects up to 16 computing nodes. Disk storage via IBIS 1 Gbyte drives with a 12 Mbyte/sec data transfer rate.

Software: The F77 and C compilers share a common back-end code generator. The compilers are designed to exploit the architecture and make efficient use of all internal parallel resources. DEC VMS extensions are supported.

CULLER CSD 4.2 UNIX is an extension of Berkeley 4.2 UNIX to integrate user processes. All normal SUN networking and development utilities are supported.

Floating-point arithmetic conforms to IEEE standard.

Performance: Up to sixteen processors (UPs, or Unit Processors) may be linked with the CULLER 8 PSC. Each UP has a peak rating of 75 Mflops in both single and double precision.

The harmonic mean of all the Livermore Loop benchmarks is expected to be well in excess of 10 Mflops for each UP. This is scaled from the established performance of the CULLER 7 architecture from which the CULLER 8 is derived. The CULLER 8 has five times the scalar and ten times the vector performance of that machine.

Additional hardware modules will accelerate specific algorithms such as FFTs at 120 Mflops and convolution at 100 Mflops.

Status: Deliveries were to commence in 1989, with full production by 2Q, 1989.

The price of a complete 37.5 Mflop CULLER 8 PSC was to be about \$90,000, with 75 Mflop UPs costing about \$80,000.

Contact:

Glen Culler Associates
Martin Beenham
Compass Systems
Bridge House
Faraday Road
Newbury
Berks RG13 2DH
England
0635-521600 Telex 846301 Compas G FAX (0635) 521268

CDC CYBERPLUS

Ring Bus Architecture

Architecture: This is a multiple parallel processor system. It grew from the Flexible Processor Project and the subsequent Advanced Flexible Processor Project (AFP), used in military applications since 1976. The machine is based on ring technology with an 800 Mbits/second transfer rate, with a read and a write possible between processors at this sustained rate.

There are two CYBERPLUS processor modes: 16-bit integer and 32- and 64-bit floating point. The integer processor has 15 independent functional units capable of 8-, 16- and 32-bit working; each processor has a 20-nsec cycle time. The floating-point processor is an extension of the integer one through the addition of three floating-point functional units capable of 32- and 64-bit precision, with rated maximum performance of 65 Mflops (103 in 32-bit mode). Each processor contains 2048 Kbytes of memory expandable to 4096 Kbytes. A crossbar architecture allows the output of one functional unit to go to any or all other functional units in one machine cycle and permits all functional units to fire every cycle. The independent functional units are as follows:

- 1 program unit
- 9 I/O units including 4 read/write 16-bit memory units
- 2 read/write 64-bit memory units, 2 ring port I/O units,
- 5 integer/Boolean units (2 add/subtract, 1 multiply, and 2 shift Boolean)

Floating point: 1 add/subtract, 1 multiply, 1 divide/square root connected by an additional crossbar. Floating-point units can run simultaneously with fixed-point ones. Each instruction can initiate multiple functional units.

Configuration: Up to 16 rings can be connected to a CYBER 800 computer (through a channel ring port) with up to 16 CYBERPLUS processors per ring. Within this ring all processors can operate autonomously and may execute each clock cycle. Processor Memory Interface allows direct reading and writing of the memory of any processor by another processor on the ring every machine cycle. Central Memory Interface (CMI) for transfer of data to host. The central memory ring is 64 bits wide with an 80 nsec cycle time, and this provides a direct transfer of 64 bits between the CYBER and a CYBERPLUS processor. Data transfers are controlled by the system ring and will be direct memory-to-memory transfers with the HPM memory on the CYBERPLUS processors. Two rings connect the processors: the system ring and the application ring. The ring packet has 13 bits of control information and 16 bits of data. A function code in the ring packet can determine whether access to other memories (one or several) is direct or indirect, the latter requiring the acceptance by the target processor.

There are three distinct memory systems:

1. 4K 16-bit data memory: 4 independent bipolar data memories with a one-cycle read/write.

2. 256K 64-bit high-performance data memory: 4 banks with 4-cycle memory access, expandable to 512K 64-bit words with 8 banks.
3. Program Instruction Memory with 4096 200-bit words. Each machine cycle, the instruction memory fetches and initiates the execution of one or all of the parallel functional units. When the floating-point option is in use, the size of these memory words increases to 240 bits. The program instruction memory is expandable to 16K words.

Software: The host CDC 170 Series 800 (under NOS 2) loads code into the processors, transmits data from host to processors, and starts and stops processor's task. Software includes a cross assembler (MICA), a CYBERPLUS instructor load simulator (ECHOS), and an ANSI 77 Fortran cross-compiler.

64-bit floating point is 14 decimal accurate with a range of 10^{-293} to 10^{+322} . 32-bit is 7 decimal accurate with range 10^{-39} to 10^{+37} .

Performance: Claimed performance of 64 CYBERPLUS systems linked to a single Control Data 170 Series 800 is 16 billion calculations per second on signal data applications. Edge detection algorithm for image processing is about 100 times faster than on a CDC 7600.

Status: Announced on October 4, 1983; Fortran compiler available for research activities fourth quarter 1984 and released April 1985.

Cost: Base processor is priced at \$470,000, which includes a 16-bit integer unit and 2048 Mbytes of memory. With all available options the price is \$1.6 million.

Contact:

Martin Ferrante
 Control Data Corporation
 CYBERPLUS Marketing
 P.O. Box O
 HQS09B
 Minneapolis, MN 55440

B. Lawrence
 Control Data Limited
 3 Roundwood Avenue
 Stockley Park
 Uxbridge
 Middlesex UB11 1AG
 England
 01-848-1919

Cydrome CYDRA 5 (formerly AXIOM Systems)

The company is no longer marketing this product.

VLIW Directed Dataflow Architecture

The main memory has a sustained transfer rate of 400 Mbytes/sec and consists of 8 to 256 Mbytes with up to 64-way interleaving. There is also a support memory of 8 to 64 Mbytes that is optimized for rapid data access. The virtual address space is 2 Gbytes.

Peripheral devices supported include disks of capacity 830 Mbytes/drive with a 2.5 Mbyte/sec transfer rate. There are RS-232C and Ethernet connections and a 6250 bpi tape drive with 75 ips in start/stop mode.

Configuration: The central system unit has dimensions 66" x 61" x 34" and weighs 1350 lbs. The central system unit has dimensions 66" x 61" x 34" and weighs 1350 lbs.

Software: The Cydrix 5.3 operating system is a compatible extension of AT&T UNIX System V.3. Extensions include support for transparent multiprocessing, dynamic load balancing, extent-based file systems, buffered and unbuffered I/O, asynchronous I/O, disk striping, batch queue facilities, TCP/IP, remote graphics library, performance profiling tools, and a socket library compatible with Berkeley 4.2 UNIX.

Cydrix Fortran 77 incorporates the ANSI Fortran 77 standard with DEC, IBM, UNIX, and Cydrix Fortran extensions. It generates code for the Interactive Processors and schedules code execution on the Numeric Processor. There is a source level debugger.

Contact:

CYDROME Inc.
1589 Centre Pointe Drive
Milipitas, CA 95035
408-945-6300 FAX 408-262-8938

ELXSI 6400

Parallel Processor/Bus Architecture

Architecture: The System 6400 features a high-speed 64-bit bus architecture. Multiple CPUs, IOPs, and memory modules, based on high-density LSI components using ECL technology, plug directly into the bus and communicate using microcoded messages. Modules, operating in parallel, perform processing, I/O, and memory operations simultaneously.

Configuration: The system can be configured with 1-12 CPUs, 1-4 IOPs, and up to 2 Gbytes of main memory. The CPUs and IOPs have their own local cache of up to 1 Mbyte of high-speed RAM. Each IOP can support up to 32 I/O controllers. The SECDED memory is interleaved 2-way internal and up to 16-way external.

The main memory is accessed through the fast bus. The bus is a 64-bit wide channel providing a gross bandwidth of 320 Mbytes per second, giving a transfer rate of 160-213 Mbytes/second.

Other features:

M6410/M6420

Each CPU has 3 boards, rated at 7 MWhets on M6410 CPU and at 12 MWhets on M6420 CPU.

Up to 12 CPUs.

64-bit wide data paths.

50-nsec cycle time.

6410 16-Kbyte, 2-way set associative cache

6420 64-Kbyte, 2-way set associative cache (100-nsec access time).

16 sets of 64-bit general-purpose registers.

IEEE floating-point arithmetic.

M6460 (Pegasus)

Each CPU has 2 boards, rated at 57 MWhets and 10 Mflops on LINPACK.

Up to 10 CPUs

64-bit wide data paths.

31.25-nsec cycle time.

1 Mbyte of cache allocatable dynamically by the users.

16 sets of 64-bit general-purpose registers.

Five-stage pipelining with optimization.

IEEE floating-point arithmetic.

Fully compatible with 6410 and 6420.

Dual floating-point units.

External interrupt lines for real-time processing.

The system is a stand-alone system with a high-performance I/O system capable of a peak of 64 Mbytes/sec. Various controllers can be attached to the I/O processor including disks, tapes, asynchronous terminals, DRII, VME, Ethernet, X.25, and printers. Networking

is available over Ethernet using TCP/IP and/or Community (DECnet), over X.25 using Coloured Books, or over DRII to Hyperchannel.

Software: Multiple operating systems can operate concurrently on the System 6400. Virtual memory management, load balancing, and process migration are incorporated as a base for all operating systems including EMBOSS (ELXSI proprietary Message Based Operating System), ports of AT&T System V.3 and BSD 4.3, and EMS (ELXSI's VMS-like environment).

ENIX System V runs a native port of UNIX System V.3 in a multiple processor environment. It migrates UNIX processes across multiple CPUs, performing load balancing and resource allocation automatically. ENIX System V runs multiple copies concurrently on single or multiple CPUs. Shared libraries, C and Bourne shells, and TCP/IP over Ethernet are supported.

ENIX BSD runs a native port of 4.3 BSD, allows 2 Gbytes per process, and is efficient on memory-intensive applications. Again multiple copies run concurrently on single or multiple CPUs, and there is automatic load balancing and resource allocation.

The EMS system includes ECL, which interprets and executes VMS interactive commands and command files; ERT, which provides VMS applications with system and utility support during compilation and execution; and EDT, which is an interactive text editor.

From any of the operating systems, programmers can access parallel intrinsics at both the micro and macro level. Functions included are parallel execution of subroutine-level tasks, parallel execution of loops by dividing loops into microtasks and executing microtasks in parallel, automatic load balancing of parallel processes, non-cacheable or cacheable data sharing, and simultaneous parallel processing and general purpose multiprocessing.

Languages: Fortran 77, Pascal, COBOL 74, C, MAINSAIL, Franz LISP, Common Lisp, Simscript, and Ada. DEC Fortran extensions are supported. Auto-vectorizing, and a full suite of debugging facilities, including a symbolic debugger and monitoring utilities, are available.

Performance: The peak rating of the Pegasus machine is 250 mips. A single Pegasus processor performs at 10 Mflops on the LINPACK benchmark, and there are up to 10 processors in a single configuration. A 12-processor M6420 is rated at 120 mips.

Status: The first machine based on a 6410 CPU was delivered in 1983. The 6420 CPU was first delivered in 1986, and the 6460 was announced February 1988 and is scheduled for delivery in 4Q 1988. ELXSI now has over 200 CPUs installed in over 100 systems at 80 customer sites.

Contact:

ELXSI
2334 Lundy Place
San Jose, CA 95131-1873
408-942-0900; Telex 172320; FAX 408-945-5875

Joseph Rizzi - Chairman
James Dutton - President
Tony Yates - Marketing
Pat Trytten - VP Research and Development

ELXSI Bridge House
Walton-on-Thames
KT12-1A1
England
011-44-932-253081; FAX 011-44-932-247199

Chris Morrow - General Manager

John Ware - Sales Consultant

Encore Multimax 310/320

Parallel/Bus Multiprocessor Architecture

Architecture:

National Semiconductor 32332 chip set running at 15 MHz.

64-Kbyte write-through cache per CPU.

Processors connected via a fast, 64-bit wide bus (the nanobus)
with data throughput rate of 100 Mbytes/sec.

Address space of 4 Gbytes

Main memory from 16 to 128 Mbytes of RAM in dual, independent 32-bit banks,
in increments of 16 Mbytes.

Standing floating point - NS32081 hardware FPUs (one per CPU).

Single (32-bit) and double (64-bit) IEEE format.

Optional floating point: Encore proprietary control chip plus Weitek 1164/1165.

Configuration:

+ Multimax 310—11 slots: 4 memory card slots, 6 processor
and I/O card slots, 1 system control card slot

+ Multimax 320—20 slots: 8 memory card slots, 11 processor and I/O slots,
and 1 system control card slot

Terminal and unit record I/O connected via Annex 16 or 32 line terminal
concentrators attached to Ethernet, providing pre-processing.

From 2 to 20 processors.

Ethernet communications using TCP/IP.

Software: UMAX 4.2 (bsd compatible), UMAX V (AT & T System V), MACH; Parallel
Tools: Microthreads

Languages: C, Fortran, Encore parallel Fortran, Ada, Concurrent Ada, Lisp, Cobol;
Parasight debugger. **Other:** Universe, Oracle, Ingres, Informix, and more.

Performance: 4 mips to 40 mips; 5 to 50 Mflops.

Status: Currently over 200 systems worldwide. Marketed in Sweden by Erbe Data, in
Japan by Rikei, in Malaysia by BIS, and in Australia by DISC.

Contact:

Encore Computer Corporation
257 Cedar Hill St
Marlboro, MA 01752
508-460-0500

Lynne Connors- Marketing Support Manager

ETA-10

Vector Parallel Architecture

The company is no longer marketing this product.

The ETA-10 is a successor to the CYBER 205 and is manufactured by ETA Systems, a subsidiary of Control Data. ETA Systems was founded in 1983. The ETA-10 major announcement was on April 27, 1987. The company folded in 1989.

Architecture: Multiprocessor system with up to 8 processors. Very high density circuitry (20,000 gates/chip) in 1.0 μ CMOS. The top end of the range (Models G and E) use liquid nitrogen cooling; the P and Q models are air cooled. Each processor occupies a single 44-layer printed circuit board containing 240 chips, measuring 16" x 22" x .25", and containing nearly 3 million gates. A Built-in Evaluation and Self-Test (BEST) feature is present in each 20 K gate array. Each processor has 4-16 Mwords (64-bit words) of static RAM local memory with a memory bandwidth of 8 words per clock cycle.

The ETA-10 requires only 700 Watts per CPU (i.e., about 200 Watts per CYBER 205 equivalent). CPU and memory require 1.6 KWatts.

Languages:

Fortran

- ANSI 77 with vector extensions
- VAST-2 vectorizer can run as precompiler
- 32-bit half-precision data type
- Special calls to machine instructions
- Anticipated Fortran 8x array notation
- Automatic vectorization
- Scalar optimization
- Multiprocessing library and directives
- Symbolic debugger

C

- AT&T compatible
- Scalar optimization
- Symbolic debugger

Performance: 44 configurations in the product line. The specifications for each model are given in the table below:

Model	Cycle time	Processors	Peak	Shared Memory	LINPACK
G	7 nsec	2-8	10.3 Gflops	2 Gbytes	93 Mflops
E	10.5 nsec	1-8	6.9 Gflops	2 Gbytes	62 Mflops
Q	19 nsec	1-2	947 Mflops	512 Mbytes	34 Mflops
P	24 nsec	1-2	750 Mflops	512 Mbyte	27 Mflops

Some CFD customer benchmarks have rated the 10.5 nsec machine at 11 times the performance of a single-processor CRAY X-MP when 32-bit working is used on the ETA-10.

Status: Over 40 machines installed by year end 1988, with approximately 30 at customer sites. Large liquid-cooled models are installed at Florida State University, the German Weather Service, the John von Neumann Center at Princeton, U.K. Met., Tokyo Institute of Technology, and the Minnesota Supercomputer Center. The Tokyo Institute machine is the first 8-processor supercomputer installed anywhere in the world.

Cost: The ETA 10 product line ranges in price from under \$1M to \$22M. A four-processor Model G (7 nsec) with 64 Mwords of shared memory and 15 Gbytes of disk is priced at around \$14M. An ETA 10-P (24-nanosec) with disks and software costs \$995,000. A special program is available for universities for 21-nanosec ETA-10P systems at the \$995,000 price. Future developments expected are increased memory capacities, faster CPUs, additional standard networks, and higher speed peripherals.

Contact:

ETA Systems, Incorporated
1450 Energy Park Drive
St. Paul, MN 55108
612-642-3400

FLEX/32 MultiComputer Parallel Bus Architecture

Architecture: This machine is a true 32-bit multicomputer with variable architecture structure and is an MIMD machine. It uses National Semiconductor 32032 chips at 10 MHz, with an independent self-testing system using a Z80 micro. The local memory cycle time is 145 nsec. The claimed limit on the number of CPUs is 20480.

Each processor is on one PCB with full 32-bit data bus and full 32-bit address capability, with speed capacity of approximately of 1 mip using the 32032. Each card has a hardware floating-point processor and hardware memory management and memory protection with a local bus interface and a 32-bit VMEbus I/O interface. Also, each processor board has 1 Mbyte or 4 Mbytes of ECC RAM in addition to cache memory and 128 K of ROM. An optional 1 Mbyte of RAM (later planned to have up to 8 Mbyte) with integral error detection and correction code logic is available. Also, an optional floating point accelerator (1 Mflop) is available on each processor. The company envisages attaching array processors that are VME compatible such as a SKY Warrior.

Other features:

Standard VME bus open architecture supporting Eurocard standard.

Communication rates on local buses of 160 Mbit/sec each.

Communication rates on common bus of 380 Mbit/sec each.

Time to get on local bus - 1 msec.

Time to do an an arbitrated read/write through high-speed (45 nsec)

common memory - 170-185 nsec

Direct message passing into another processor's memory is via global memory.

Configuration: The machine can have flexible configuration of local (145 nsec) and common memory (45 nsec). Mass memory cards (local memory) contain from 1 to 8 Mbytes RAM connected by local and/or 32-bit VMEbus I/O interface and can be used in any combination or permutation with CPU cards (these memory cards also have a microprocessor for SelfTest diagnostics and fault isolation). The system can be dynamically configured and reconfigured using the SelfTest mechanism.

There are two computers in two 19-in. standard cabinets:

- one cabinet (the peripheral control cabinet PCC) for the SelfTest System and VME Eurocard card cage (with room for further 19-in. card cages for peripherals)
- the other cabinet (the MultiComputer Cabinet MCC) with a 30-slot card cage partitioned into three 10-slot sections. The backplane contains 2 common buses, 10 local buses, and 20 VMEbus interfaces. The MCC also houses a local bus to common bus interface (common control card) with fair arbitration mechanism, up to 9 common access cards with 128 Kbytes to 512 Kbytes of common memory (45 nsec) each, and a

universal card with 128 Kbytes ROM, 1Mbyte or 4 Mbytes of ECC RAM, 1 mip processor, and VME interface with a separate microprocessor for the SelfTest System.

Cabinet size is 24"x76"x36".

Each cabinet can include up to 20 32-bit processors or 160 Mbytes of memory.

Software: A full UNIX System V can run on each processor, with extensions for concurrent processing. FLEX has a 4.2 license. The software license is for 32 users, with optional software license for unlimited users.

FLEX's own multicomputing multitasking operating system (MMOS) is for real-time operating system support providing all the tools for interprocessor communication and signalling, synchronization, event management, etc.

Ethernet-supported TCP/IP

Languages:

Fortran 77 with ISA S61.1 extensions

Ratfor

C

Concurrent C and Fortran by using a preprocessor

Assembly

Ada under development

Status:

Cost: Price starts at approximately \$100,000

\$36,000 list price/CPU + 1Mbyte RAM with 128 Kbytes ROM, FPP and MMU.

Contact:

Larry Samartin

FLEX Inc.

Dallas, TX 75229

214-929-6000

Distributors in Germany, France, and Norway.

**Floating Point Systems MP32 SERIES
MIMD**

Architecture: Both local and global-shared memory. Range of memory sizes available is 1 Mword to 16 Mwords (32-bit)

Configuration: Basic chip used M68000 (Control Processor), AMD & Weitek Chips (arithmetic processor).

Front ends: DG MV Series, Perkin-Elmer, Microvax II, VAX

Peripherals: I/O port, IOP-32. Bus connectivity.

Software: Own. IEEE standard 32-bit.

Languages: MAX 68 control language, XPAL assembler, MPFORTRAN, XPFORTRAN.

FORTRAN characteristics: F77 (MPFORTRAN and XPFORTRAN, which are F77 less I/O and character data type support). Extensions: Calls to coprocessor programs and MAXL. Debugger: MPFORTRAN debugger. Vectorizing/parallelizing capabilities: Horizontal microcode synthesis that allows up to 10 operations to execute simultaneously.

Performance: Peak: 18 to 72 Mflops. Benchmarks: 2D CFFT 1024 x 1024 pts - 1.89 sec. (3 coprocessors).

Applications: Runs on prototype or on front-end simulator. Software available includes several math libraries: Basic math, Signal, Image, Seismic, and Parallel Processing Constructs.

Status: Available since 8/85

Cost range: \$57,500 to \$125,000

Proposed market: Signal processing, Image processing, and Computational physics.

Contact:

Jim Christiansen
FPS Computing Inc.
3601 SW Murray Blvd.
Beaverton, OR 97005
503-641-3151

Duncan Hamilton
FPS Computing U.K. Limited
Apex House
London Road
Bracknell
Berks RG12 2TE
England
0344-56921 Telex (851) 849218 FPS UK G

Floating Point Systems FPS-5000 SERIES MIMD

Architecture:

Basic chip used: AMD Chips, Weitek Chips on coprocessor.
Local and global-shared memory.
Bus connectivity.
Range of memory sizes available: 256K to 1024K (38-bit words) data,
4K to 64K (64-bit words) program memory
Floating point unit: 32-bit IEEE (coprocessor)

Configuration: Front ends: VAX; PDP-11; Microvax; Perkin-Elmer 3200; Gould 32; IBM 4300, 3080, 3090; Norsk Data; Data General; Prime 750, 9950; Harris 800, HP 1000E.

Peripherals: 300Mbytes and 80Mbytes Disks, Programmable I/O processors

Software: Own system.

Languages: CP FORTRAN, XPFORTRAN, MAXL control language (FORTRAN-like); APAL and XPAL assemblers

Fortran characteristics: F77 (CPFORTRAN and XPFORTRAN, which are F77 less I/O and character data type support). Extensions: Calls to coprocessor programs and MAXL. Debugging facilities: Symbolic debugger. Vectorizing/parallelizing capabilities: Horizontal microcode synthesis that allows up to 10 operators to execute simultaneously.

Applications: Runs on prototype or simulator on front end

Software: Math Libraries—basic and advanced math, signal and image processing, simulation and geophysical, graphics, and parallel processing constructs.

Performance: Peak: 8 to 62 Mflops. Benchmarks on codes and kernels: 2D convolution 31x31 operations - 33 Mflops (FPS-5430)

Status: Date of delivery of first machine, beta sites, etc.: October 1983

Cost range: \$45,000 to \$99,000 for 256Kword system + standard software

Proposed market: 350+ units per year in signal processing, image processing, geophysical analysis, computational physics, real-time simulation, and computer graphics

Contact:

Jim Christiansen
FPS Computing Inc.
3601 SW Murray Blvd.
Beaverton, OR 97005
503-641-3151

Duncan Hamilton
FPS Computing U.K. Limited
Apex House
London Road
Bracknell
Berks RG12 2TE
England
0344-56921 Telex (851) 849218 FPS UK G

Floating Point Systems M64/40, M64/140, M64/50, M64/60
Scalar Pipelined Architecture

Architecture:

Basic chip used: Proprietary (CPU), Weitek chips (MAX)
Range of memory sizes available: .5 to 16 Mwords (64-bit words)
Floating point unit: IEEE Standard compatibility

Configuration: Front-end connection to IBM 4300, 308x, 303x, 309x under MVS, MVS/XA, VM/CMS; all DEC VAX under VMS; Sperry 1100 Series; Apollo Domain.

Peripherals: Disk subsystems (1-2 controllers, 1-8 drives), P64/40 (850 Mbytes to 6.8 Gbyte), P64/20 (255 Mbytes to 2.0 Gbyte removable). Mass storage subsystem: P64/110 (128Mbytes to 15.7 Gbyte at up to 22 Mbytes/sec). I/O subsystem: P64/210 (High-speed interface to disks, tapes, graphics terminals, allowing shared files with VAX front-end).

Software: System Job Executive, Math Library routines (500+), Fast Matrix Solution Library (FMSLIB), NAG, IMSL, BCSLIB, LINPACK, and over 150 other third-party software packages available.

Languages: FORTRAN, ASSEMBLY, C

Fortran characteristics: F77 ANSI 77 optimizing compiler, 5 levels of optimization. Extensions: DOE Extensions for asynchronous I/O, and VMS Fortran extensions. Debugging facilities: Symbolic debugger Vectorizing/parallelizing capabilities: Takes advantage of architecture through horizontal microcoding allowing 10 different operations to occur in 8 separate functional units per machine/cycle.

Performance: Peak: 11 Mflops (M64/40), 33-187 Mflops (M64/140), 19 Mflops (M64/50), 38 Mflops (M64/60). 1000 x 1000 matrix multiply - 189 seconds (M64/40); 54 seconds (M64/60)

Status: Available since July, 1985.

Cost: \$298,000 to \$950,000.

Proposed market: Computational Chemistry/Physics, Electronic Circuit Design, Oil Reservoir Simulation, Structural Analysis

Contact:

Pat Moore
3601 SW Murray Blvd.
Beaverton, OR 97005
503-641-3151

David A. Tanqueray
Apex House
London Road
Bracknell
Berks RG12 2TE
England
0344-56921 Telex (851) 849218 FPS UK G

Floating Point Systems M64/20 and M64/30 Pipeline scalar processor with attached processor

Architecture: The M64 Series includes the three entries just described. Here we discuss features of the other machines in this series.

Basic chip used: Proprietary using CMOS (CPU)

Range of memory sizes available: 1 Mwords to 4 Mwords (64-bit words)

Floating-point unit: IEEE Standard compatibility

Configuration: M64/20-Technological continuation (in CMOS) of 164/264 range. 6 Mflops 64-bit floating-point arithmetic. M64/30-Similar to above but with 12 Mflop floating-point arithmetic. The models M64/20 and M64/30 can also be supplied with TCP/IP, NFS, (IEEE 802.3) Ethernet interfaces and are then termed the M64/220 and M64/230. They can also be supplied with a VAX Station II/GPX, disk, cartridge tape, mouse and keyboard, colour monitor, and graphics processor, they are then termed the M64/320 and M64/330, respectively. Front-end connection to all DEC VAX under VMS and SUN workstations.

Peripherals: P64/30 disk subsystem (1-6 controllers, 2-24 drives) 500Mbytes to 6.0GB, plus 2 x 250 Mbytes internal disks.

Software: System Job Executive

Languages: Fortran, Assembly

Fortran characteristics:

F77 ANSI 77 optimizing compiler, 5 levels of optimization

Extensions: DOE Extensions for asynchronous I/O and VMS Fortran extensions.

Debugging facilities: Symbolic debugger

Vectorizing/parallelizing capabilities: Takes advantage of architecture through horizontal microcoding allowing 10 different operations to occur in 8 separate functional units per machine/cycle.

Applications: As per M64/40, M64/140, M64/50, M64/60 (q.v.)

Software: As per M64/40, M64/140, M64/50, M64/60 (q.v.)

Performance: Peak: 6 Mflops (M64/20), 12 Mflops (M64/30).

Status: Cost range: M64/20-30, \$65,000 to \$160,000; M64/220-230 \$70,000 to \$165,000; M64/320-330 \$95,000 to \$200,000

Contact:

Pat Moore
3601 SW Murray Blvd.
Beaverton, OR 97005
503-641-3151

David A. Tanqueray
FPS Computing U.K. Limited
Apex House
London Road
Bracknell
Berks RG12 2TE
England
0344-56921 Telex (851) 849218 FPS UK G

Floating Point Systems FPS T Series

The company is no longer marketing this product.

Hypercube architecture - Vector processors

Architecture: The Inmos T414 Transputer is a 32-bit CMOS processor, rated at 7.5 mips, with 2 Kbytes of on-chip RAM with one-cycle access that serves as a large register set. There are 4 links which can sustain .7 Mbytes/sec in each direction and can be multiplexed four ways to give 16 links for the maximum hypercube configuration. Aggregate external bandwidth for a single node is 5 Mbyte/sec when 4 input and 4 output channels are active simultaneously.

The Mark II version of the T Series machines (due in 1988) will use the T800 Inmos Transputer which is a 10 mips processor, with 4 Kbytes RAM and 1.7 Mbytes/sec bidirectional links. It also has a 1.5 Mflop floating-point unit.

Memory: Each node has a local memory of 1Mbyte of dual-ported RAM that will be increased to 4 Mbytes in the Mark II machine, with further upgrades to 16 Mbytes later.

Vector processor: The vector processor is a proprietary machine with its own instruction stream, which incorporates a 6-stage 8-Mflops adder and a 7-stage 8-Mflops multiplier. using the Weitek floating-point chip set with a cycle time of 125 nsec. The bandwidth to/from memory is 192 Mbytes/sec.

On the Mark II machine, the vector processor will be upgraded to an 18 Mflops (64-bit arithmetic) engine, with a bandwidth to/from memory of 320 Mbytes/sec.

Maximum number of nodes that can be connected is 2^{14} (16384), giving a peak potential execution rate of 262 Gflops for 64-bit operands.

Eight nodes with one system node and disk make up a module. Two modules make up a cabinet. The maximum configuration has 1024 cabinets.

Configuration: The system is hosted by a DEC MicroVAX II which is included as an integral part of each T Series system.

A cabinet contains two system disks which the user may reference through the system node network. I/O peak transfer rate 80 Mbytes/sec for a 16-node cabinet system.

Interconnection to other systems is through an Ethernet interface on the MicroVAX although work is in progress to provide a VME bus interface.

The minimum system is a single cabinet model T 20 comprising 16 processing nodes with a maximum peak performance of 192 Mflops. It weighs around 300 lb, consumes 1.7KW, and has a footprint of 5 sq. ft, with dimensions 24.1" w x 24" d x 58" h. The largest model is a T 40000 with 1024 cabinets although the largest so far delivered is the T 200 (128 processors) at Los Alamos.

Software: The T Series runs under the ULTRIX operating system on the MicroVAX front-end. Comprehensive libraries are included for data partitioning and distribution, dynamic

configuration of processing nodes into application topologies, structured asynchronous communications and vectorized mathematics.

Languages: Fortran, C, and OCCAM 2.

Performance: The T 100 can perform a matrix multiply at 596 Mflops and can solve a linear system at 135 Mflops. A quantum Monte Carlo benchmark on the T 20 at Daresbury ran only 1.7 times slower than a CRAY 1S.

Contact:

Tom Bauer
FPS Computing Inc.
Beaverton, OR 97005
1-800-547-1445

David A. Tanqueray
FPS Computing U.K. Limited
Apex House
London Road
Bracknell
Berks RG12 2TE
England
0344-56921 Telex (851) 849218 FPS UK G

**Galaxy YH-1 (Chinese Supercomputer)
Vector Register Architecture**

Architecture: China has built its first supercomputer, as was revealed by *China Pictorial*. The development of this machine, which has the appearance of a CRAY-1 computer, started in 1978 at the University of Defense, Science and Technology in Changsa.

Performance: The YH-1 (Galaxy), as it is called, can execute 100 million operations per second.

Status: According to *China Pictorial*, the YH-1 was finished two years ahead of schedule and at only one-fifth of the planned budget.

Loral MPP

Bit-Slice Parallel Architecture

Architecture: The MPP is the product of research and development designed to evaluate the application of a computer architecture containing thousands of processing elements, all operating concurrently.

The major elements are the array unit, the array control unit, and the staging buffer. The 128x128 processing element has nearest neighbour connection with full-edge closure. The 16,384 processing elements, not including the extra columns for reliability, are simple bit-serial processors, each with a 32 element on chip shift register.

The heart of the array unit is a custom integrated circuit containing eight processing elements. A total of 2112 chips have been combined with commercial memory on control chips to give the capability to perform 400 million floating-point operations per second.

The array control unit contains all the logic to provide a pipeline of commands to the array unit, an I/O controller, and a custom-built 16-bit high-performance microprocessor for program management. The staged buffer is a 16-Mbyte, multidimensional I/O buffer. This unit has the capability necessary to reformat input date into the bit plane format of the MPP I/O system. The staging buffer has an external input rate of 40 Mbytes and an internal transfer rate to and from the array unit of 160 Mbytes in each direction.

Languages: Parallel Pascal

Status: The Massively Parallel Processor was delivered to NASA Goddard Space Flight Center in May 1983.

Contact:

Loral Defense Systems Division
1210 Massillon Road
Akron, OH 44315
216-796-4511

GOULD NP1

The company is no longer marketing this product.

Vector Register, Parallel, Shared-Memory Architecture

Architecture: The NP1 is a stand-alone multi-CPU minicomputer. It can have up to four CPUs, 2 Gbytes of memory, up to 500 Gbytes disk space, and all the usual peripherals. Each CPU is independent. Coarse-grain parallelism is supported using FORK and VFORK.

The CPU is based on Gould proprietary ECL gate array technology. Two CPUs can be attached to a 64 bit wide, 154 Mbytes/sec system bus. Each CPU consists of a scalar and an optional vector processor (arithmetic accelerator). Up to four system buses can be tightly coupled.

Up to 512 Mbytes of 16-way interleaved memory can be attached to a single system bus (in modules of 16 or 64 Mbytes) using 1 Mbit DRAMs. Each CPU can access memory on any of the system buses, giving a total of 1 Gbyte of real physical memory. The virtual memory space is 4 Gbytes.

Each CPU has a 32 Kbyte cache memory. The NP1 arithmetic format has less precision but greater range than the IEEE-754 standard.

I/O is handled through Universal I/O Micro-engines (UIOM) providing 80 Mbytes/sec links to the system bus. Each UIOM will support multiple Intelligent Peripheral Interface (IPI) buses. Each IPI will support up to eight I/O modules, which could be Gould's Disc Intelligent Module (DIM) controlling up to four disc drives (with SMD, SMD-E or XMD interfaces) at 3 Mbytes/sec or VME for peripherals (line printer, magnetic tape, 8-line asynchronous controllers, Ethernet etc.). New parallel connections are being developed for supercomputer connections (100 Mbyte/sec) and high-speed disk (20-40 Mbytes/sec). Ethernet TCP/IP is supported, as is X25 with color book, and 3270 BSC, and Cray HSX. DECnet support under TCP/IP is expected soon, and Gould is committed to OSI.

Software: Gould's UTX/32 operating system combines all features of AT&T version V and BSD 4.3 with extensions for simultaneous real-time support. UTX/32 on the NP1 rates as a 'C2' class operating system. Gould supports IEEE 1003.1 - POSIX standard.

Languages: Highly optimized C and Fortran 77 compilers are available, and Ada, Pascal, LISP, COBOL, BASIC, and Assembler are also supported. Gould Common Fortran (GCF) is an extension to Fortran 77 that includes most VMS extensions, an 8x intrinsic function library, and interlanguage communication support.

Application software includes DI-3000 CORE and GK-2000 graphics, PVI metafile system, INGRESS, Q-Office, Q-Office+, NAG, oracle, 20/20, Informix, and Unify.

Performance: NP1 peak performance is 14 mips, 40 Mflops (20 in double precision) using the vector processor. The following benchmarks were obtained on a single CPU NP1 system with a arithmetic accelerator and 64 Mbytes of 8-way interleaved memory: 14.1 Whetstone (single precision), 10.2 Whetstone (double precision), 5.3 Mflops LINPACK (double precision coded BLAS) and 167000/sec Dhrystones. Compilation speeds average 17800 lines/min for GCF and 8000 lines/min for the GCF vector pre-processor.

Status: The Gould NP1 was launched in March 1987, and the first NP1 was delivered to Purdue University in June. The first UK order from a major aerospace company has just been received. The University of Bath has placed an order for a 3-processor system to provide central teaching and research facilities. The University of Edinburgh has also installed a single-processor system. The cost varies from around \$250K for a single CPU system to \$1.8M for a fully configured 4 CPU system with 1 Gbyte of memory.

Contact:

Gould Inc.
Computer Systems Division
15378 Avenue of Science
San Diego, Calif. 92128-3407

Northern Headquarters
Regent House
Heaton Lane
Stockport Cheshire SK4 1BS
England
061-480-0907

Graeme Boydell

Hitachi S-810 and S-820 Vector Register Architecture

Architecture: The S-810 comes in three models: the S-810/5, the S-810/10, and S-810/20 (not available in the United States, only for the Japanese market). The S-820 comes in two models: the S-820/60 and the S-820/80. These may be available outside Japan.

Hitachi's approach has been to employ independent scalar and vector processors. The S-810/20 relies on their current top-of-the-line mainframe (the M280H) for their scalar processor, with a cycle time of 28 nsec, and runs the complete IBM 370 instruction set. The vector unit was designed with a cycle time of 14 nsec. The main memory capacity of the S-810/20 is 256 Mwords.

The model 20 has four floating-point add/logical units and eight combination multiply/divide-add units. In addition, there are two load pipes and two load/store pipes to/from memory, each capable of loads/stores at a rate of two words (64 bits) per cycle.

The vector register capacity is 32 registers, each with a fixed length of 256 elements (64 bits). A unique feature of the Hitachi design is that vectors greater than 256 elements are managed automatically by the hardware.

The more recent S-820 incorporates several enhancements to the S-810 range. Among these are a cycle time of 4 nsec, a change of density from 2K to 5K gates/chip with the delay time reduced from 250 to 200 picoseconds. The PCBs now have 22 instead of 14 layers, contain 100K rather than 50K gates, and have their cable delay reduced from 5nsec/m to 3.8 nsec/m. The number of pipes has been increased to 18 (4 add/logical, 4 multiply/add, 1 divide, 1 mask, 4 load, and 4 load/store) on the model 80. On the 60 all sets of 4 pipes are reduced to 2. There are 512 32 (64-bit) word vector registers.

Configuration: A memory of 512 Mbytes is available, and an expanded memory of up to 12 Gbytes with a transfer rate of 2 Gbytes/sec can be supplied. There are 64 I/O channels, rated at 6 Mbytes/sec for an overall transfer rate of 288 Mbytes/sec.

All machines are air cooled.

Performance: The scalar speed of the Hitachi S-810 may be slower than either the CRAY X-MP or Fujitsu VP-200. A peak performance of 3 Gflops is claimed for the model 80 and 1.5 Gflops for the model 60.

Software: Advanced Editor System for Programming Environment (ASPEN) and interactive timing aids.

Application: Application packages include the structural packages MATRIX/HAP and ISAS II/HAP, GRADAS for graphics, HICAD3D for CAD. A partial differential equation solver DEQSOL is integrated with the SGRAF 3-D graphics system.

Status: The first installation of an S-810 was an S-810/20 at the University of Tokyo in November 1983. The S-820 range was announced in June 1987, with first shipments in November 1987.

Contact:

Yoshihiro Koshimizu
Hitachi America Ltd.
Computer Division
950 Elm Ave.
Suite 100
San Bruno, CA 94066-3094
415-872-1902

IBM 3090/VF

Vector Register, Parallel, Shared-Memory Architecture

Architecture: The IBM 3090 is the top-end system available from IBM. It uses the System/370 Extended Architecture for scalar operations. The vector facility option (VF) was announced in October 1985.

The current models suffixed by S, as in 600-S, have a cycle time of 15 - 18.5 nsec. All machines are constructed using TCM (thermal conduction module) circuitry. Memory is organized in wafers, each containing 110 1 Mbit chips.

3090 Model 120S is a uniprocessor with 18.5 nsec machine cycle, up to 64 Mbytes of central memory, and up to 256 Mbytes of expanded storage.

3090 Model 150S is a uniprocessor with 17.75 nsec machine cycle, up to 64 Mbytes of central memory, and up to 256 Mbytes of expanded storage.

3090 Model 180S is a uniprocessor with 15 nsec machine cycle, up to 128 Mbytes of central memory, and up to 256 Mbytes of expanded storage. 32 Mbytes or 64 Mbytes of central memory and 64 up to 256 Mbytes of expanded storage

3090 Model 200S is a dyadic processor with 15 nsec machine cycle, 256 Mbytes of central storage, and up to 1024 Mbytes of expanded storage.

3090 Model 250S is a two-way multiprocessor with 17.75 nsec machine cycle, up to 128 Mbytes of central storage, and up to 512 Mbytes expanded storage.

Model 280S is a two-way multiprocessor with 15 nsec machine cycle, up to 256 Mbytes of central storage, and up to 512 Mbytes of expanded storage.

3090 Model 300S is a three-way multiprocessor with 15 nsec machine cycle, 256 Mbytes of central storage, and up to 1024 Mbytes of expanded storage.

3090 Model 400S is a four-way multiprocessor with 15 nsec machine cycle, 512 Mbytes of central storage, and up to 2 Gbytes of expanded storage.

3090 Model 600S is a six-way multiprocessor with 15 nsec machine cycle, 512 Mbytes of central storage, and up to 2 Gbytes of expanded storage.

All 3090 15 nsec models have a high-speed cache of 128 Kbytes per processor. The cache is system controlled. The 120, 150, 170, and 250 models all have 64 Kbytes cache per processor.

Vector Facility (VF):

Optional feature to all 3090 models.

Pipelined vector processor with vector registers.

Each VF has 8 vector floating-point registers for 128 64-bit elements or 16 vector floating-point registers for 32-bit elements.

The 120, 150, 170, and 250 models have 256 elements in a register.

For the VF, 171 vector instructions are added.

In the VF, 32-bit operands are treated as 64-bit operands.

Fixed-stride addressing on vectors is allowed, as well as indirect addressing or mask control.

Models 120, 150, 170, and 180 can have 1 VF added.

Model 300 can have one, two, or three VFs added.

Model 400 can have one, two, three, or four VFs added.

Model 500 can have one, two, three, four, or five VFs added.

Model 600 can have one, two, three, four, or six VFs added.

ES/3090 S models offer up to 40 percent increased numerically intensive computing vector performance over comparable ES/3090 E models through doubling the vector register section size on the larger models, enhancement of the vector divide instruction (3-5 times), increased high-speed buffer, decreased machine cycle time (15 nsec), and a faster scalar multiply.

ES/3090 provides the virtual address capability of up to 16 terabytes. MVS/ESA permits access to this virtual address capability both at the assembler level and through the use of the language-independent data windowing services interface.

Configuration:

Power:

7.8 KW

Footprint 171 sq ft.

Closed water/air cooled.

Performance: Each machine with 15 nsec machine cycle has a theoretical peak performance of 133 Mflops, giving a peak performance for the 600S of 800 Mflops.

IBM's philosophy on performance is to build a very fast scalar machine with a vector facility that typically gives a speedup of about four over the scalar code. This is based on the premise that most scientific codes are not greater than 80 percent vectorizable.

Software:

MVS/XA

VM/XA

VM/SP High-Performance Option

AIX/370

Languages:

Assembler H Version 2

VS Fortran 2 including Library Program Multitasking Facility and Interactive Debug.
The Fortran compiler will automatically vectorize existing codes.

Fortran characteristics: Multitasking features are available to the MVS Fortran programmer via explicit library calls. IBM has recently made available an automatic parallelizing compiler (PFP) on a limited basis. This PFP compiler implements a rich set of functions which includes microtasking, loop detection, task creation, and subroutine dispatching with synchronization locks.

Applications: Engineering and Scientific Subroutine Library.

On ES/3090 Models 180S through 600S, more than 90 vector application programs offered by more than 50 vendors and program owners in such areas as structural analysis, computational chemistry, fluid dynamics, seismic/reservoir modeling, quantitative analysis, computational techniques/mathematical analysis and simulation.

Status: Prices range from around \$1.5M for a 120 model with a VF. The 200E is around \$3M, and the 400E is around \$8M. The cost of the first VF is around \$.28M and subsequent VFs are around \$.17M. Roughly speaking, a VF option is 10 percent per processor additional cost.

Future plans include extending the vector processor facility to more machines, more processors on the 3090, optical linking of two 3090 model 600s (at Cornell), and reduction of cycle time. IBM is spending much effort and a great deal of publicity on the Numerically Intensive Centers but is, as usual, very tightlipped about their long-term plans for parallel machines. IBM Europe will be funding at least five major centers, and will also provide VF support for about twenty-five other centers (about five in UK).

Contact:

Mr. Troy Wilson
IBM Kingston New York Laboratory
962/263
Neighborhood Road
Kingston, NY 12401

Adrian Futado
IBM United Kingdom Limited
PO Box 18
Normandy House
Alencon Link
Basingstoke
Hampshire RG21 1EJ
England
0256-56144 ext 5974

Intel iPSC/2

Hypercube Architecture

Architecture: The original iPSC/1 machine was developed from Caltech work on the Cosmic Cube. The iPSC/2 is Intel's second generation of concurrent computer. Although physically connected in a hypercube topology, the iPSC/2 incorporates Direct-Connect routing hardware that avoids the intermediate node store-and-forward overhead of the earlier systems and significantly reduces the message start-up time. Thus processors are not interrupted for routing messages and, from a programming and performance viewpoint, the configuration acts like a fully connected graph rather than a conventional hypercube.

The iPSC/2 is available with 4, 8, 16, 32, 64, or 128 nodes. Each node consists of an 80386 CPU, an 80387 floating-point coprocessor, and from 1 to 16 Mbytes of memory. The 80387 has IEEE arithmetic with 32, 64, and 80 bit formats and a speed of about 350 Kflops. Each node also has 8 bidirectional communication channels rated at 2.8 Mbytes/sec per channel in each direction. One of the channels is reserved for I/O (or for connection to the System Resources Manager in the case of node 0), and the others are used for inter-node communication.

The basic system may be enhanced by adding memory modules, the SX scalar extension (Weitek 1167) providing 1.1 Mflops at 32-bit (0.625 Mflops at 64-bit) per node and the VX vector extension providing 20 (6.6 at 64-bit) Mflops peak per node. Nodes can support both the SX and the VX options simultaneously.

The vector extension, or VX, boards consist of two 100 nsec cycle, pipelined floating-point units, one for addition/subtraction and one for multiplication, an additional Mbyte of 250 nsec data memory, and 16 Kbytes of 100 nsec fast data memory. The speed of vector operations is determined largely by the memory speed. For example, a DAXPY involving long-precision vectors in the large, main memory has a peak rate of 2.6 Mflops on a single node, while a dot product involving short-precision vectors in the small, fast memory can approach 20 Mflops.

The computational facilities of the iPSC/2 system can be extended with large-scale mass storage with the Concurrent I/O facilities for the iPSC/2. The initial system will support as much as 40 Gbytes of formatted data. The I/O facilities expand the iPSC/2 architecture to include parallel arrays of 5 1/4 in. Winchester disk drives, or other peripheral devices, managed by I/O processing nodes which are part of the Direct-Connect communications network. The architecture supports up to 7 Winchester disks for each I/O node, connected via an SCSI bus, each with formatted storage of 676 Mbytes, in configurations numbering up to 127 I/O nodes. Thus, the maximum capacity for the architecture is 343 Gbytes.

The Concurrent I/O facilities are paired with the Concurrent File System (CFS) software, which makes the parallel disks transparent to the user. User files are automatically partitioned across the disks, allowing very large files and concurrent file access. CFS caches file blocks (4 Kbytes each) in the 4 Mbyte memory on each I/O node. The Unix standard I/O library of file system calls is supported by CFS.

The I/O facilities can also be extended with VME or MULTIBUS II devices with a bus

interface card that can connect to each I/O node. The facilities are housed in a new iPSC/2 cabinet ("full-size") with standard 19-in. racks, each holding up to 16 disks and 8 I/O nodes.

Configuration: The iPSC/2 includes a front-end System Resource Manager (SRM). The SRM acts as a gateway to TCP/IP Ethernet networks, as a compile server for remote hosts, and as a prime host for directly connected terminals. The SRM includes an 80386/80387 processor, 8 Mbytes of memory, a 140 Mbyte Winchester disk, 1.2 Mbyte floppy, 60 Mbyte cartridge tape, a keyboard, and a monitor, in addition to the TCP/IP Ethernet connection.

Software: The operating system for the SRM is Unix V.3, and the message-passing operating system on each iPSC/2 node is Intel's NX/2 node executive. The Concurrent File System is co-resident with NX/2.

Languages available: The languages supported by Intel for the iPSC/2 system include Fortran 77, C, Common Lisp, and Ada. Also available is VAST-2, a Fortran vectorizer from Pacific Sierra Research for VX systems; DECON, a concurrent debugger for the iPSC/2; and the iPSC/2 Simulator (hosted on any Unix system). The collection of software tools is part of the Concurrent Workbench, a development environment hosted on Sun workstations. The environment supports dynamic allocation of subcubes of the larger system, as well as remote compilation and loading.

Independent developers have implemented other environments for the iPSC/2, including Interwork-II, a discrete event simulation environment from Block Island Technologies, and Strand, a concurrent logic language from AI Limited in England.

Numerical libraries are available for many basic operations including the solution of linear equations (LINPACK, 130 Mflops on a 64-node iPSC/2 VX in 64-bit precision) and FFTs (276 Mflops for a 1024 x 1024 FFT on the same model, in 32-bit precision). The peak performance of the maximum 128-node VX configuration is 2560 Mflops.

Applications: Applications software includes NEKTON (fluid dynamics and heat transfer), PASSAGE (flow through complex stationary and rotating passages), ADDS (die design system for the metal forming industry), FLO 87 (successor to the FLO 57 aerodynamics code from Princeton), VSAERO (subsonic aerodynamic flow), and HYPERNEWTON (molecular dynamics). Numerous other packages are in development, and many non-commercial applications and utilities are available through the user group.

Status: Systems are offered in two sizes of cabinets: 16" x 16" x 49", 3.4 KW, 215 lb (compact), and 21" x 26" x 62", 4.6 KW, 475 lb (full-size).

iPSC/2 systems are available from \$100,000, with discounts available for qualifying educational establishments. Sample prices for some iPSC/2 configurations are given in the following table:

System	Nodes	Memory per node	Price
iPSC/2 D4	16	1 Mbyte	\$ 187K
iPSC/2 D4	16	16 Mbytes	\$ 433K
iPSC/2 D6	64	16 Mbytes	\$ 1.6M
iPSC/2 D7	128	8 Mbytes	\$ 2.1M
iPSC/2 SX D7	128 scalar	8 Mbytes	\$ 3.5M
iPSC/2 VX D3	8 vector	1 Mbyte	\$ 234K
iPSC/2 VX D4	16 vector	4 Mbytes	\$ 464K
iPSC/2 VX D7	128 vector	8 Mbytes	\$ 4.1M

	I/O Nodes	Disk Capacity	Price
Concurrent I/O	1	2 disks, 1.3 Gbytes	\$ 61K
Facilities	2	4 disks, 2.6 Gbytes	\$ 94K
	8	16 disks, 10.4 Gbytes	\$ 300K
	32	64 disks, 41.6 Gbytes	\$ 1.2M

The VAST-2 vectorizer is priced at \$10K, Lisp varies from \$10K to \$30K, and Ada from \$15K to \$60K. The iPSC simulator is \$395.

The first deliveries of the iPSC/1 were made in July 1985; the first deliveries of the iPSC/2 were made in December 1987; the first deliveries of Concurrent I/O facilities were made in September 1988.

Current research environments using iPSC/2 systems include Oak Ridge, MIT, SERC (England), University of Wisconsin, Yale, Cornell, University of Illinois, SUNY, BRI (Canada), ONERA (France) and GMD (Germany).

Contact:

Intel Scientific Computers
 15201 NW Greenbriar PW
 Beaverton, Oregon 97006
 503-629-7600 FAX 503-629-9147

Applications Technology: David Scott
 Marketing Manager: Charlie Bishop

EUROPE

Intel Scientific Computers
 Intel Corporation (UK) Ltd
 Pipers Way
 Swindon SN3 1RJ
 England
 0793-696578 Telex 444447/8 FAX 0793-641440

International Parallel Machines Inc. (IP-1) Parallel Cross-Bar Architecture

Architecture: The IP-1 has between 1 and 33 proprietary MOS technology CPUs which have access to a common memory through an interconnection switch. The combination of cross-bar switch and a multi-access memory (developed from work on the Goodyear Aerospace Staran system) avoids the bottleneck associated with bus-based systems. The memory bandwidth is 213 Mbytes/sec even with a full configuration. Floating-point performance is obtained through an optional floating-point accelerator based on Weitek chips.

Standard systems have from 8 to 264 Mbytes of memory although the 32-bit address can be extended to 48 bits for a potential of 256 Terabytes of memory. Data paths are 64 bits wide. Systems are available with from 170 Mbytes to 11 Gbytes of disk space.

Configuration: I/O is handled asynchronously from the CPUs. There can be up to 52 I/O ports. The IP-1 can be used in stand-alone mode but can have front-end machines, for example a VAX or a Sun. Other equipment can be interfaced through standard VMEbus connections, including any UNIX-based workstations or Prolog or Lisp based symbolic processing workstations. Peripherals include 1/2-inch tape drives, multiple disk drives running in parallel, plotters and printers, and a close-coupled high-speed communication interface to other CPUs. Ethernet (TCP/IP) is supported.

Software: The operating system is compatible with UNIX System V.3, supporting up to 64 users, with real-time extensions available. There is a C-callable parallel math routine library, vectorizing C and Fortran 77 compilers, and an IP-1 virtual machine package for software developers.

Applications: Applications software includes tools for database management, printed circuit board layout, oil reservoir simulation, seismic data analysis, chemical modelling, computational fluid dynamics, aerospace simulations, and structural analysis. International Parallel Machines also specialize in turnkey systems and will port any serious application depending on market potential.

Performance: Peak performance of the 9-processor system is 150 Mflops double-precision IEEE, while that of the 33 CPU machine is just under 600 Mflops double-precision IEEE.

Status: The IP-1-9 system (with 9 CPUs) costs \$220,000 with peak performances of 100 mips and 150 Mflops. The entry level IP-1-1, rated at 8 mips/16 Mflops, in a cabinet with a 21 slot Eurocage, 8 Mbytes memory, a 128 Kbyte WCS, a (formatted) 170 Mbyte disk, 1 CRT, IPOS, and C compiler, costs \$49,000. The 33 CPU 264 mips/528 Mflops IP-1-33, with 264 Mbytes of memory costs just under \$1M.

The dimensions of the standard 9-processor configuration are 49" x 20" x 22". The weight is 200 lb, and the power consumption a little over 1 KW.

The first machine was delivered in October 1985. Many sales through multiple OEM contracts.

Contact:

Technical:

Dr. Robin Chang
President
International Parallel Machines, Inc.
700 Pleasant Street
New Bedford, MA 02740
508-990-2977; Telex 888648

Sales:

Robert E. Hydrisko
Strategic Markets Director
International Parallel Machines Inc.
700 Pleasant St.
New Bedford, MA 02740
508-990-2772; Fax: 508-9966726

Isis

The company is no longer marketing this product.

Pipelined vector processor with multiple processing elements

Architecture: The Isis supercalculateur was designed after an extensive survey of French computational scientists to ascertain their principal requirements for a supercomputer. Although constructed and marketed by the Bull company, it can be thought of as a French national project.

RISC architecture.

Central processing unit consisting of four independent scalar elements connected to one vector unit. Each scalar processor has its own units for floating-point and integer operations, 256 general registers with three simultaneous accesses and a cache of 256 instructions to handle branch conditions. The cycle time for these scalar units is 15 nsec, and they are rated at 33 mips. The scalar unit can initiate (or spawn off) tasks to the vector unit.

There are 8 to 32 elementary processors, which can function simultaneously, in the vector unit. Elements of a vector are normally (and automatically) assigned across the processors in a wraparound (or folded) fashion. Each processor can do a floating-point or integer operation every 30 nsec and has 256 registers to store scalars or vectors. Logic is in 100K gate arrays, with 2,500 gates on chip. Delay time is 350 picoseconds. Each vector processor has a peak performance of 33 Mflops.

Main memory of 8 Mwords (64-bit words), arranged in 16 interleaved banks with a throughput rate of 266 Mwords/sec. Memory is in static MOS with a 35 nsec cycle time. There is hardware indirect addressing.

The secondary memory has up to 64 Mwords (64-bit words) with a possible extension to 256 Mwords. Its speed is the same as the main memory but with a latency time four times as great. It is accessed in blocks and used as an I/O cache. Secondary memory is in dynamic MOS with a 120 nsec cycle time.

Configuration: The I/O controller is built around an SPS7 machine and manages the 4 Gbytes of fast-memory disks. Its peak speed is 100 Mbytes/sec. The I/O controller communicates with the outside world through a hyperchannel.

The Isis is connected to its front-end machine through the hyperchannel.

Software: Own version of UNIX. Batch oriented, Math library under development. BLAS, IMSL, and NAG are available.

Languages: Fortran 8x, Assembler. Pacific Sierra VAST-2 is available.

Applications: Principal applications covered by the survey included finite-element techniques, finite-difference methods, spectral methods, and Monte Carlo calculations.

Contact:

Claude Timsit
Bull DRTG
Rue Jean Jaures
78340 Les Clayes-sous-Bois
France

**Kendall Square Research
Shared Memory Parallel Processor**

Architecture: The computing system being designed at KSR is a parallel architecture with hierarchically managed shared memory. Main design goal of the architecture is scalability of memory capacity and performance.

Contact:

Jim Rothnie
President
Kendall Square Research
170 Tracer Lane
Waltham, MA 02154
617-895-9400

Loral System 500

Data Dataflow Architecture

The company is no longer marketing this product.

Architecture: The Loral System 500 is a real-time data acquisition and parallel processing system. Communication between processors is via dataflow tokens. Communication is handled through a 12 Mwords/sec 48-bit time multiplex bus. This bus, the MUXbus, is used to broadcast dataflow tokens that have 16 bits of tag and 32 bits of data. The processing subsystem for data manipulation and compression consists of from 1 to 64 Field Programmable Processors (FPPs).

Each single-card FPP includes .25 Mbytes of data memory and .25 Mbytes of program memory. Each FPP has a throughput of approximately 738K parameters per second.

Configuration: Real-time I/O controllers (e.g., telemetry, analogue to digital converters, MIL-STD-1553 and RS232 serial data, disk, and tape) share the data flow bus with the parallel processors.

The System 500 is connected to color graphics workstations by a 10 Mbit/sec Ethernet network for command, control, and display.

Applications: A full complement of standard telemetry processing algorithms, including bit matching, engineering unit conversion, fast Fourier transforms, and power spectral density is available.

Applications program development and system control is through an Ethernet link attached to a VAX station 2000 under Ultrix and an X Window System.

Languages: FPP programs may be developed in Fortran 77, C, or Assembler on the VAX and downloaded to the FPPs. Parallelizing, vectorization, and optimization techniques are used by the compilers, and there are utilities including linkers, loaders, debugging tools, and simulators. Heavy use of chaining is made by the compilers.

Performance: Each FPP is based on a Weitek XL numerical processor with a rating of 8 mips and 20 Mflops. The performance on the single-precision LINPACK benchmark on a single FPP is 1.5 Mflops.

Contact:

Loral Instrumentation
8401 Aero Drive
San Diego, CA 92123-1720
619-560-5888 Telex 695222

Paul J. Friedman - marketing manager

MEIKO

Parallel MIMD Architecture

Meiko was founded in 1985 to exploit the availability of low-cost, high-performance microprocessors to build parallel computers. Its first product, "The Computing Surface," is a flexible parallel computer based on the Inmos transputer. The company founders include those originally responsible for transputer implementation at Inmos with combined expertise in VLSI processor design, system design, compiler writing, and application programming.

Architecture: The Computing Surface is an MIMD parallel processor. The number of processors in a system is variable, with no upper limit. Entry level systems have 4 processors; the largest operational machine to date has 300.

The basic compute node has the following specification:

Processor	20 MHz IMS T800, T414 processor.
Memory	4 Kbytes high-speed on-chip memory tightly coupled to processor. 1M, 2M, 4M, 8M, 16M, 32M, 48M byte error-checked local memory. Direct mapped memory.
Floating point	IEEE standard, single and double length.

Communication between nodes is by high-performance serial links. Connectivity is flexible and may be manually or electronically configured. Manual configuration requires the use of a patch panel to wire up configuration. Electronic configuration establishes required connectivity for a given program automatically. Four point-to-point links per processor allow various topologies such as rings, grids, low-order hypercubes, and pipelines to be constructed. Microcoded communications instructions in the processor give very low set-up costs for message transfers.

Communication performance is as follows:

Baud rate	10 Mbits/s, 20 Mbits/s
Data rate	2.8 Mbytes/s full duplex 1.7 Mbytes/s unidirectional (20 Mbits/s)

Configuration: Machine may be hosted from VAX, microVAX, Sun workstation, or IBM PC. Self-hosted systems are also available. Basic system has dual RS232 lines plus GPIB interface as standard. Additional peripherals may be added with appropriate controller boards. Controller boards share common architecture with compute nodes and use same processor.

Available boards:

Mass store	4 Mbytes/s DMA SCSI interface, T800 or T414 processor, 8 Mbytes RAM. Allows any SCSI peripheral to be connected. Range of SCSI devices available includes disk drives up to 600 Mbyte capacity and various tape and cartridge drives.
Graphics	Programmable graphics controller. Supports various screen formats. Maximum pixel rate 110 MHz. PAL/NTSC broadcast quality graphics available. 2 Mbyte video memory, configurable as 8 or 24 bits per pixel. T800, T414 local processor. 4 Mbytes local memory.
Frame Grabber	1 Mbyte dual ported frame store. Inputs RGB/monochrome. Sample rate up to 80 MHz, line rate up to 70 kHz. Local T800, T414 processor.
Data Port	T800 or T414 based processing node with 0.5 Mbytes dual ported memory and 80 Mbytes/sec parallel I/O port.
Local Host	System processor self-hosted systems. T800 processor with 8 Mbytes local memory. Ethernet interface, DMA SCSI interface and dual RS232 ports.

Software: Basic, remote hosted systems run the Meiko Development System, MDS. This provides editor, parallel linker and loader, and run-time monitor for a single-user system. Multi-user capability available under MVCS, Meiko Multiple Virtual Computing Surfaces, which allows the machine to be partitioned into multiple independent domains. Self-hosted systems under beta test use UNIX-like environment and file system.

Languages available: C, FORTRAN 77, PASCAL, and OCCAM II.

Fortran characteristics: FORTRAN 77 is standard and does not require vectorizing. Various parallelization strategies are used depending on data parallelism available in the problem.

Applications: Current applications running on Computing Surfaces include finite-element analysis, lattice gauge theory simulations, ray tracing, molecular modelling, seismic data processing, reservoir simulation, image processing, and fingerprint recognition.

Performance: Integer performance of T800- and T414-based machines is 10 mips per node. Floating-point intensive applications require T800 processors, with 1.5 Mflops achievable in

32-bit floating-point arithmetic. LINPACK benchmark on a single node yields 0.5 Mflops in single precision.

Status: First machine installed March 1986. Installed user base as of March 1988 is 120 machines. Largest machine to date is Edinburgh Concurrent Supercomputer, at the University of Edinburgh, which when fully populated will have 1 Gflop of processing power and 4 Gbytes main store. Other customers include GE; Automation and Robotics Research Institute, Fort Worth, Texas; and several UK universities.

Price dependent on configuration. Entry level systems from around \$30,000.

Contact:

Meiko Limited
650 Aztec West
Bristol BS12 4SD
England
0454 616171 Fax (0454) 618188

Meiko Scientific Corp.
400 Oyster Point Blvd., Suite 523
South San Francisco, CA 94080
415-952-9900 Fax 415 952 7092

Contact: Moray McLaren (England)

Multiflow TRACE 7/200, 14/200, and 28/200

Scalar VLIW (Very Long Instruction Word) Computer

Architecture: Designed for Trace Scheduling Compacting Compilers. Each model within the TRACE family of processors executes instructions ranging to 1024 bits in length. The entry-level TRACE 7/200 is capable of initiating seven operations (two floating-point, four integer/load/store, and one branch operation) on each 130 ns instruction cycle. The TRACE 14/200 and 28/200 can initiate fourteen and twenty-eight operations per cycle, respectively.

The TRACE design includes no microcode and very little synchronization logic. Rather, the very wide horizontal architecture is directly exposed to TRACE Fortran and C language compilers. This allows the TRACE compilers to generate wide instructions automatically, without programmer intervention, based upon unusually wide-scope analysis of application code.

As a result, the TRACE hardware/software system presents itself to the programmer (and to the source-level program) as a conventional computer that executes a single instruction stream at high speed.

Configuration: The TRACE 7/200 and its basic peripheral complement are housed in a single equipment bay 28" w x 40" d x 60" h. The 7/200 is directly field-upgradable into the wider word TRACE models by inserting additional integer and floating-point modules into its backplane; expansion cabinetry is necessary only for additional disk and tape storage.

Attached VME I/O processors handle low-level I/O functions, allowing the CPU to operate with minimal interruption.

Register Complement: 160, 320, and 640 general-purpose 32-bit registers for the 7, 14, and 28/200, respectively

Access bandwidths of 984, 1968, and 3692 Mbytes/sec, respectively

Instruction cache: 8K instructions, independent of instruction width

Technology: CMOS functional units, 8K CMOS gate arrays, and Schottky TTL logic

Memory: 16 to 512 Mbytes capacity, with ECC (all models); up to 64-way interleaved (all models).

Four-Gbyte demand-paged virtual address space per process

Floating Point: IEEE 32- and 64-bit formats

Hardware add, multiply, divide, and square root

Power and Cooling: Requires 2.5 to 7 KVA, depending upon configuration. Operates from 10 to 35 C with air cooling.

Software: Operating System: Multiflow's adaptation of UNIX bsd 4.3, including fast file system, TCP/IP, network file system, disk striping, asynchronous disk I/O, shared libraries, and copy-on-write.

VAX Compatibility: A suite of tools that allow the TRACE to perform as an adjunct to an existing VAX/VMS-based computing environment.

Trace Scheduling compacting compilers from Multiflow exploit the fine-grained parallelism present throughout nearly all applications.

Languages: TRACE Fortran compatibility spans ANSI Fortran 66, Fortran 77, and proposed 8X; VAX/VMS Fortran; Cray Fortran; and IBM Fortran.

TRACE C improves the performance of C codes generally, including the UNIX operating system and related utilities.

Applications: Though users often focus upon highly visible, mathematics-intensive applications, O.S. kernel and library functions are comparably important.

Performance: Multiflow submits that processors in the mini-supercomputer price range ought to be evaluated in the same terms as their minicomputer cousins: by interactive performance across a broad range of applications, not by isolated vectorizable applications alone. With the advent of the TRACE, Multiflow contends that there is no reason why mini-supercomputers should not be treated simply as faster super-minicomputers.

Benchmark Performance (TRACE 7/200, Software Release 1, April 1987):

LINPACK 100 x 100 Compiled at Full Precision – 6.0 Mflops
Double-Precision Livermore Loops Harmonic Mean – 2.3 Mflops
Double-Precision Whetstone – 12605 KWhets

ANSYS Test Cases (CPU Seconds per Job): SP-1 – 31, SP-2 – 84, SP-3 – 135, and SP-4 – 51

ANSYS Large Cases (Elapsed Seconds per Job): M1 – 40, M2 – 516, M3 – 3989; S1 – 374, S2 – 1955, S3 – 5842, S4 – 13209, S5 – 12499

Peak Performance:

TRACE 7/200 – 53 VLIW MIPS, 30 Mflops / SP, 15 Mflops / DP
TRACE 14/200 – 107 VLIW MIPS, 60 Mflops / SP, 30 Mflops / DP
TRACE 28/200 – 215 VLIW MIPS, 120 Mflops / SP, 60 Mflops / DP

Status: Multiflow announced the TRACE product family April 21, 1987, for first production shipment in mid-1987. System packages start below \$300,000.

Contact:

Multiflow Computer Inc.
175 North Main Street
Branford, CT 06405

Voice and FAX: 203/488-6090

Voice only: 800/777-1428

UUCP: decvax!yale!mfci!smith

Technical Contact: John O'Donnell, V.P. of Engineering

Sales Contact: Robert Smith, V.P. of Sales and Marketing

MYRIAS

Shared-Memory Parallel Architecture

Architecture: Myrias offers the Parallel Application Management System (PAMS) and the Scalable Parallel Supercomputer (SPS-2). PAMS is a software environment that facilitates the development of applications on a shared-memory parallel architecture, and provides run time optimization of the system during execution of several programs. The SPS-2 is the hardware system in which PAMS is implemented.

The SPS-2 consists of processing elements (PE's), a Master Controller, Disk subsystems, and associated peripheral devices. A PE provides processing resource (Motorola MC68020 CPU, 68851 MMU, 68882 FPU) and memory resource (4 MBytes of SECDED DRAM). PE's are connected hierarchically: four PE's share a single bus on a multiple-processing element board with 16 Mbytes of memory; up to 16 multiple-processing (MP) element boards share two 33 MByte/sec backplane buses in a card cage; each card cage has five 11 MByte/sec communications channels for interconnection between cages or to the Master Controller. MP boards can be exchanged within a cage for input/output (I/O) boards that each connect to high-speed I/O controllers at up to 20 MBytes/sec.

PAMS defines a virtual machine which presents a constant interface to applications. The virtual machine contains a transparent control mechanism that automatically schedules parallel tasks on PEs, enables tasks to access data, levels loads across programs and PE's shared by a program, and merges the results of parallel computation. The system uses virtual memory (32 bit) addressing, within which each PE can address a 1024-Gbyte address space. The hierarchically interconnected PE's provide a transparent hierarchical memory cache for each parallel task. Thus, although there is no shared central memory, each parallel task can access a large address space.

Independent parallel tasks inherit memory images from their parent (the task that invokes parallelism), and execute in distinct memory spaces. Sibling tasks do not generally affect each others' memory spaces, although a mechanism is provided that enables communication between them.

Configuration: 64 PEs and 256 MBytes minimum. There is no known technical maximum.

Disk controllers support striped disks that transfer at up to 20 Mbytes/sec into cache memory. Controllers can connect to up to 4 I/O boards that may be in the same or different cages. Transfer rates to the I/O boards can take place at up to 20 Mbytes/sec.

Software: UNIX (POSIX)

Languages available: Myrias Parallel Fortran (MPF) and Myrias Parallel C (MPC)

Fortran characteristics:

- upwards compatible with ANSI Fortran 77
- single PARDO extension provides access to parallelism
- interactive source-language level debugger

C characteristics:

- upwards compatible with proposed ANSI standard

single PARDO extension provides access to parallelism
interactive source-language level debugger

Applications: General physical modelling (Monte Carlo and Particle-in-Cell methods, computational fluid dynamics, drug design, geophysical applications, image processing, text retrieval, and VLSI design.

Performance: Will exceed current supercomputers on an economically significant set of applications.

Status: Cost ranges from \$750K to over \$10M

Contact:

Mr. Peter A. Gregory
Myrias Computer Corporation
124 Myrtle Street
Boston, MA 02114
617-723-5727

National Advanced Systems AS/91X0

Integrated Vector Processor

Architecture: The NAS 91X0 is the top-end system available from National Advanced Systems. It uses the System/370 Extended Architecture for scalar operations.

AS/9140/50 are uniprocessors with 48 Mbytes of central memory.

AS/9160 is a uniprocessor with 64 Mbytes of central memory.

AS/9170/80 are dyadic processors with 64 Mbytes of central memory.

Each processor has a high-speed cache for scalar operands. The cache is system controlled.

Vector Processing Facility (VPF): Optional feature to the 91X0. 46 vector instructions are added for the VPF. 32-bit operands in the VPF are treated as 64-bit operands.

Pipelined vector processor using memory to memory operations (no vector registers).

Fixed stride addressing on vectors is allowed as well as indirect addressing or mask control.

Based on the Hitachi S-9 plus IAP.

Configuration: Closed water/air cooled.

Software: MVS/XA, VM/XA, VM/SP High-Performance Option

Languages: Assembler H Version 2

The Fortran compiler will automatically vectorize existing codes using Pacific Sierra's VAST.

Status: Rough cost is \$3M.

Contact:

Claud Stoudmeyer
National Advanced Systems
800 East Middlefield Rd.
PO Box 7300
Mountain View, CA 94039
415-962-6100

NCUBE

Hypercube Architecture

Architecture: Node Processor

Custom VLSI

11 Interrupt driven DMA channels at 2 Mbytes/sec

10 channels for hypercube; 1 for system I/O

VAX style 32-bit byte addressable architecture

16 general registers (32 bits)

complete, orthogonal 2-address instruction set

8, 16, 32-bit integer and logical operations

32, 64-bit IEEE floating point operations

17 addressing modes (eg. autoincr, autodecr, autostride)

Performance (8 Mhz: approx. VAX 780 with fl.pt. accelerator)

1-2 mips (32 bits); .5 Mflops (32 bits); .3 Mflops (64 bits)

Memory: 512 Kbytes SECDED in 1Mbit chips.

Processor Board (16"x22") contains 64 nodes + 8 Mbytes SECDED memory

Host Board (16"x22") contains

Intel 80286/80287 with 4 Mbytes SECDED memory

1 ESMD Disk Interface for up to 4 disks (160 330 500 Mbyte)

8 serial RS-232 channels

1 parallel Centronics compatible interface

3 iSBX interfaces

16 Node processors with memory; provide small cube for starter system or 128 DMA channels for larger system

Performance: up to 180 Mbytes/sec bandwidth to hypercube

Graphics Board (16"x22") contains 2Kx1Kx8 frame buffer (768x1024 displayed

60 Hz); color table (16 M color); 180 Mbytes/sec data bandwidth

(30 frames/sec); zoom; pan; text processor;

RS-343 RGB output

Intersystem Link Board: Connects multiple NCUBE/ten systems together

Open System Board: Allows user-defined interfaces to the hypercube.

Disk Farm Board: Allows direct disk connection to hypercube nodes.

Configuration:

NCUBE/ten: 16 to 1024 Nodes; 3 ft cube; 220 V; 8 KW max; air cooled; 24 slot backplane: 8 for I/O options, 16 for Processor Boards; 160, 330 or 500 Mbyte disk drives and 60 Mbyte cartridge tape

NCUBE/seven: 16 to 128 Nodes; 14" wide by 29" by 29"; 110 V; office environment; 4 slot backplane: 2 for I/O options,

2 for Processor Boards; 160 or 330 Mbytes disk, 16 Mbytes tape drive
NCUBE/four: 4 to 16 Nodes; PC-AT Accelerator (4 Nodes+AT bus
interface); up to 4 Boards per AT; for software development
plus workstation.

Software:

Axis (Host): Unix-style multiuser; distributed file system;
EMACS style screen editor with up to 4 windows; cube managed
as a device that can be allocated in subcubes; parallel
symbolic debugger.

Vertex (Nodes): Message passing primitives including automatic
routing; message typing; process debugging support

Fortran 77 and C are available.

Axis, Vertex and Compilers run on the NCUBE/four (PC-AT).

Status:

NCUBE/ ten or seven: \$40K (cabinets + peripherals) + \$60K/Host
Board + \$100K/Processor Board (University discount available)

NCUBE/four: \$10K/board (4 nodes) + \$4K OS licence.

First complete shipments in December 1985.

Contact:

1815 NW 169th Place
Suite 2030
Beaverton, OR 97006
602-839-7545

NEC SX-JA, SX-1EA, SX-1A and SX-2A

Vector Register Architecture

Architecture: The SX-A Series has two types of processor: up to two Central Processors (CP) and one Arithmetic Processor (AP) sharing the main memory. The CP is a front-end mainframe processor where system control programs and user programs run. The AP is a kind of Fortran engine dedicated to execution of user programs. Although the SX runs in stand-alone mode, NEC supports its ACOS series mainframes and also IBM mainframe connections.

AP Architecture

- RISC-based scalar architecture
- 16 vector arithmetic pipelines: four identical sets each with an add, multiply, logical, and shift pipe.
- 1000 gate LSIs with 250 picosecond gate delay.
- Circuits are packaged 36 to a module, and 12 modules to a board.
- 1 Kbit bipolar memory with 3.5 nsec cache memory access time.
- 1024 Mbyte memory (512-way interleaving) with 8 Gbyte extended memory.
- 256 Kbit static MOS memory chip with 40 nsec access time, giving a memory-to-register rate of 11 Gbyte/sec.
- Register-to-register machine with 80 (on the SX-2) Kbytes of vector registers and four sets of four decimal units yielding a maximum result rate of 8 flops per cycle.

Scalar arithmetic is pipelined (128 scalar registers) and operates in parallel with vector units. The NEC scalar cycle time is the same as the vector, and is segmented and pipelined to allow more than one pair of operands to progress through the same functional unit concurrently.

CP Architecture

- An extension of the NEC mainframe computer, rated at maximum 37 mips (4 types of CP prepared).
- Virtual storage support.

Configuration: A summary of the four machines is given in the following table:

	SX-1EA	SX-1A	SX-2A	SX-JA
Number Pipes	4 v-pipe	8 v-pipe	16 v-pipe	4 v-pipe
Length regs	20K v-reg	40K v-reg	80K v-reg	20K v-reg

Can be used as stand-alone machine or can link to other equipment through a hyperchannel. Separate I/O processors with an overall rate of up to 192 Mbytes/sec are included in the standard configuration.

Software:

- Uses the NEC standard operating system and UNIX called SX-UX is also supported.
- Does not run the IBM instruction set (unlike other Japanese computers)

Languages:

- Fortran 77 with automatic vectorization. Performance tuning tools available are VECTORIZER/SX and ANALYZER/SX. The compiler vectorizes IF statements, intrinsic functions, and indirect addressing using vector gather and scatter instructions (into temporaries).
- Other languages supported (but not vectorized) include ALGOL, PL/I, BASIC, PASCAL, C, LISP, PROLOG, and COBOL.

Performance: Maximum rating of the SC-JA is 250 Mflops, of the SX-1EA is 330 Mflops, of the SX-1A is 665 Mflops, and of the SX-2A is 1.3 Gflops. It appears to be the most powerful of the Japanese supercomputers, and the only one to aggressively address the scalar bottleneck.

Status: First delivery date in the US was July 1986. The NEC machine is available for benchmarking. NEC has sold 18 of its supercomputers in Japan, in the USA, and in Europe. The USA machine is situated at HARC in Houston, Texas. A machine was delivered to the Netherlands in 4Q 1987.

Contact:

Mr. S. Adams
NEC Information Systems
1414 Massachusetts Ave.
Boxborough, MA 01719
617-264-8800

Garry Foley
Manager - Marketing
Communications Systems Division
NEC Business Systems (Europe) Ltd.
NEC House
1 Victoria Road
London W5 6UL
England
01-993-8111 Telex 261914 NEC LDN

Prevec

Shared-Memory, Parallel Vector Processor

Architecture: VLIW parallel system. General-purpose 32-node supercomputer, expandable to as many processors as needed.

Configuration:

Cross-bar

Uses BIT chip for floating point

25 nsec cycle time

Performance: Livermore Loops, LINPACK - 50 Mflops per node

Status: \$50K for a single processor. Product to be available second quarter of 1990.

Contact:

J. Yoon, President
Prevec Computer Co.
3713 S. George Mason Dr.
Suite C 1 W
Forest Church, VA 22041
703-845-1800

PS 2000 (Russian supercomputer) Parallel Architecture (SIMD)

Architecture: In the Soviet Union there is assembly-line production of PS-2000 computers with a capability of up to 200 million operations per second.

The PS-2000 complex is classified as SIMD. The complex includes an SM-2 and the PS-2000 processor. The complex was first commissioned in 1980. Unspecified type of addition speed is 0.3 microseconds, with a memory access or cycle time (source gives both in heading without saying which the number applies to) of 0.64 microseconds.

The structure of the PS-2000 computer consists of 8, 16, 32, or 64 processor elements (PE). They are connected to each other in an identical fashion, are located under a unified control, and are of a single type. Each processing element has its own (local) direct access semiconductor 12 or 48 Kbyte memory. This makes it easy to upgrade the system and thus change its performance within wide limits. The performance of the minimum PS-2000 8-processor computer configuration is approximately 25 million short operations per second. The maximum PS-2000 64-processor computer configuration permits a performance of about 200 million short operations per second. The PS-2000 operates on 12, 16, and 24-bit words and can work in both fixed and floating-point modes.

Configuration: The processors of the PS-2000 can be connected under program control into a ring structure. It is possible to form two identical rings, each consisting of 8, 16, or 32 processors. These processors are controlled by the PS-2000 CPU, which uses 64-bit instructions from its own semiconductor memory. A basic 8-processor configuration fills a 28" rack. A full 64-processor 40-Mflop configuration fills 5 such racks. By comparison, the US-made 30-Mflop Numerix 432 fills half of a 22" rack.

Languages: The basic programming language for the PS-2000 is assembly, which reflects the PS-2000 microinstruction set.

Applications: While the bulk of the applications of the PS-2000 appear to be seismic data processing, other problems such as near-sonic gas flow studies and nuclear reactor simulations have been reported.

Performance: The PS-3000 array processor is designed to augment the computing capability of the SM-1210 computer, which is either a new machine or an upgraded SM-2. The PS-3000 probably is not yet in production. It will be a multiprocessor superior to the PS-2000 and capable of 100-Mflop computing rates. The PS-3000 will apparently have four parallel processors, each of which has three arithmetic units that run in parallel.

Status: retails at 800,000 rubles

Saxpy MATRIX 1

The company is no longer marketing this product.

Reconfigurable systolic architecture

Architecture: Five basic components: (1) system control unit (DEC Micro VAX II), (2) matrix processing unit (Systolic processor) with 2 to 8 Mbytes of local memory, (3) mass storage system, an I/O interface system for access to high-speed data storage peripherals, (4) system memory (64 to 2048 Mbytes SECDED), and (5) Saxpy Interconnect - a control and data bus (320 Mbytes/sec transfer rate).

Configuration: Possible configurations for the Matrix Processor are shown in the table below.

Model name	Mflops	No. of cabinets	Matrix Processor zones
MATRIX 1/250	250	3	8
MATRIX 1/500	500	3	16
MATRIX 1/750	750	4	24
MATRIX 1/1000	1000	4	32

Size: 95.2" wide x 78.2" high x 40.4" deep

Weight: 3500 lb.

Power: 15 KW at 220 VAC (60 Hz).

Air cooled

Software: VAX/VMS Operating system. MATRIX 1 Fortran is Fortran-77 with some 8x extensions. There are several Saxpy-supplied libraries. The Standard Subroutine Library provides basic system and data manipulation functions in addition to some simple BLAS level calculation subroutines; there are also some simple synchronization subroutines. The Engineering and Scientific Subroutine Library includes high-level matrix arithmetic subroutines (utilizing block algorithms for efficiency on the matrix processing unit), fast Fourier transform routines, convolution and correlation subroutines, and further utility subroutines. The Signal-Processing Subroutine Library includes spectral analysis, digital filtering, beam forming, and direction-finding subroutines.

Languages: HiMAT language in C-based and Fortran 77-based version.

Applications: Major application areas include signal and seismic processing, image processing, and numerical analysis.

Status: Machines are priced from \$695,000 to \$1.9M. A fully configured system including a MATRIX 1/1000 would cost around \$3M.

Beta units installed in April 1987. Customers include Martin Marietta Baltimore Aerospace. First production delivery made in March 1988.

Contact:

Saxpy Computer Corporation
255 San Geronimo Way
Sunnyvale, CA 94086
408-732-6700

President and CEO: Tony Yates
VP, Marketing: Joseph E. Straub
VP, Chief Scientist: Dr. Robert Schreiber
Director, Advanced Technology Group: Dr. Ben Friedlander

SCS-30/XM and SCS-40/XM

The company is no longer marketing this product.

Vector Register Architecture

Architecture:

- register-to-register CRAY-compatible architecture (all CRAY software should run on this machine)
- microcode driven emulator to emulate the CRAY X-MP instruction set.
- 64-bit scientific computer with pipelined, asynchronous functional units.
- multiple pipelined functional units.
- 45 nsec cycle time.
- 5 vector, 1 scalar, and an address calculation can execute concurrently.
- transfer rate from registers to functional units of up to 6 words/ clock cycle (1.07 Gbytes/sec).
- 256 word buffer between memory and instruction decode logic allows execution of one instruction per cycle (two cycles for conditional branch).
- supports flexible hardware chaining of functional units and memory references.

Configuration:

- 32, 64, 128, 256, 512, and 1000 Mbyte field-upgradable memory configurations with 4-16 banks.
- four ports to memory (like the CRAY X-MP, i.e., 2 vector loads and a store can be going on at the same time.)
- will interface to a front end, either VAX 11/780, VAX 11/750, or Hyperchannel.
- 2-10 programmable I/O channels, each with 16 Kbyte buffer and a transfer rate of 20 Mbyte/sec. Transfer rate of buffers to central memory is 1 word/clock period (178Mbytes/sec).
- DD-550 disk drive holds 550 Mbytes and can sustain read/write data transfer rate of 10 Mbyte/sec with an average access time (seek plus latency) of 24 msec
- maximum of eight drives can be attached to each of the eight optional I/O channels.

Other features:

- Size: 55 x 55 x 60 inches
- Forced air cooling.
- Power consumption: 208 3-phase 11-16.5 KVA
- Weight: 1 ton

Software:

- SCENIX (UNIX V5.3 compliant), SCS/COS and CTSS.
- software licensing agreement with CRAY.
- I/O and network connections performed by SCS I/O Network Nodes (IONNs) which have Direct Memory access to any SCS processor memory setting on the SCS 178 Mbyte network.
- multiuser, multiprogramming OS supports interactive job execution.

Languages:

- Fortran 77. Fortran compilation expected at 20,000 to 40,000 lines per minute. Fortran vectorizing compiler. Interactive debugger.
- Assembler.
- Vectorizing C compiler.

Applications:

- MSC/NASTRAN
- GAUSSIAN 86
- ANSYS
- ABAQUS
- FIDAP
- over 150 more applications ported from the CRAY X-MP.

Performance:

- SCS-40/XM peak vector rate of 44 Mflops and peak scalar rate of 22 mips in 64-bit arithmetic.
- SCS-30/XM peak vector rate of 33 Mflops and peak scalar rate of 16 mips in 64-bit arithmetic.
- LINPACK timings around 1/4 the performance of a single CPU X-MP.

- Matrix vector operations (subroutine SMXPY).
- around 37.6 Mflops (simulated).

Status: Prototype available 11/85; first customer shipment 4/86

Cost: Base system \$295,000.

Market target is to provide a CRAY-compatible general-purpose scientific computer that computes at 1/4 the CRAY X-MP, but has the price of a super-mini and thus the price/performance of a supercomputer.

Contact:

Scientific Computer Systems
10180 Barnes Canyon Road
San Diego, CA 92121
619-546-1212

President: Barry Rosenbaum

Pierre Hassid
Scientific Computer Systems Corporation
5 Villa Alexandrine
92100 Boulogne Billancourt
France
+33-1-48.25.73.47

Sequent Balance 8000 and Balance 21000

Parallel Bus Architecture

Architecture: Two products, Balance 8000 and Balance 21000, employing same system components and differentiated only by capacity.

Family has 2-30 NS 32032 processors running at 10 MHz, each with floating-point unit, memory management unit, and 8-Kbyte cache sharing a global memory via a 64-bit wide pipelined packet bus supporting multiple, overlapped memory and I/O transactions with a sustained data transfer rate of up to 53.3 Mbyte/sec enhanced to support future generations of system modules.

Memory: The machine has up to 28 Mbytes of physical memory, a 4-Mbyte I/O address space, and a 16-Mbyte virtual memory address space for each user process. Memory can be two-way interleaved, and there can be up to 4 memory controllers which each manage 2 to 8 Mbytes using 256 Kbit RAM components. Processor and memory boards can go in any slot on the SB8000 bus.

Configuration: A Sequent-designed IC chip (SLIC, System Link, and Interrupt Controller) resides on each board to manage interprocessor communication, synchronization, interrupts, diagnostics, and configuration control. There is an extensive diagnostic subsystem.

Industry-standard I/O, interfaces:

MULTIBUS - has terminal multiplexor and other controllers. Ethernet - at 10 Mbits/sec. Connection to PC as virtual disk through Ethernet.

SCSI - at 2.5 Mbyte/sec. Offers 5-1/4 in. disk drives (72 and 150 Mbytes formatted) and streamer tape drives with adaptor boards for the SCSI bus.

DCC - a very-high-performance SMD and SMD-E disk controller supporting up to 8 disks allowing two simultaneous data transfers at up to 3 Mbyte/sec and overlapped seeks on all drives. Performance enhanced by rotational position sensing and slip sector bad block handling.

Peripherals include 1/2" 1600 and 6250 bpi tape drives and 396, 264, and 540 Mbyte disk drives.

The Balance 8000 packaged system includes a 9-slot SB8000 bus backplane and an 8-slot MULTIBUS backplane and can take up to six dual-processor boards (12 processors).

Other features:

Table height packaging.

Dimensions 30.5" h x 23.25" w x 28.625" d

SB800 chassis 15.5" x 10.5" x 13.5"

MULTIBUS chassis 14.2" x 6.68" x 8.5"

11 amps max at 60Hz 115VAC.

Maximum configuration dissipates 1500 Watts

The Balance 21000 packaged system with a 26 slot backplane and a 12 slot MULTIBUS

supports a full configuration of 30 processors (15 dual processor boards). Dimensions are 67" h x 27.5" w x 38" d

Software: The operating system, called DYNIX, is a version of UNIX supporting a dual universe for System V.2 and BSD 4.2 UNIX applications, enhanced for application-transparent multiprocessing and user-controlled parallel processing. Among the enhancements are a completely reentrant kernel, tuneable virtual memory, user-level shared memory, and synchronization services. All processors run a shared copy of the operating system. The configuration is symmetric, and load balancing is automatic and dynamic.

Supports X.25 and ARPANET TCP/IP protocols plus all the networking facilities of UNIX 4.2. Support is also available for customer-provided application accelerators.

Languages: Ada, C, Fortran 77, ANSI-standard Pascal, Assembly language, Lisp, Prolog, Modula-2, Cobol, Mumps, Basic. Parallel programming library callable from any language. Extensions to Fortran, C, and Pascal to allow shared common blocks. Preprocessor for Fortran to parallelize DO-loops.

Performance: Entry-level Balance 8000 similar in power to Microvax II. Fully populated B21000 seen as three times a VAX 8650 (21 mips) in power. Designed as a high throughput system, with support for parallel processing at user level.

Status: Incorporated in January 1983 (original name of company was Sequel). European subsidiary established in UK in March 1986. Sequent Europe now (January 1988) has subsidiaries in the Netherlands (Amsterdam), West Germany (Munich), and United Kingdom (London). Shipments began 12/84.

Entry price for Balance 8000 system (2 processors) of \$60,000; ranging to \$500,000 for a large Balance 21000.

Contact:

Sequent Computer Systems, Inc.
15450 SW Koll Parkway
Beaverton, Oregon 97006-5903
503-626-5700 800-854-0428 Telex 296559

Casey Powell and Scott Gibson, co-founders.
Technical: David Rodgers and Gary Fielland

Sequent Europe Limited
1 Martindale Road
Hounslow
Middlesex TW4 7EW
Endland 01-570-2066 Fax 01-577-5834 Telex 946114 SQNTUK

European General Manager : Stuart Bagshaw
UK Research and Academic : Ian Blagg
UK Sales Director: Peter Winder.
UK Technical Support: Steve Wanless.

Sequent Symmetry S27 and Symmetry S81

Parallel Bus Architecture

Architecture: Two products, Symmetry S27 and Symmetry S81, employing same system components and differentiated only by capacity.

The S81 has from 2 to 30 32-bit INTEL 80386 microprocessors (2 to 10 on the S27), each with an INTEL floating-point coprocessor, 64 Kbyte two-way associative cache, memory management unit, sharing a global memory via a 64-bit wide pipelined packet bus supporting multiple, overlapped memory and I/O transactions with a sustained data transfer rate of up to 53.3 Mbyte/sec (Peak rate is 80 Mbyte/sec). Optionally, each processor can be fitted with a floating-point accelerator based on the Weitek 1167 chip. The cache uses a copy back control scheme implemented in custom designed VLSI. Each CPU board contains 2 Intel 80386 and 80387 microprocessors per board running at 16 MHz with plans to increase rate to 20 MHz.

It is possible to upgrade from the earlier Sequent machines simply by swapping boards. A buy-back scheme for the old boards is currently in operation.

Each memory controller has either 8 or 16 Mbytes of memory, while expansion boards contain 24 Mbytes each. 1 Mbyte DRAM devices are used and the maximum physical memory is 240 Mbytes. Each process has a maximum virtual address space of 256 Mbytes. Processor and memory boards can go in any slot on the bus.

A Sequent-designed IC chip (SLIC, System Link, and Interrupt Controller) resides on each board to manage interprocessor communication, synchronization, interrupts, diagnostics, and configuration control. There is an extensive diagnostic subsystem.

Configuration: Industry-standard I/O, interfaces:

MULTIBUS - has terminal multiplexor and other controllers and provides a link for up to four IEEE 796 standard MULTIBUS systems. Also MULTIBUS-based controllers for connecting RS-232 terminals, 1/2" 1600 and 6250/1600 bpi tape drives and parallel line printer controllers.

Ethernet - at 10 Mbits/sec. Connection to PC as virtual disk through Ethernet.

SCSI - at 2.5 Mbyte/sec. Offers 5-1/4 in. disk drives (72 and 150 Mbytes formatted) and streamer tape drives with adaptor boards for the SCSI bus.

DCC - a very high performance SMD and SMD-E disk controller supporting up to 8 disks allowing two simultaneous data transfers at up to 3 Mbyte/sec and overlapped seeks on all drives. Performance enhanced by rotational position sensing and slip sector bad block handling.

NFS - a version of SUN Microsystem's Network File System. TCP/IP - a standard network protocol. X-Windows, X.25, Colored book software available by 2Q 1988.

Peripherals include 1/2" 1600 and 6250 bpi tape drives and 396, 264, and 540 Mbyte disk drives.

Other features:

Dimensions 30.5" h x 23.3" w x 26.8" d (S27)
67.0" h x 38.0" w x 27.5" d (S81)
S27 16 amps max at 60Hz 115VAC.
8 amps max at 50Hz 220VAC.
Maximum configuration dissipates 1500 Watts
S81 15 amps max at 50Hz 415 VAC three phase.

Software: The operating system, called DYNIX 3, is a version of UNIX supporting a dual universe for System V.2 and BSD 4.2 UNIX applications, enhanced for application-transparent multiprocessing and user-controlled parallel processing. Among the enhancements are a completely reentrant kernel, tunable virtual memory, user-level shared memory, and synchronization services. All processors run a shared copy of the operating system. The configuration is symmetric, and load balancing is automatic and dynamic.

Languages: C, Fortran, and Pascal are fully supported and have support for parallel programming. A parallel programming library is callable from any language. There are extensions to Fortran, C, and Pascal to allow shared common blocks. There is a preprocessor for Fortran to parallelize DO-loops. Pdbx, a version of the dbx source level debugger that has been enhanced to support the debugging of parallel programs, is available.

Performance: From 8 to 108 mips. The S27 runs at 5 Mflops (single precision) on the LINPACK benchmark (with a one-line compiler directive), and the S81 at 13 Mflops (single precision) when configured with the Weitek floating-point accelerator.

Status: Symmetry machines being delivered since December 1987. By January 1988, over 50 Sequent Symmetry machines had been delivered.

Cost: Cheapest system is \$60,000, and smallest increment is an 8 mips board at \$21,000.

Contact:

Sequent Computer Systems, Inc.
15450 SW Koll Parkway
Beaverton, Oregon 97006-5903
503-626-5700 800-854-0428 Telex 296559

Casey Powell and Scott Gibson, co-founders.
Technical: David Rodgers and Gary Fielland

Sequent Europe Limited
1 Martindale Road
Hounslow
Middlesex TW4 7EW
England
01-570-2066 Fax 01-577-5834 Telex 946114 SQNTUK

European General Manager : Stuart Bagshaw

UK Research and Academic : Ian Blagg

UK Sales Director: Peter Winder

UK Technical Support: Steve Wanless

Sales offices throughout Europe

Silicon Graphics IRIS 4D/70 High Performance Graphics Workstation

Architecture: The 4D/70 CPU is based on a RISC architecture MIPS R2000 chip running at 12.5 MHz. The CPU has a 64 Kbyte instruction cache, a 32 Kbyte write-through data cache, a write buffer, and a MIPS R2010 floating-point coprocessor. Up to 16 Mbytes of memory are accessed through a high-speed bus. The CPU can optionally have a floating-point accelerator based on Weitek parts. Communication with the graphics processors is through a VME bus.

The Graphics Subsystem operates independently from the CPU. It includes a number of proprietary VLSI processors and resides on three to five triple-high by quad-wide VME boards. Conceptually, graphics processing is performed by three sections of the Graphics Subsystem: the Geometry Subsystem, the Rendering Subsystem, and the Display Subsystem.

The Geometry Subsystem, implemented in 2.0μ NMOS VLSI technology operating at 10 MHz, includes a 16 MHz 68020 Graphics Manager with 1 Mbyte local memory for the distributed processing of graphics tasks. A pipeline of seventeen 10 MHz Geometry Engines handles object rotation, translation and scaling, six-plane clipping, perspective or orthographic viewing, and scaling to screen coordinates at over 400,000 3D coordinates/sec. The IRIS 4D/70 renders 60,000 Z-buffered, Gouraud-shaded four-sided 100-pixel polygons per second.

The Rendering processor generates pixel addresses, and performs hardware parallel interpolation of color intensities and depth values.

The Display Subsystem includes a three-domain frame buffer with image planes of 1280 x 1024 24-bit pixels, 8-bit deep window planes, optional depth planes for rapid hidden surface removal, and proprietary multi-mode graphics processors which can read the contents of the frame buffer in five parallel streams.

Configuration: The system is housed in twin towers. One contains a 12-slot card cage for the CPU, the graphics subsystem, and peripheral controllers; the other houses the power supply and up to four stacking storage peripheral modules such as 170 Mbyte hard disks or streamer tape drives.

Power requirement is 1 KW, and the system is air cooled. The monitor used is a 19" Hitachi 1280 x 1024 monitor running at 60 Hz. The 19" monitor weighs 84 lb and has dimensions 18.5" h x 20" w x 21.5" d. The dimensions of the 185 lb twin-tower chassis are 26" h x 24" w x 27" d.

Software: The operating system is an enhanced version of UNIX V.3 incorporating many features of BSD 4.3 and local enhancements to support real-time graphics.

Languages: Fortran 77 and C optimizing compilers are available. Tools include the IRIS Edge, a window-based graphical interface to DBX, enabling concurrent viewing of source code execution and results.

Applications: The major application areas are MCAE, animation, industrial design, visual

simulation, and various scientific applications such as molecular modelling and computational fluid dynamics.

Performance: Peak rate is 100 Mflops; sustained rate is 40 Mflops.

Contact:

Forest Baskett
Silicon Graphics
2011 Stierlin Rd.
Mountain View, CA 94043
415-960-1980

Gareth Jones
Windrush Court
Blacklands Way
Abingdon Business Park
Abingdon
Oxon OX14 1SY
England
0235-554444 FAX 0235-554440

Star Technologies VP Series Pipeline Floating-Point Architecture

Architecture: The VP-2 has five independent programmable processors. A separate processor is dedicated to each of the following functions: external data flow, internal data flow, and synchronization; two are dedicated to arithmetic processing. A hierarchical memory system consists of external storage devices, a large main memory, a high-speed random access partitioned data cache, and a universal register set.

The main memory comprises a 320 nsec memory, 8-way interleaved, composed of four 256K dynamic RAMs with SECDED. It is expandable to 64 Mbytes in increments of 8 Mbytes. All main memory is byte addressable (address range 4 Gbytes) and can be partitioned and protected at multiples of 16 Kbytes. Memory access time is 40 nsec (per 32-bit word). The random access data cache memory consists of 6 banks of 32K 32-bit words for a total of 768 Kbytes. During each machine cycle, four cache references are permitted: three by the arithmetic processor and one by the storage/move processor. Information flow is from host to main memory to cache to functional unit to cache to memory to host.

Other features:

80 nsec clock cycle.

2 μ CMOS.

32-bit floating-point arithmetic, pipelined functional units, both with
2 adders,

2 multipliers, and a 480 nsec divide/square root functional unit.

Ambient air cooled

Size 19" x 21" x 29"

A data interchange unit permits one of 16 operands to be selected for each arithmetic input register. During each machine cycle, three cache banks may be referenced, one loop control operation computed, four arithmetic operations started, and a conditional branch executed.

The 25 Mbyte I/O channel supports 3 device adapters; 12.5 Mbyte/sec data transfer rate.

Configuration: The VP-2 Series of array processors are designed to attach to a more general-purpose computer or host via bus.

Software:

Fortran-like control language (APCL)

Macro assembler

Simulator/debugger and Linker

Library Maintenance Program

Applications Library available.

Performance: 100 Mflops peak in single-precision (32-bit) arithmetic for convolution and matrix operations.

Status: \$95,000 base price.

Contact:

Star Technologies Inc.
515 Shaw Road
Sterling, VA 22170
703-689-4400

Technical: Phil Cannon

Stellar GS1000

Vector Register, Shared-Memory, Parallel Architecture - graphics supercomputer

Architecture: Custom-designed Application-Specific Integrated Circuits (ASICs) are used both for processor and graphics hardware. There are 11 distinct modules with approximately 2 million 1.5μ CMOS gates in 61 physical components. Silicon is fabricated by LSI Logic foundry in California, and boards are assembled by TI in Tennessee.

A central feature of this computer is its DataPath architecture which acts as a switch and multiplexor/demultiplexor rather than a conventional bus. The main processing unit, the SPMP (Synchronous-Pipeline Multiprocessor) is a custom multi-stream architecture providing up to four instruction-execution streams. Streams share functional units, on a pipelined basis, and each has its own large register files with dedicated integer, scalar floating-point, and vector floating-point registers. The clock cycle time is 50 nsec. The multi-stream processor (MSP) is implemented as a single unit thus enabling 100 nsec synchronization through concurrency registers within the MSP. The four streams of the MSP are interleaved onto a single 12-stage MSP pipeline. In the steady state, an instruction finishes on every cycle, for an instruction throughput of 20 mips. Note that each stream is completing an instruction every 200 nsec because of the length of the pipes. The use of a technique called packetization, whereby a single stream can execute 2 non-conflicting instructions simultaneously, can increase the performance to 25 mips.

In addition to the multi-stream processor, there is a special-purpose functional unit for scalar and vector floating-point instructions using Weitek 2264/2265 chips.

Minimum main memory is 16 Mbytes expandable to 128 Mbytes with a data transfer rate of 320 Mbytes/sec and a memory bandwidth for graphical operations of 640 Mbytes/sec (possible because pixel data is accessed in 128 rather than 64 byte blocks). Memory cycle time is 200 nsec. All four streams share a single 1 Mbyte static RAM cache, thus avoiding the coherency problems of multiple cache machines. The cache line is 64 bytes, and one line can be accessed every clock cycle for a transfer rate of 1.28 Gbyte/sec.

The Main Data Path also manages DMA I/O, using four I/O channels each with a capacity of 16 Mbytes/sec. Multiple controllers and disk striping are supported.

A PC-AT compatible integral Service Processor, based on a 80386 microprocessor, controls booting and console functions in addition to managing scan-path and remote diagnostic systems built into each circuit. PC-AT software can be run on the Service Processor under a window of the display.

One feature of this computer is the tight integration of general-purpose and graphical computation. Central to this is the Rendering Processor, a custom-built special-purpose SIMD engine which executes 320 million graphics-specific operations per second and implements high-level machine instructions for high-performance rendering of complex shaded and solid images, including lighting, Gouraud and Phong shading, depth-cuing, and anti-aliasing. Using virtual pixel maps, images are rendered into virtual memory which allows n-way buffering and does not restrict image size to that of the display devices. Images are trans-

ferred from main memory to the frame buffer at 640 Mbytes/sec. A 16- or 32-bit frame buffer is available which allows both hardware double-buffering and stereo viewing. An enhanced X-window system and the Programmer's Hierarchical Interactive Graphics System (PHIGS) are supported. Stellar is assisting in the development of PHIGS+ and will support these extensions to PHIGS in hardware as well as software. Main display device is a 1280 x 1024 19" color monitor running at 74 Hz.

Configuration: The GS1000 can be used as a stand-alone machine. However, TCP/IP, NFS, and Ethernet are supported, and support is planned for ISO/OSI protocols, Pronet-80, and FDDI (Fiber Distributed Data Interface) when available or defined. 80 Mbyte hard disks, a 380 Mbyte 5 1/4 inch disk drive (with double density option at 766 Mbytes), a 600 Mbyte 8" disk drive, a 120 Mbyte cartridge tape drive, and a 1/2 inch tape drive are supported. There is support for up to three VME buses, and a PC-AT compatible bus.

Software: The operating system, called Stellix, is based on Unix System V Release 3.1, with enhancements for multiprocessing, I/O, Berkeley 4.3, etc.

Languages available: The Fortran-77 (with extensions) and C compilers automatically detect parallelism and use vector processing. The Fortran compiler has most of the popular VMS extensions. An execution profiler and a multi-stream symbolic debugger are available. A concurrency library is available for manual control of program concurrency. A Stellar Assembler language is available. Ada and LISP compilers are under development.

Applications: Major applications targeted include computer-aided design and engineering, molecular modelling, computer animation, image processing, geophysical modelling, simulation and analysis, fluid dynamics, aerodynamics, astrophysics, and meteorology. Stellar has reached agreements for the porting of over 45 third-party software applications, and discussions are ongoing with over 70 application software vendors.

Performance: Peak rates of 20-25 mips and up to 40 Mflops in double precision (64-bit words). Graphics maximum rate of 600,000 3D vectors/sec and 150,000 Gouraud shaded polygons/sec. It is planned to offer over 100 mips processing by 1990.

Status: The cost of a configuration which includes the three processing units, 16 Mbytes of memory, a 380 Mbyte disk, cartridge and high-density floppy drives, PC/AT controller, a 1280 x 1024 monitor, and operating system is \$104,900. The parallelizing Fortran compiler costs \$4K.

The first shipment was to the NIH in March 1988. European shipments commenced in June 1988.

Contact:

Stellar Computer Inc
85 Wells Ave.
Newton, MA 02159
617-964-1000

Chairman and CEO: Dr. John William Poduska, Sr.
President and COO: Arthur Carr

VP Sales: Wallace E. Smith
Technical Support: Timothy Stewart
VP International Sales: Dan Murray

Ian Gilbert
UK Marketing and Sales
Little Eastwick
Lower Farm Road
Effingham, Surrey KT24 5JJ
England
0372-58707

Hans Holler
Germany Marketing and Sales
Hagenauer Strasse 42
6200 Wiesbaden
WEST GERMANY
49-61-22037

Makoto Yamada
Japan Sales and Marketing
Kihoh Bldg. 1F
2-2 Koji-machi
Chiyoda-ku
Tokyo, JAPAN
81 3 237 0131

European marketing divided into three regions centered on UK, France, and Germany.

Supertek S-1
Vector Processor

Architecture: Basic chip used: Proprietary CPU based on TTL/CMOS logic chips

Memory: Shared memory

Connectivity: four-port, 16-way interleaved

Memory sizes: up to 128 Mbytes real memory

Floating-point unit: Cray format

Configuration:

Stand-alone or networked via Ethernet or HYPERChannel

Peripherals: disks (800 MB each, 2.5 MB/sec SMD)

(680 MB each, 12.5 MB/sec PTD)

Tape: 9-track, 6250 BPI

Software: CTSS (Cray Time Sharing System)

VECTRIX (IEEE POSIX compliant, Supertek proprietary Unix)

Languages: CFT 1.13 (Cray Fortran Compiler)

Supertek Fortran Compiler

Supertek C Compiler

Fortran characteristics: Fortran-77 vectorizer with interactive debugging facilities

Applications:

Approximately 200 packages

Performance: Peak: 36 Mflops, Benchmark: 26.5 Mflops (LINPACK 300 x 300, all-Fortran)

Status: Delivery of first machine: July 1988

Cost: \$250,000 base price

Proposed markets: mainly scientific/engineering

Thinking Machines Connection Machine Model CM-2 Data Parallel Supercomputer (SIMD)

Architecture: The Connection Machine architecture assigns a processor to each element of the program's data. For example, a 4096 x 4096 array has 16 million elements of data. Hence it requires 16 million processors. Virtual processing architecture allows the systems software to subdivide physical processors into the requisite number of virtual processors. CM-2 systems have a maximum of 65,536 physical processors. A proprietary chip implements 16 physical processors. Each processor has 8K bytes of local memory, for a system-wide total of 512 Mbytes. Information is passed among processors by a very high speed (3 Gigabits/second) communications path. All physical processors may send messages in parallel. Systems may be configured with either 32-bit or 64-bit floating-point hardware.

Configuration: CM-2 systems include 4K, 8K, 16K, 32K, and 64K physical processors, and at least one front-end system (maximum of four frontends in a configuration). The front end provides program control and user interaction. VAX, Symbolics 3600, and Sun front ends are supported. Large data files are stored on the 10 Gbyte (expandable to 20 Gbyte) Data Vault. Data Vaults sustain transfer speeds above 20 Mbytes per sec. The results of computations may be displayed on a high-speed graphic display system.

Software: Systems with VAX and/or Sun front ends use the Unix environments of these systems. Systems with Symbolics front ends use the Lisp environment of these systems.

Languages: System languages are C*, *Lisp, and Fortran.

Fortran Characteristics: Connection Machine Fortran is Fortran 77 with array extensions from the ANSI 8x language proposal. Array extensions allow computations to be carried out on every element of an array at once. They include array intrinsics, which are a compact way of specifying global operations that may change the dimensionality or otherwise alter an array. Since there is a single program (that operates on all data at once), no special synchronization commands or debugging techniques are required.

Applications: Applications running currently on CM-2 include molecular dynamics, fluid flow, 3-D elastic wave simulation, document retrieval, medical imaging, finite element stress analysis, object recognition, cellular automata, VLSI simulation, and fundamental physics simulation.

Performance: Peak hardware performance is 31 Gflops. Rated performance is 2500 Mflops (the speed at which the machine multiplies two large 64-bit matrices).

Status: The first CM-2 system was delivered in September 1987. System prices range from \$1M to \$7M.

Contact:

Thinking Machine Corp.
245 First St.
Cambridge, MA 02142-1214
617-876-1111

Unisys Integrated Scientific Processor System ISP 1100/90 Vector Parallel Architecture

Architecture: Heterogeneous processing system - up to four processors (IPs), two of which can be vector processors (ISPs). A completely integrated system. LSI and MSI integration scale ECL logic is used. The ISP has independent scalar modules (SM) and a vector module (VM) with a multiply, and add, and a move pipe and 16 vector registers, each holding 64 integers, 64 single precision floating-point numbers (36-bit), or 32 double precision floating-point numbers. The clock cycle time is 30 nsec. The SM and VM share a local cache of 4 Kwords with a one cycle access time. The cache is principally used by the SM although the VM can address into the cache. Hardware gather/scatter is supported. The operating system will run on the IPs in parallel with computation on the ISPs.

There are up to 16 Mwords (36-bit words) of memory available in increments of 4 Mwords. Memory is interleaved with a bank size of .5 Mwords and a bank cycle time of 90 nsec. The peak transfer rate from memory to an ISP is 133 Mwords/sec. This main memory is termed the Scientific Processor Storage Unit (SPSU).

Configuration: A basic Integrated Scientific Processing system consists of a Unisys 1100/90 CPU with one I/O Unit, the ISP, and a 4 Mwords SPSU.

Software: OS/1100 and UNIX available. Vectorizing compiler UFTN with symbolic debugger, many 8x extensions, and fork/join parallelization primitives. A program execution evaluation routine (PEER) is available. Common Math Library includes functions like SIN and COS. Extended Math Library includes BLAS, LINPACK, EISPACK, and FFTs.

Performance: The peak performance of a single ISP is 133 Mflops in single precision (36-bit word) and 67 Mflops in double precision (72-bit word). The sustained performance is 20 to 30 Mflops in double precision and may double for single precision.

Status: First delivery was June 1986.

Contact:

Dave Deak
Unisys Corporation
Information Systems Group
P.O. Box 500
Blue Bell, PA 19424
215-542-5216