.

ARGONNE NATIONAL LABORATORY
9700 South Cass Avenue
Argonne, Illinois 60439

Advanced Architecture Computers

Jack J. Dongarra and Iain S. Duff

Mathematics and Computer Science Division

Technical Memorandum No. 57 (Revision 1)

January 20, 1988

Advanced Architecture Computers *

Jack J. Dongarra and Iain S. Duff

(dongarra@anl-mcs.arpa and na.duff@su-score.arpa)

Mathematics and Computer Science Division
Argonne National Laboratory
Argonne, Illinois 60439-4844

Computer Science and Systems Division
Building 8.9
Harwell Laboratory
Oxfordshire OX11 ORA
England

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. It is intended that this report will be continually updated, and to this end the authors welcome comments.

Keywords

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

1. Introduction

In the last few years several machines have been announced that use some form of parallelism to achieve a performance in excess of that attainable directly from the underlying technology used in the design 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. We give a list of such chips in Appendix A.

After listening to a great number of both technical and sales presentations on these new computers, we quickly became overwhelmed and confused with the characteristics of each product and its relative strengths and weaknesses. In an effort to clarify our understanding, we have written this report summarizing the principal features of each machine. We hope that the publication of this report will provide similar assistance to other computational scientists and will clarify what architectures are currently being employed and the range of machines

* Work supported in part by the Applied Mathematical Sciences subprogram of the Office of Energy Research, U. S. Department of Energy, under Contract W-31-109-Eng-38. During preparation of the original report, the second author was on leave from Harwell Laboratory.

This version was typeset on January 20, 1988.

available.

In Section 2 we list the computers considered and discuss the criteria we have used to select these computers. We present a rough classification based on architectural features and use this in our list of machines. We also summarize principal features of the machines in two tables: one for the expensive supercomputers and the other for cheaper machines. More detailed information on the machines is provided as Appendix B of this report.

The guidelines used in preparing the detailed descriptions are given in Section 3. 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 the document. We would be grateful for comments and criticisms that might help to remedy these deficiencies. We intend to update this report from time to time to reflect both the changing marketplace and further information on currently listed machines.

2. Summary and Classification of Machines Considered

In recent months there has been an unprecedented explosion in the number of 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 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, be delighted to be alerted to ongoing activities.

Some machines not commonly thought of as multiprocessors can be used as such. For example, the IBM 3081, 3084, and 3090 are multiple-processor machines. Most installations use this feature to increase the throughput, but it is possible to use them as multiple processors (with multiplicity up to 2, 3, and 4 for the two machines, respectively) using the IBM Program Product MTF which runs under MVS. We do not, however, give further details of these machines in Appendix B. In addition, we include information only on attached processors whose performance is the supercomputer range.

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 (1986). 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 on examining current machines that the 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.

With this caveat on the difficulty of classifying machines, we list below the machines considered in this report. We group those with similar architectural features. We have not included the machines from American Information Technology, Cydrome (Axiom), Data Technology Corporation, and Vitesse in this list since the documentation we have on these machines has insufficient technical details for us to classify them.

scalar

pipelined (e.g., 7600, 3090)

parallel pipelined

wide instruction words

CHoPP

FPS 164

FPS 264

Multiflow

STAR ST-100

vector

memory to memory

CDC CYBER 205

register to register

Convex C-1

CRAY-1

CRAY X-MP-1

Amdahl 500,1100,1200,1400

(Fujitsu VP-50,100,200,400)

Galaxy YH-1

Hitachi S-810

NEC SX-1E, SX-1, SX-2

Scientific Computer Systems

cache-based r-to-r

Alliant FX/1

parallel

global memory

bus connect

Alliant FX/8 (vector capability)

Culler 7

Elxsi 6400

Encore Multimax

FLEX/32

IP-1

Sequent Balance 21000

direct connect

CRAY-2 (vector capability)

CRAY-3 (vector capability)

CRAY X-MP-2/4 (vector cap.)

Denelcor HEP-1

IBM 3090/VF (vector capability)

NAS AS/91X0 (vector capability)

Sperry 1190/ISP (vector capability)

Banyon network connect

BBN Butterfly

local memory

hypercube

Ametek System 14

Connection Machine

FPS T-Series

Intel iPSC

NCUBE

ring-bus

CDC CYBERPLUS

lattice

Goodyear MPP

Active Memory Systems (DAP)

dataflow

Loral DATAFLO

user configurable

Meiko

multilevel memory

ETA-10 (vector capability)

Myrias 4000

systolic

SAXPY

high-performance graphic workstation

Dana Group

Silicon Graphics Inc

Stellar

A more empirical subdivision can be made on the basis of cost. We split the machines into two classes: those costing over $1 million and those under $1 million. The former group is usually classed as supercomputers, the latter as high-performance engines. With this subdivision, we can summarize the machines in the following tables. Since we do not have sufficient technical information on the the Galaxy YH-1, Vitesse machines, PS-2000, and MIPS, we have excluded them from these summary tables.

Table 1

Machines Costing over $1M (base system)

Machine	Word Length	Maximum Rate in MFLOPS	Memory in Mbytes	OS	Number of Proc.
Amdahl 1400 (Fujitsu VP-400)	32/64	1142	256	Own	1
CHoPP	64	?	?	Own	16
CRAY-1	64	160	32	Own	1
CRAY X-MP	64	235/proc	128	Own/UNIX	1,2,4
CRAY-2	64	488/proc	2048	UNIX	4
CRAY-3	64	1000/proc	16000	UNIX	16
CYBER 205	32/64	800(f)	128	Own	1
CYBERPLUS	32/64	100/proc	4(a)	Own	256
Denelcor HEP-1	32/64	10/PEM	16/PEM	UNIX	16(b)
ETA-10	32/64	1250/proc	2048(c)	Own	1,2,4,6,8
FPS T-Series	32/64	16/proc	16384	Own	8 - 16384
Hitachi S-810/20	32/64	840	256	Own	1
IBM 3090/VF	32/64	108/proc	256	Own	1,2,4
Myrias 4000	32/64/128	???	512/Krate	UNIX	1024/Krate
NAS AS/91X0	32/64	???	64	Own	1 or 2
NEC SX-2	32/64	1300	320(d)	Own	1
SAXPY	32	32/proc	512	Own	32
Sperry 1190/ISP	36/72	133/proc	64	Own	1,2,4 (e)

(a) Memory per processor.
(b) 64 processes possible for each PEM; however, effective parallelism per PEM is 8-10.
(c) Also 32 Mwords of local memory with each processor.
(d) Also a 2-Gbyte extended memory.
(e) Only 1 or 2 ISPs can be attached.
(f) 800 MFLOPS for 32-bit arithmetic / 400 MFLOPS for 64-bit arithmetic.

The actual price of the systems in Table 1 is very dependent 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-1 in SSI, CRAY X-MP, and 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. For the multiple-processor systems, the form of synchronization varies: event handling on the CRAYs, asynchronous variables on the HEP, send/receive on the CYBERPLUS. The CRAY-3 and ETA-10 are not yet available. Both Amdahl and Hitachi systems are IBM System 370 compatible.

In Table 2 we summarize machines in the lower price category. The data presented in Table 2 differ from that of Table 1. Full details for all the machines are given in Appendix B. Because of the widely differing architectures of the machines in Table 2 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 either as local or global. A value for the maximum speed is even less meaningful than in Table 1, since a high Megaflop rate is not necessarily the objective of the machines in Table 2, and the actual speed will be very dependent upon the algorithm and application. In the other aspects quoted in Table 1, all the machines in Table 2 are similar. All machines, except the FPSs and the SCS (all 64 bit), the DAP, MPP, and Connection (all bit-slice, supporting variable-precision floating point), the Star and SAXPY (32 bit), and Sperry with 36 and 72 bit, have both 32- and 64-bit arithmetic hardware, with most of them adhering closely to the IEEE standard.

Table 2

Machines costing under $1M

Machine	Chip	Parallelism	Connection
Active Memory (DAP)	ECL	1024	near-neighbor
Alliant FX/8	WTL 1064/1065 plus 10 gate arrays	8+vector	cross bar (reg to cache) and bus (cache to memory)
Ametek System 14	80286/80287	256	hypercube
Analogic	MC68000/VLSI	Vector	(scalar)
BBN Butterfly	68020/68881	256	Banyon network
TMI Connection	VLSI	64000	hypercube
Convex C-1	Gate array	Vector	(vector)
Culler 7	Gate array	4	bus
Cydrome (Axiom)	LSI	VLIW	(scalar)
Dana Group	Gate array	vector	(vector)
Elxsi 6400	ECL	12	bus
Encore Multimax	32032/32081	20	bus
Flex/32	32032/32081	20	bus
FPS-364	LSI	VLIW	(scalar)
FPS-264	ECL	VLIW	(scalar)
FPS-164/MAX	VLSI	16	bus
FPS-5000	VLSI	4	bus
FPS MP32	VLSI	3	bus
Intel iPSC	80286/80287	128	hypercube
IP-1	????	8	cross-bar
Loral DATAFLO	32016/32081	256	bus
Goodyear MPP	VLSI	16384	near-neighbor
Meiko	Transputer	157	user-configurable
Multiflow	Gate array	VLIW	(scalar)
NCUBE	Custom VLSI	1024	hypercube
Numerix	VLSI	Vector	(scalar)
SCS-40	ECL/LSI	Vector	(vector)
Sequent Balance 21000	32032/32081	30	bus
Silicon Graphics	Gate array	vector	(vector)
Star ST-100	VLSI	VLIW	(scalar)
Stellar	Gate array	vector	(vector)

VLIW - Very Long Instruction Word

3. 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 3, we give the template used in presenting the data in the appendices.

Table 3
Template for Description of Machines

- Name of machine, manufacturer, backers, etc.
- Contact: technical and sales
- Architecture
  - Basic chip used
  - Local, global-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)

Reference

Flynn, M. J. (1966) Very high-speed computing systems. Proc IEEE, vol. 54, pp. 1901-1909.

APPENDIX A

LIST OF BASIC CHIPS USED

General-Purpose Floating-Point Processors

Intel 8087/80287
National 32081
Motorola 68881
Zilog 8070
AMD 9511A/9512
Fairchild F9450

Building-Block Floating-Point Processors

Weitek WTL1032/1033
TRW TDC 1022/1042
Weitek WTL 1064/1065
AMD 29325
Analog Devices ADSP2310/2320

General-Purpose Building-Block Floating-Point Processors

Weitek WTL 1164/1165 (Fandrianto and Woo 1985)

Memroy, control, and communication chip

INMOS T414 transputer
INMOS T800 transputer (integral floating point)

Reference

Fandrianto, J. and Woo, B.Y. (1985), VLSI floating-point processors. IEEE Proceedings of the 7th Symposium on Computer Arithmetic, pp. 93-100.

APPENDIX B

DETAILS OF MACHINES CONSIDERED

ALLIANT FX/1 and ALLIANT FX/8

Alliant Computer Systems Corp.
42 Nagog Park
Acton, MA 01720

617-263-9110

In Europe:
Peter Smith
Sales Manager DPS9000 Products
Apollo Computer (UK) Ltd
Oriel House
26 The Quadrant
Richmond
Surrey TW9 1DL
UK
01-948-6055   Telex 8953944   Fax 01-948-5845

Contact: Technical: Craig J. Mundie, vice president of software
Contact: Sales: David L. Micciche, vice president marketing, sales
and customer servies

Backers: Venrock
Hambrecht and Quist
Kleiner, Perkins, Caulfield and Byers

Formerly, the company was called Dataflow.

Vector Register Parallel Shared Memory Architecture

Computational elements (CEs) execute applications code using vector instructions. An FX/1 has one CE. An FX/8 has 1-8 CEs. The CEs transparently execute the code of an application in parallel. CEs may be added in the field, increasing performance without recompilation or relinking.

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/8 has 1-12 IPs.

Basic chip used: Weitek 1064/1065 plus ten different gate array types with 2600 to 8000 gates. In addition, the

Motorola 68012 is used in the IP. The cycle time is 170 ns.

CEs are cross-bar connected on the backplane to a 64 Kbyte/128 Kbyte write-back computational processor (CP) cache (FX/8). Bandwidth is 376 Mbyte/sec.

Each 32-Kbyte IP cache is connected to 1-3 IPs (FX/8) or 1-2 IPs and a CE (FX/1). The FX/8 has 1-4 IP caches; the FX/1 has one IP cache.

The CP and IP caches are attached by two 72-bit busses to the main memory. Memory bus bandwidth is 188 Mbyte/sec.

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

Range of memory sizes available: 8-16 Mbytes (FX/1), 8-64 Mbytes (FX/8), all 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: Standalone. TCP/IP network support.

Size (inches): FX/1 system - 28h x 13w x 25d (the FX/1 I/O expansion cabinet is the same size); FX/8 system - 43h x 29w x 34d (the FX/8 I/O expansion cabinet is 22w and same height and depth).

Cooling: Both the FX/8 and FX/1 are air-cooled. The FX/8 system consumes 4950 watts (max. configuration), the FX/1 system 1155 watts (max. configuration).

Peripherals:

800/1600/6250 BPI start-stop tape drive
67, 134, and 379 Mbyte (formatted) Winchester disk drives
45 MBbyte 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. Compiler runs on production hardware and software.

Languages: Fortran, C, Pascal

Fortran characteristics:

F77 - Conforms to 1978 ANSI standard.
Extensions - Most of VAX/VMS extensions and Fortran 8x array extensions.
Debugging facilities - Yes.

Vectorizing/parallelizing capabilities - Automatic detection of vectors and parallelism.
Feedback to user via diagnostic messages.
User control of transformations via directives in the form of Fortran comments.
Does interprocedural dependency analysis for automatical determination of parallel subroutine calls in loops.

Performance:

Scalar 32 bit - 4.45 MIPs / CE. (4450 Kwhetstones)
Scalar 64 bit - 3.63 MIPs / CE. (3630 Kwhetstones)
Vector 32 bit: 11.8 MFLOPS / CE.
(1 chime multiply-add triad at 170ns/chime)
Vector 64 bit: 5.9 MFLOPS / CE.
(2 chime multiply-add triad at 170ns/chime)
(64-bit multiply is 2 chimes; 64-bit add, subtract, and move are 1 chime).

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.

Expected cost: FX/1 - $132,000 to $200,000; FX/8 - $270,000 to $750,000

Amdahl Vector Processors (Fujitsu VP)

John Roberts
Amdahl Corp.
1250 East Arques Ave.
P.O. Box 3470
Sunnyvale, CA 94088

408-746-6880

In Europe:
AMDAHL UK
Dr. Horst-Peter Rother
Producy Manager Amdahl Vector Processor
International Management Services Limited
Dogmersfield Park
Hartley Wintney
Hampshire RG27 8TE ENGLAND
(0252)-24555 Telex 858486

Vector Register Architecture

The Amdahl 500, 1100, 1200 and 1400 Vector Processors are marketed by Amdahl Corp. in the U.S., Canada, and Europe. 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.

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 plus 10 unique instructions. The vector performance varies according to model as follows:

Model	Peak MFLOPS
500	133
1100	267
1200	533
1400	1142

The scalar processor cycle time is 14 ns (VP 1400 only) or 15 ns (compared to the X-MP’s 9.5 ns), but a sampling of scalar instructions indicates that the VP operations may be slightly faster than the X-MP’s. There is, moreover, a difference in the pipelining between the X-MP and VP. Each VP scalar instruction is pipelined in three stages: fetch, decode, and execute. However, unlike the X-MP, the execution stage in the VP is not segmented. Thus, there is less potential purely scalar overlap in the VP than in the X-MP. (Note that all scalar work can overlap vector operations.)

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. 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 256 Mbytes (4 to 32 M 64-bit words).

Configuration	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 (ns)	15	15	15	14
Vector cycle time (ns)	7.5	7.5	7.5	7
# concurrent arith pipes	1	2	2	2
# 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)	128	128	256	256
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

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

Technology:
- 400 and 1300 gate ECL, 350-picosecond delay
- main memory - 64 Kbit, 55 ns, MOS static RAM
- 380-470 square feet
- 36-62 KVA power consumption
- air cooled

Software:
- Automatic vectorizing Fortran compiler
- Scalar Fortran compiler
- Interactive debugger
- Performance measurement tools
- Interactive vectorizer
- Scientific subroutine library (223 routines)

AMETEK System 14

Ametek Computer Research
610 North Santa Anita Avenue
Arcadia, California 91006

Technical Contact: Dr. Jeff Fier
Sales: John C. Wyckoff, IL
818-445-6811

Hypercube Architecture

This is the first generation of AMETEK Concurrent Processing Systems. Each node is based on a 80286/80287 Applications Processor/Floating Point Co-processor with a separate 80186 Communication Processor. Each node has 8 bidirectional communications channels at 3 Mbits/sec connected to the host machine through a 1 Mbyte/sec parallel interface. Effective node-to-node throughput is 100 Kbyte/channel. Software overheads per message are about 300 microseconds.

Local memory - 1 Mbyte per node.

Connectivity - 16 to 256 nodes are connected in hypercube to form a System 14.

Floating Point Unit - IEEE Standard Floating Point Arithmetic

Configuration: Front-end machines (host) are DEC VAXs (MicroVAX II through VAX 8600). Support is available for the host running either UNIX 4.2bsd or VMS. A copy of the AMETEK Operating System, XOS, runs in each node. XOS supports automatic message buffering, message forwarding, process creation, and machine partitioning for multiple users.

Language: C

Software: Consisting of a simulator, single and multi-process debuggers, and user interfaces, the AMETEK Development Environment (ADE) is specifically designed to be the most complete set of software development tools for parallel program development.

ADE allows the programmer to develop, compile, and link programs that run on the simulator and/or the hypercube. Only one copy of source code exists for debugging on the simulator and running on the hypercube. The ADE allows the user to switch between the simulator mode and the hardware mode with a single command - automatically locating the correct libraries, using the correct compilers, and generating the executables for either mode.

The simulator enables the programmer to simulate and debug parallel processes on a sequential computer. While the single process debugger allows the debugging on one task at a time, the multi-processes debugger enables the debugging of many concurrent processes. The programmer has the ability to shift on command between processes at any time.

The user interface will automatically assign the type of topology requested by the programmer. The choices consist of the nodes being defined as ring, 2-D nearest neighbor, and 3-D nearest neighbor. This enables the programmer to spend time where it is most important - writing and debugging the program.

ADE training classes have shown that the experienced sequential programmer will be running successful parallel programs in two to three days.

STATUS:

Production shipments since first quarter 1986.

ANALOGIC AP500

Analogic Corporation
Audubon Road
Wakefield, MA 01880
(617) 246-0300

In Europe:
Analogic Limited
68 High Street
Weybridge
Surrey KT13 8BN ENGLAND
(0932) 56011

Pipelined Array Processor

Control processor uses Motorola MC68000

Cycle time 160 nsec.
Pipelined adder can deliver a result each clock cycle, whereas pipelined multiplier produces a result every other cycle for a maximum rate of 9.375 Mflops.

32-bit words but arithmetic performed in 40-bit pipeline.

Program memory of 256K bytes and data memory of 912K words.

I/O:
DMA/PIO Host Interface.
RS-232 serial port with user-settable transmission rate to 19.2K baud.
Two 6.25 MHz auxiliary I/O ports (optional)
IEEE-796 standard multibus (optional)

Software includes :
Linker
Assembler
Debugger
Diagnostics
Program optimization
Function libraries

Applications can be written in :

Host high-level language
Host assembly language
AP assembly language

Arithmetic:
32-bit DEC floating-point arithmetic.
multiple-precision capabilities.

1024-point complex FFT in 4.7 msec. 100 x 100 matrix inversion in 649 msec.

size: 5.25"h x 19"w x 21"d (rack-mountable) weight: 55lbs. power: 200 Watts for basic system

BBN Butterfly Parallel Processor

Bolt, Beranek and Newman; Advanced Computer Inc.

Gary Schmidt
BBN Advanced Computers Inc.
Cambridge, MA 02238

617-497-3931

Parallel Butterfly Network Architecture

The Butterfly Parallel processor is a tightly coupled, shared memory multiprocessor housing up to 256 processor boards, each with an MC68000 microprocessor or, optionally, an MC68020 microprocessor and MC68881 floating point coprocessor. Every processor board includes either 1 or 4 megabytes 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 1- or 4- megabyte share of the global memory pool.

Other features:

Tightly coupled, shared memory, symmetrical multiprocessing.
Multiple instruction, multiple data (MIMD) architecture.
Up to 256 Mips of processing power in 1-Mip increments.
All processors have equal access to as much as 1024 megabytes
(ie, one gigabyte) of main memory.
Memory bandwidth up to 1024 megabytes/sec (one gigabyte/sec).
Memory access time less than 1 microsecond typical,
4 microseconds worst case (without contention).
Distributed I/O system supports RS-232, RS-449, Ethernet, and
Multibus.
Field expandable in single processor increments.

Each processor node is a separate circuit board with its own MC68000 (or MC68020 with MC68881 floating point coprocessor), an AMD2901 bit slice processor that extends the MC68000 instruction set, an onboard switching power supply, and either 1 or 4 megabytes of memory. Processors access their onboard "home" memory directly in less than 1 microsecond; they can access the home memory of any other processor through the Butterfly switch in about 4 microseconds. Providing true parallel access to memory, the Butterfly performs up to 256 simultaneous reads or writes and automatically resolves contention for memory.

Software includes the Chrysalis Operating System (somewhat like UNIX) with full C and Fortran support. A Lisp system is being developed. Extensions to all languages simplify parallel programming. Any of several "front end" processors, such as Sun Microsystems or VAX family computers, provide the familiar Berkeley UNIX development environment where parallel programs for the Butterfly can be written, maintained, and partially debugged.

Cost varies from $40,000 to $2,5000,000 depending on size.

CHoPP

Sullivan Computer Corporation
1012 Prospect Street
Suite 300
La Jolla, California 92037
(619) 454-3116

Lee Higbee, VP Research

VLIW (Very Long Instruction Word) Architecture

The computer is under development by Sullivan Computer Corporation. The single processor is claimed to be several times faster than current supercomputers and will not require special coding techniques such as those required for vector processors, hypercubes, or other highly-parallel systems. The machine under current development is the Demonstration Unit (DU), a single processor version of the CHoPP 1. The CHoPP 1 will include up to 16 parallel processors. The features listed below highlight some features of the DU.

A superinstruction that includes up to 9 instructions is executed each clock cycle, providing one of the highest instruction issue rates available today.
Four address arithmetic and logic units (ALUs) and four computational functional units, each an ALU and floating point unit, support the 8 concurrent computations in each superinstruction.
A zero delay branch is the ninth executable instruction. The central processing unit has multiple register sets to support many tasks in concurrent execution (multiprogramming).
The memory bandwidth is approximately 200 MWDS/sec or 1600 MB/sec and the I/O bandwidth is approximately 16 MWDS/sec or 130 MB/sec.
On the Livermore Loops, the DU is expected to perform over two times the CRAY X-MP. Delivered performance to price ratio is expected to be over 4 times the CRAY X-MP/12.

The machine is small and air cooled; it is compatible with most computer environments; it does NOT require special cooling systems. Much of the lowest level of the Operating System is in hardware, providing much lower O/S overhead.
Optimizing compilers are easy because there is no need for special techniques such as automatic vectorization or parallelization. This implies that it will be easy for Sullivan to support many languages with very high quality code from their compilers. Porting will be easy.
The Fortran will accept the common extensions, both those that extend Fortran's functionality and those that allow for improved optimization of the compiler's output.

Plans
The CHoPP 1, which is essentially a multiprocessor version of the DU described above, allows from four to 16 processors (and will be about four to 16 times as fast) because of their (patented and proprietary) conflict-free, crashless memory and memory interconnect design.

The CHoPP 2, which is the CHoPP 1 with very high speed circuitry (ECL), is expected to allow from four to 32 CPUs, each running at about five times the clock speed of the CHoPP 1. The CHoPP 2 is projected to provide ten times the performance of the CHoPP 1.

Connection Machine

Thinking Machines Inc.
245 First St.
Cambridge, Mass. 02142-1214

617 876-1111

James Bailey - Director of Marketing

Parallel Hypercube Architecture

The Connection Machine is a very fine grain parallel computer with an architecture suitable for artificial intelligence applications. The 64000 node processor prototype will have 1000 times the logical inference performance capabilities of current LISP workstations.

The processing elements are one-bit machines having 4096 bits of memory connected so that each processor can communicate with any other through a fast message-routing system that forms a hypercube network. All linkages are software controlled with system-wide message flow being handled by a 3 Gigabit per second message routing system. All memory is dual ported and is hence directly accessible by both the Connection Machine and the front end.

Configuration: The Connection Machine system has 65536 physical processors but may be configured for a much larger number of logical processors by means of the global-reset and configure commands.

Access is through a front-end processor, currently either a VAX or a Symbolics 3600. The front-end provides the operating system environment, including terminal interaction and file management.

The clock rate may range up to 10 MHz, giving an expected performance of 2 billion 32-bit integer additions per second in the 64K (65536) node configuration. Average instruction mixes are expected to exceed 1000 Mips.

I/O can be through the front end or direct to a 1.2 Gigabyte disk at the rate of 500 Megabits per second.

Languages: Applications programs reside in the host and can be written in CM-C (a Connection Machine extension of C), CM-Lisp, or an assembly language REL-2.

Applications: One of the principal applications is expected to be image processing. Other applications include VLSI simulation and FFT’s.

The prototype currently available uses a conservative VLSI technology of 10000 gate CMOS gate arrays.

Convex C-1 (XL and XP).

Convex Computer Corporation
701 N. Plano Rd.
Richardson, Texas 75081

Phone: 214-952-0200

Technical: Steve Wallach
Sales: Bob Shaw

In Europe:
CONVEX Computer Limited
Hays Wharf
Millmead
Guildford GU2 5BE
England
0483-69000 Telex 858136 Fax 0483-36775

Vector Register Architecture

The machine is based on CMOS VLSI gate array 8000 gates/chip (24 different chips in the machine). The c1-XP also uses two 20000 gates/chip CMOS VLSI gate arrays. It uses vector architecture, register to register, with pipelined functional units (each of which operates asynchronously - 3 present). The machine is based on a 100-ns major cycle time, 50-ns minor cycle time, with virtual memory (page size 4096 bytes) and 1024 bytes logical cache between memory and registers. Also a 64-Kbyte, 50-ns access. physical cache.

Vector operations bypass the cache (cache bypass). Scalar operands are encached.

Physical memory - up to 1024 MB (1 billion) dynamic RAM(32-way interleaved).
Virtual address space - 4 Gbytes
User address space - 2 Gbytes.
Memory - on a 32-Mbyte board(256 kbit Dram), 128 -Mbyte board(1 Mbit
Dram), 2 banks per board, each 4-way interleaved.
Transfer rates between memory and CPU - rated at 80 Mbytes/sec.
Single memory pipe between memory and registers.
Note: 64-bit vector references that are aligned on 32-bit boundaries will bypass the cache.
Vector registers - 8, each with 128 elements (64-bit elements).

VL and VS registers
0.512 Mbyte IOP buffer. IOP 68K based with event-driven monitor
I/O transfer rates of 80 Mbyte/sec

Floating point IEEE Standard format.
5 independent I/O processors each rated at 80 Mbyte/sec.
Concurrent operation of scalar and vector units (fixed and float).
Mask/merge and compress operations supported.
Reduction operators max,min,sum,prod,any,all, and parity supported.
Degradation for indirect addressing not specified.

```
A(i) = B(C(i)) ...
        LD VL
        LD C,V0
        SHF 4,V0,V1
        LD B,V0,V1
        STORE A
```

Byte-addressable with integer*1, integer *2, *4, and *8 arithmetic supported.
Also, real *4 and real*8, Logical *1, *2,*4, and *8, and complex *8
and complex *16.

Configuration: Designed as a stand-alone multiuser machine.

Software: UNIX 4.2 bsd operating system.

Languages: Fortran 77 and C (accepts VMS Fortran), with excellent Fortran vectorizing compiler. Fortran compiler accepts VAX VMS Fortran. C compiler (VC) automatically vectorizes scalar code.

Performance: Peak performance 20 MFLOPS in double precision (64-bit arithmetic), 40 MFLOPS in single precision (32-bit arithmetic). LINPACK timings - expect around 3-4 MFLOPS.

Note: Convex rates their machine as 1/6 of a CRAY 1-S, 600 ns per subroutine call, 9 cycles latency (cf. 11 on CRAY, 30 on FACOM VP)

Basic system: two 19-in. racks and 16-Mbytes memory, 1 I/O processor, service processor, 414 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
Cost: XL base system $350,000, XP base system $500,000

16 Mbytes	414-Mbyte disk	one IOP [16 lines]	$495,000
32 Mytes	828-Mbyte disk	one IOP "	$545,000
64 Mbytes	828-Mbyte disk	two IOP [32 lines]	$745,000
128 Mbytes	3312-Mbyte disk	two IOP "	$1,400,000

3312 Mbytes = 8 Fuji eagles
can have 3 asynchronous 16 line ports

F77 compiler $24.5K
VC compiler $24.5K (PCC comes with OS)
(has GPROFF, PROF, and BPROF run-time profilers)
Networking package $15K

CRAY-1

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

In Europe:
CRAY UK
Malcolm Hammerton
Cray Research (UK) Ltd
Cray House
London Road
Bracknell
Berkshire RG12 2SY ENGLAND
(0344) 485971 Telex 848841

Vector Register Architecture

This machine is no longer being produced, although when first introduced in 1976 (Los Alamos), it was undisputedly 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 second-hand CRAY-1s at knockdown prices.

Features:

A uni-processor.
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 M words 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 1-M) and the memory access time (latency) is 12 clocks (150 nsec).
No virtual memory
Register-to-register machine
8 registers of length 64 (64-bit) words each
Word addressable (64-bits).

No half precision.

Double precision is through software and is extremely slow (factors of about 50 times single precision 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.

Performance:

Low vector startup 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.

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.

CRAY-2

Cray Research Inc.
1440 Northland Drive
Mendota Heights, MN 55120

Phone : 612-452-6650

1100 Lowater Rd.
Cray Research Inc.
Chippewa Falls, Wisconsin 54701

Phone : 715-726-1211

In Europe:
CRAY UK
Malcolm Hammerton
Cray Research (UK) Ltd
Cray House
London Road
Bracknell
Berkshire RG12 2SY ENGLAND
(0344) 485971   Telex 848841

Vector Register Parallel Shared Memory 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 is compatible with the X-MP.

The system has a 4.1-nsec clock cycle time.

Memory is 256 M words 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

The machine is liquid cooled using inert fluorocarbon.

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

Performance: Max. quoted at 500 MFLOPS per processor.

Cost: $15M - $20M

Delivered: NMFECC, NASA Ames, University of Minnesota, Stuttgart, Ecole Polytechnique (Paris). Orders placed by AERE Harwell.

CRAY-3

Cray Research Inc.
1440 Northland Drive
Mendota Heights, MN 55120

612-452-6650

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

In Europe:
CRAY UK
Malcolm Hammerton
Cray Research (UK) Ltd
Cray House
London Road
Bracknell
Berkshire RG12 2SY ENGLAND
(0344) 485971 Telex 848841

Vector Parallel Architecture

The machine is essentially a GaAs version of the CRAY-2 being developed by a team under Seymour Cray at Chippewa Falls.

Architecture:

16 processors
2-nsec cycle time
4 logical functions/clock period
Memory twice as fast as CRAY-2.
Speed about 8 times CRAY-2.

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

Boards reduced from the 4 x 8 x 1 of the CRAY-2 to 1 x 1 x .1.
Only 1 cu ft in size, with power dissipation of 180 kW as in CRAY-2.
Power supplies take 10 cu ft and liquid coolant 100 cu ft.

Status: 1988 production version; 1990 sales

CRAY X-MP

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

Steve Chen
Chris Hsiung
1100 Lowater Rd.
Cray Research Inc.
Chippewa Falls, Wisconsin 54701
715-726-1211

In Europe:
CRAY UK
Malcolm Hammerton
Cray Research (UK) Ltd
Cray House
London Road
Bracknell
Berkshire RG12 2SY ENGLAND
(0344) 485971 Telex 848841

Vector Register Parallel Shared Memory Architecture

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

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 can be chained to 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 ns and a memory access time of 17 clocks.)
X-MP-4 ECL. (Bank busy time on the ECL machine is 34 ns and a memory access time of 14 clocks.)
ECL logic with .35-.5 ns gate delay and 16 gates/chip.
Main memory - ECL 4K RAMs with 25-ns 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 M word of secondary MOS memory. Data transfer to high speed (1200 Mbyte) DD-49 disk takes 10 Mbytes/sec.

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

Performance: Max. per processor is 235 MFLOPS.

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

Culler 7

Culler Scientific Systems Corporation
100 Burns Place
Santa Barbara, CA 93117
805-683-5631

Ward Davidson
Vice President, Sales and Support

Parallel Array Processor

Up to four processors. Each processor is a proprietary 64-bit high-performance computational processor.

Global data memory of 96 MBytes real memory of 120 nsec access time.

Local memory consists of program memory up to 256 KB and array memory of 4 x 16 KB with 40 nsec access time.

Each processor rated at 18 MIP and around 11 MFLOPS.

Software is an enhanced version of 4.2 BSD UNIX.

Fortran and C

The Fortran and C compilers generate instructions in parallel streams which employ all the computational function units to achieve execution concurrency within a processor.

Cost: $275K - $750K

Culler PSC

Connects to a front end workstation like a Sun.
Designed as a network compute server product architecture and performance similar to a single processor Culler 7 unit.

Cost: $98.5K (order quantity one, discounts for OEM’s).

CDC CYBER 205

ETA Systems, Incorporated
1450 Energy Park Drive
St. Paul, MN 55108

612/642-3400

Charles D. Swanson - Account Support

In Europe:
CDC and ETA UK
D. Swanston
Control Data Limited
Genesis Centre
Garrett Field
Birchwood Science Park
Birchwood
Warrington
Cheshire WA3 7BH ENGLAND
(0925) 824757 Telex 629900

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-nanosecond 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 nanosecond 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 M bits/second for each port

Maximum I/O port bandwidth of 3200 M bits/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.

CYBERPLUS

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

Martin Ferrante

800-828-8001 ext 88

In Europe:
CDC and ETA UK
D Swanston
Control Data Limited
Genesis Centre
Garrett Field
Birchwood Science Park
Birchwood
Warrington
Cheshire WA3 7BH ENGLAND
(0925) 824757 Telex 629900

Ring Bus Architecture

This is a multiple parallel processor system. It grew from the Flexible 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 Megabits/second transfer rate with a read and a write possible between processors at this sustained rate.

There are two CYBERPLUS processor models: 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 which can be expanded 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. There are 15 independent functional units:

- 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 (each connected 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 nanosecond 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-memory transfers with the HPM memory on the CYBERPLUS processors. There are two rings connecting 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 bi-polar 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 words.

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}.

Water cooled

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. Change detection algorithm for image processing is about 100 times faster than on a CDC 7600.

Software: Floating point hardware and software delivered in first quarter 1985. Fortran compiler available for research activities fourth quarter 1984 and released April 1985.

Cost: Entry-level CYBERPLUS base processor is priced at $735,000, which includes a 16-bit integer unit and 2.048 Mbytes of memory. With all available options the price is $1.6 million.

Status: Announced formally on October 4, 1983; deliveries started in the first quarter of 1985.

Cydrome (formally AXIOM Systems)

1589 Centre Pointe
Milpitas, California 95035

Richard Lipes
Bob Rau
408-943-9460

Ross Towle (compiler person, student of Kuck)
Bob Rau (Architect from University of Illinois and Elxsi)

Dataflow Architecture

Dana Group

Ben Wegbreit
Dana Group
550 Del Ray
Summyvale, CA 94086
408-732-0400

Very High Performance Integrated Graphics Workstation
Very High Performance Integvrated Graphics /workstation

Company founded by Allen Michels (from Convergent Tech)

Vector register architecture

Heavy emphasis on interactive graphics for large computational problems.

48 MFLOPS peak performance

UNIX Fortran C

Availability: 1987

Markets: CAD/CAM/CAE Molecular Modeling Image Processing Scientific Engineering Research and Development

Cost: $50 - 75K

DAP-3

Bruce Apler
Active Memory Technology Inc.
6600 Peachtree Dunwoody Road
300 Embassy Row
Suite 670
Atlanta, GA 30328
404-399-5633

In Europe:
S. MacQueen/I. Merry
International Computers Ltd
ICL Defence Systems
Lovelace Road
Bracknell
Berkshire RG12 4SN
England
0344-24842 Telex 22971

Professor Dennis Parkinson
DAP Support Unit
Computer Centre
Queen Mary College
Mile End Road
London E1 4NS
01-980-4811

Active Memory Technology Limited
Eggington House
25-28 Buckingham Gate
London SW1E 6LD
England
01-630-9811 Telex 296923 (ADVENT G) Fax 01-828-4919

Bit Parallel Architecture

Configuration: This is an SIMD lockstep machine which operates on multiple data one bit at a time. It has variable-length arithmetic. Configuration is as a grid of processing elements with nearest neighbor connections. There are also row and column data highways (not present on the ILLIAC IV) so that broadcasts can be used to sum efficiently the entries of an array or to find the maximum entry, for example. The other main advantage over the ILLIAC IV lies in the far greater memory for each processing element and the greater reliability of the components.

Three versions of the machine have been produced to date. The first, the prototype 32 x 32 machine, was followed by a larger 64 x 64 version which had an ICL 2900 host. The DAP was configured as one of the host's store modules. This resulted in no communication costs between the two machines when a common data to memory mapping format was used. The standard machine had 2 Megabytes of store, but the QMC (Queen Mary College) machine was later upgraded to 8 Megabytes (i.e., it can be visualized as a cube of dimensions 64 x 64 x 2048 bytes). Six of these machines are in use.

The third version of the machine, the one currently marketed, has returned to the 32 x 32 array size, and has 8 Megabytes of array storage. The machine is approximately two orders of magnitude smaller, (it now fits under a desk) and can run without a host. The only architectural change has been the provision of a 40 Megabyte/sec I/O subsystem to permit real time processing. The instruction cycle time has also been reduced from 200 to 150 nsec.

Software: 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 Megabyte/sec parallel interface.

Language: The principal programming language used is DAP Fortran, an augmented Fortran that includes most of the array features proposed for Fortran 8X.

Applications: Some of its main applications are in lattice gauge theory and molecular dynamics. It is particularly powerful on the Ising model because of its bit arithmetic. It is also used in many Monte-Carlo calculations and in image processing where the major problem is in data movement rather than processing speed. For some specialized applications, the DAP will outperform a CRAY-1. The new mini DAP has also been used to implement a high-performance military radar system. the Micro-VAX II and development software.

Basic System Configuration:

32 x 32 p;rocessor array
8 MBytes of array memory
1 MByte of MCU code memory
10 MHZ instruction rate
Micro-Vax II host
Single Caninet, approx: 17 x 13 x 20 inches

Cost: The DAP-3 is currently priced at around $150,000, including the Micro-Vax and development software.

Status: Work has already begun on a new machine that will use VLSI to achieve further improvements in integration levels and heat dissipation, with a dramatically improved arithmetic performance.

Elxsi System 6400

Len Shar
Elxsi
2334 Lundy Place
San Jose, CA 95131

408-942-1111

Harvey Goldman - Marketing
Len Shar - Research

Parallel Processor/Bus Architecture

This machine uses ECL technology high-density LSI components. The system can be used as a multiprocessor for multitasking of a single Fortran program, or as a loosely coupled architecture with no parallel processing capability executing independent programs or processes, or both ways.

The system can be configured with 768 Mbytes of memory and many disk drives (474 Mbytes each). Up to 12 processors can be configured with this machine, with up to 64 Kbytes of cache on each processor.

Global memory architecture is via a fast bus. The bus is 64-bit wide channel providing a gross bandwidth of 320 Mbytes per second, giving a transfer rate 160-213 Mbytes/second. All major components are connected to the bus. Up to 768 Mbytes of MOS memory are available (4 Gbytes virtual).

Other features:

Each CPU 3 boards, rated at 6 MIPS on M6410 CPU and at 10 MIPS on M6420 CPU.
64-bit wide data paths.
50-nsec cycle time.
64-Kbyte, 2-way set associative cache (100-nsec access time).
16 sets of 64-bit general-purpose registers.
IEEE floating point arithmetic.

Software: The operating system, called EMBOS, is a message-based OS. There is also Elxsi's version of UNIX, a port of AT&T System V.2 and 4.2 BSD.

Size: The 5-CPU system fits in a single cabinet, 32 in. deep by 59 in. wide.
Languages: Fortran 77, Pascal, COBOL 74, C, MAINSAIL

Cost: A single-processor system is in the range of $400,000.

A new model the 6420 CPU outperforms the old 6410 by a factor or 1.5 to 2 times. The new CPU can exist with the old CPUs.

Encore Multimax

Encore Computer Corp
257 Cedar Hill St
Marlboro, Mass. 01752

617-460-0500

Julius Marcus - VP of Marketing

Parallel/Bus Multiprocessor Architecture

Architecture:

National Semiconductor 32032 chip set running at 10 MHz.
32-Kbyte write-through cache per processor pair.
Processors connected via a fast, 64-bit wide bus
with data throughput rate of 100 Mbytes/sec.
Address space of 4 Gbytes
Main memory 32 Mbytes of RAM in 4 independent banks,
in increments of 4 Mbytes.

Configuration:

Terminal and unit record I/O connected via Annex 16 line terminal
concentrators attached to Ethernet, providing pre-processing.
Is compatible with 19-in. Encore workstation.

Note: The company plans successor chips using best microprocessors,
including RISC architectures.

20 processors maximum configuration.
Ethernet communications using TCP/IP.

Performance: Range quoted from 1.5 MIPS to 15 MIPS by adding processors per module.

Languages: UNIX 4.2 with C, Fortran, and Pascal.

Status: November 1985 with a product

ETA-10

ETA Systems, Incorporated
1450 Energy Park Drive
St. Paul, MN 55108

612/642-3400

Charles D. Swanson - Account Support

In Europe:
D. Swanston
Control Data Limited
Genesis Centre Garrett Field
Birchwood Science Park
Warrington
Cheshire WA3 7BH ENGLAND
0925-824757   Telex 629900

Vector Parallel Architecture

The ETA-10 is a successor to the CYBER 205, designed to operate at 10 GFLOPS by the end of 1986.

Architecture:

Central Processors

Multiprocessor system with 2, 4, 6, or 8 CPU's
(a one CPU system will also be available)
Very high density CMOS circuitry (20,000 gates/chip)
Liquid nitrogen cooling for performance and reliability
CYBER 205 instruction compatibility
Each CPU with a scalar and vector processor, and 4 million words
of local memory

Scalar unit
Independent, segmented functional units
256-word high-speed register file
64-word instruction stack
Vector unit
2 vector pipelines

Memory

Up to 32 million words of CPU memory (4Mw/CPU)
MOS semiconductor Shared Memory using 256K VLSI chips
Shared Memory sizes: 32, 64, 128, 192, or 256 million words
1 million word communication buffer for interprocessor communication
Virtual memory addressing
SECDED on each 32-bit half word
48-bit address (address space of 4 trillion words/user)

I/O

Up to 18 400-Mbit/sec Input/Output units for accessing disks, tapes, front-end systems and networks

Miscellaneous

Very low power requirement: 700 Watts/CPU (i.e., about 200 Watts per 205 equivalent)
Liquid nitrogen cooling
Compact packaging
High reliability: 100 per cent functional availability

Software:

Virtual operating system
Kernel operating system for basic processes
User environments for control languages and utilities:
    VSOS (CYBER 205 OS - provides CYBER 205 software compatibility)
    UNIX

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

Languages:

Fortran
    ANSI 77 with vector extensions
    32-bit half-precision data type

Special calls to machine instructions
Support for anticipated FORTRAN 8X array notation
Automatic vectorization
Scalar optimization
Multiprocessing library
Pascal
C

Performance: Too early to say. The performance of the product line is claimed to range from 2 to 4 times faster than the CYBER 205 for a single processor entry level system to 40 times faster at the high end (8 processors). The vector unit has been designed to reduce start-up times ($n_{1/2}$) relative to the CYBER 205; however, performance will still be degraded for noncontiguous vectors.

Status: Complete system checkout by early 1987, with initial beta site deliveries in December 1986. The fully configured high-performance machines shipping by third quarter of 1987.

FLEX/32 MultiComputer

Flexible Computer System

Larry Samartin
Flexible Computer Corporation
1801 Royal Lane
Bldg 8
Dallas, TX 75229

214-869-1234

President/Chairman Larry B. Samartin
President/CEO Dr. M. Nicholas Matelan

William T. Walker
National Manager
Flexible Computer Corporation
5 Great Valley Parkway
Suite 226
Malvern, PA 19355
215-648-3916

Parallel Bus 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 PC 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 SKY Warrior.

Other features:
  Standard VME bus open architecture supporting Eurocard standard.
  Communication rates on local 10 buses 160 Mbit/sec each.
  Communicatoin rates on common bus 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 messaging to another processor’s memory 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.

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 signaling, 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

Base system:

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

Cost: Price starts at approximately $100,000

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

Floating Point Systems MP32 SERIES   MODEL 3000

MP32 Series, Model 3000,
Floating Point Systems, Inc.

Steve Cannon
3601 SW Murray Blvd,
Beaverton, OR,
503-641-3151 x1883

In Europe:
David A. Tanqueray
Floating Point Systems U.K. Limited
Apex House
London Road
Bracknell
Berks RG12 2TE
England

Architecture: MIMD

Basic chip used M68000 (Control Processor), AMD & Weitek Chips (arithmetic processor)
Local, global-shared memory, or both: Both

Connectively (for example, grid, hypercube): Bus

Range of memory sizes available, virtual memory: 1Mword to 7Mword (32-bit)

Floating point unit (IEEE standard?): IEEE standard 32-bit

Configuration
Stand-alone or range of front-ends Front ends: DG MV Series, Perkin-Elmer, Microvax II, VAX

Peripherals: I/O processors

Software: Unix or other? Other

Language available: MAX 68 control language, XPAL assembler

FORTRAN characteristics: N/A
F77
Extensions
Debugging facilities
Vectorizing/parallelizing capabilities: Horizontal microcode synthesis that allows up to 10 operations to execute simultaneously.

Applications:
Run on prototype: Yes, or on front-end simulator
Software available: Math Libraries: Basic math, Signal, Image, & Geophysical

Performance:
Peak: 18 to 54 MFLOPS
Benchmarks on codes and kernels: 2D CFFT 1024 x 1024 pts - 1.89 sec.

Status:
Date of delivery of first machine, beta sites, etc.: Available since 8/85
Expected cost (cost range): $57,500 to $125,000
Proposed market (numbers and class of users): Signal processing, Image processing, and Computational physics

Floating Point Systems
FPS-5000 SERIES

FPS-5000,
Floating Point Systems Inc.

Steve Cannon,
3601 SW Murray Blvd.,
Beaverton, OR,
503-641-3151, x1883

In Europe:
David A. Tanqueray
Floating Point Systems U.K. Limited
Apex House
London Road
Bracknell
Berks RG12 2TE ENGLAND

Architecture: MIMD

Basic chip used: AMD Chips, Weitek Chips on coprocessor
Local, global-shared memory, or both: Both
Connectively (for example, grid, hypercube): Bus
Range of memory sizes available, virtual memory: 256K to 1024K (38-bit words)
Floating point unit (IEEE standard?): 32-bit IEEE (coprocessor)

Configuration:
Stand-alone or range of front-ends: Front ends: VAX; PDP-11; Perkin-Elmer 3200; Gould 32; IBM 4300, 3080, 3090; Prime 750, 9950; Harris 800, HP 1000E
Peripherals: 300MB and 80MB, Disks, I/O processors

Software: UNIX or other? Other

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

FORTRAN characteristics:
F77 (CPFORTRAN, which is F77 less I/O and character data type support)
Extensions: Calls to coprocessor programs

Debugging facilities: Symbolic debugger
Vectorizing/parallelizing capabilities: Horizontal microcode synthesis that allows up to 10 operators to execute simultaneously

Applications:
Run on prototype: Yes, or run on simulator on front end
Software available: Math Libraries: Basic & advanced math signal and image processing, simulation and geophysical

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.: Oct. 1983
Expected cost (cost range): $45,000 to $99,000 for 256Kword system + standard software
Proposed market (numbers and class of users): 350+ units per year in signal processing, image processing, geophysical analysis, computational physics, and real-time simulation

Floating Point Systems
FPS-164/MAX

FPS-164/MAX,
Floating Point Systems Inc.

Dave Vickers (Technical),
Mike Saunders (Sales)
3601 SW Murray Blvd.,
Beaverton, OR,
503-641-3151

In Europe:
David A Tanqueray
Floating Point Systems U.K. Limited
Apex House
London Road
Bracknell
Berks RG12 2TE ENGLAND

Pipeline Scalar Processor with Atttached Processor

Architecture:
  Basic chip used: Proprietary (CPU), Weitek Chips (MAX)
  Local, global-shared memory, or both: Both
  Connectively (for example, grid, hypercube): Bus
  Range of memory sizes available, virtual memory: .5Mwords to 15Mwords (64-bit words) or 4Mbytes to 120Mbytes
  Floating point unit (IEEE standard?): IEEE Standard compatibility

Configuration:
  Stand-alone or range of front-ends: Front-end connection to IBM 4300, 308x, 303x, 309x under MVS, MVS/XA, VM/CMS; DEC VAX under VMS; Sperry 1100 Series; Apollo Domain
  Peripherals: FD64 Disk subsystem (1-6 controllers, 4-24 drives), 680MB to 16.2GB

Software: UNIX or other? System Job Executive

Language available: FORTRAN, ASSEMBLY

FORTRAN characteristics:

F77 ANSI '77 optimizing compiler, 5 levels of optimization

Extensions: DOE Extensions for asynchronous I/O

Debugging facilities: Symbolic debugger

Vectorizing/parallelizing capabilities: Takes advantage of architecture through horizontal micro-coding allowing 10 different operations to occur in 8 separate functional units per machine/cycle. The matrix algebra accelerator (MAX) modules allow up to 15 concurrent vector operations at any one time.

Applications:

Run on prototype.

Software available: Math Library routines (500+), Fast Matrix Solution Library (FMSLIB) over 40 third party software packages available.

Performance:

Peak: 33-341 MFLOPS

Benchmarks on codes and kernels: 1000 x 1000 Matrix multiply - 66 seconds with 1 MAX module; - 10 seconds with 15 MAX modules

Status:

Date of delivery of first machine, beta sites, etc.: Available since 4/1/85

Expected cost (cost range): $435,000 to $1,900,000

Proposed market (numbers and class of users): Computational Chemistry/Physics, Electronic Circuit Design, Oil Reservoir Simulation, Structural Analysis

Floating Point Systems

FPS-264

FPS-264,
Floating Point Systems Inc.

Dave Vickers (Technical),
Mike Saunders (Sales),
3601 SW Murray Blvd.,
Beaverton, OR,
503-641-3151

In Europe:
David A. Tanqueray
Floating Point Systems U.K. Limited
Apex House
London Road
Bracknell
Berks RG12 2TE ENGLAND

Pipelined Scalar Architecture
Basic chip used: Proprietary ECL implementation
Local, global-shared memory, or both: Both
Connectively (for example, grid, hypercube): Bus
Range of memory sizes available, virtual memory: .5MW to 4.5MW (64-bit words), or 4Mbytes to 36Mbytes
Floating point unit (IEEE standard?): IEEE standard compatibility

Configuration:
Stand-alone or range of front-ends: Front-end connection to IBM 4300, 308x, 303x, 309x under VMS, MVS/XA, VM/CMS; DEC VAX under VMS; Sperry 1100 Series; Apollo Domain
Peripherals: FD64 Disk subsystem (1-6 controllers, 4-24 drives), 680MB to 16.2GB

Software: UNIX or other? System Job Executive

Language available: FORTRAN, ASSEMBLY

FORTRAN characteristics:

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

Applications:
Run on prototype:
Software available: Math Library routines (500+), Fast Matrix Solution Library (FMSLIB) over 40 third party software packages available.

Performance:
Peak: 38 MFLOPS
Benchmarks on codes and kernels: 1000 x 1000 Matrix multiply 53 seconds
Proposed market (numbers and class of users): Computational Chemistry/Physics, Electronic Circuit Design, Oil Reservoir Simulation, Structural Analysis

Expected cost: $640,000 to $1,350,000

Status:
Date of delivery of first machine, beta sites, etc.: Available since July

Floating Point Systems
FPS-364

FPS-364,
Floating Point Systems Inc.

Dave Vickers (Technical)
Mike Saunders (Sales)
3601 SW Murray Blvd.,
Beaverton, OR,
503-641-3151

In Europe:
David A. Tanqueray
Floating Point Systems U.K. Limited
Apex House
London Road
Bracknell
Berks RG12 2TE ENGLAND

Scalar Pipelined Architecture

Basic chip used: Proprietary ECL implementation
Local, global-shared memory, or both: Both
Connectively (for example, grid, hypercube): Bus
Range of memory sizes available, virtual memory: .5MW to 9MW (64-bit words) or 4Mbytes to 72Mbytes
Floating point unit (IEEE standard?): IEEE Standard compatibility

Configuration:
Stand-alone or range of front-ends: Front end connection to IBM 4300, 308x, 303x, 309x under MVS, MVS/XA, VM/CMS; DEC VAX under VMS, Sperry 1100 Series; Apollo Domain
Peripherals: FD64 (same as MAX except capacity) 1-2 controllers, 1-8 disks, 680 MB to 5.44 Gbytes

Software: System Job Executive

Language available: FORTRAN, ASSEMBLY

FORTRAN characteristics:
F77 ANSI '77 optimizing compiler, 5 levels of optimization

Extensions: DOE Extensions for asynchronous I/O
Debugging facilities: Symbolic debugger
Vectorizing/parallelizing capabilities: Takes advantage of architecture through horizontal micro-coding allowing 10 different operations to occur in 8 separate functional units per machine/cycle.

Applications:
Run on prototype:

Software available: Math Library routines (500+), Fast Matrix Solution Library (FMSLIB) over 40 third-party software packages available.

Performance:
Peak: 11 MFLOPS
Benchmarks on codes and kernels: 1000 x 1000 matrix multiply - 189 seconds

Proposed market (numbers and class of users): Computational Chemistry/Physics, Electronic Circuit Design, Oil Reservoir Simulation, Structural Analysis

Status:
Date of delivery of first machine, beta sites, etc.: Available since Sept. 1, 1985.
Expected cost (cost range): $298,000 to $950,000

Floating Point Systems

FPS T Series

Floating Point Systems
Beaverton, OR 97005
1-800-547-1445

Hypercube architecture - Vector processors

Each node is Inmos transputer, memory, plus vector processor.

Vector processor:
The vector processor consists of 2/3 of the surface of the board, and is
a proprietary state machine with its own instruction stream and microcode.
Three of the chips are currently Weitek parts.
6 stage 8 MFLOPS adder and a 7 stage 8 MFLOPS multiplier.
Peak performance is 16 MFLOPS for 64-bit operands
and 24 MFLOPS with 32-bit operands.
IEEE arithmetic.
192 MBytes/sec to/from memory.

Inmos transputer:
32-bit CMOS processor
7.5 MIPS processor
2 KB of one chip RAM with one-cycle access that serves like a large register set.

19MB/sec between local memory and transputer.
Local memory is 1MB of dual ported RAM.

Aggregate external bandwidth for a single node 8 MB/sec. 4 input and 4 output channels may be active simultaneously.
Maximum number of nodes that can be connected is 2^{14} (16384).
Maximum execution rate of 262 GFLOPS for 64-bit operands.

Eight nodes make up a module.
Two modules make up a cabinet.
Maximum of 1024 cabinets.

I/O peak transfer rate 80 MB/sec for a 16-node cabinet system.

Stand alone system.

A cabinet contains two system disks the user may reference through a system manager network.
Direct disks, up to 1 GByte/node, are planned for July 1987.

Software:
Occam is the language used on the Transputer.
Occam is enhanced with a library of mathematical subroutines.
Sequential languages C, Fortran, and Pascal can run on each node,
but Occam is still needed to manage concurrancy.

Each cabinet is air cooled, requires 1000 watts of power
and has a footprint of 5 sq. ft.

Delivered:
Cornell University, one cabinet 2nd quarter 86
Northrop, one cabinet
Michigan State University, two cabinets
Caltech, one cabinet.

Galaxy YH-1

China

Vector Register 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 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.

HEP

Denelcor, Inc.
17000 E. Ohio Place
Aurora
Colorado 80017

8-303-337-7900

Dr. Burton Smith - architect

Shared Memory Multiprocessor

The Heterogeneous Element Processor (HEP) is an MIMD machine with two levels of parallelism. Each Process Execution Module (PEM) can run asynchronously, and all can have access to the common storage through a proprietary switch. Although the HEP has been designed for use with up to 16 PEMs, the largest built was a 4-PEM machine. Each PEM is itself an MIMD machine with parallelism achieved through an instruction execution pipeline. Up to 64 user-defined tasks can be executing concurrently, but the length of the pipeline on a 1-PEM machine effectively limits the degree of parallelism to between 8 and 16, depending on memory accesses. The memory accesses are also pipelined. An instruction progresses to the next stage of the pipeline every clock cycle of 100 nsec, although a memory fetch or store can be proceeding simultaneously.

The CPU uses MSI ECL, mostly ECL 10 K with a gate delay of 3 ns, although some critical circuits use ECL 100 K with a .75-nsec gate delay. SECDED memory is used throughout.

Parallelism is obtained in Fortran by explicit task creation (with minimal overhead), and synchronization is by means of asynchronous variables.

Program, constant, register, and data memories all use 64-bit words.

- Program memory size is from 32 Kwords to 1 Mword.
- There are 2048 registers, and the minimum size of the read-only constant memory is 4096 words.
- The data memory is separate from the CPU and can be expanded in 128-Kword increments to a maximum of 1M words (8 Mbytes) per PEM. Memory access time is 50 nsec, and half and quarter word and byte addressing is possible.

Configuration:

The HEP switch that connects memory with CPUs is a flexibly configured, programmable network which uses packet switching techniques to route messages. Each node on the switch network has three full-duplex ports. Arbitration is through a priority system based on longevity. The propagation time through a node is 50 nsec.

Although designed as a stand alone system, it is probably best to front-end the machine with a machine with good interactive capabilities like a VAX.

Software: A version of UNIX III is used as the operating system, although not all utilities are available. The debugging and diagnostic capabilities are poor. Floating point uses IBM-compatible 32- and 64-bit formats. Little software outside of linear algebra kernels is available.

Languages: Fortran 77, C, and Pascal are available in addition to HEP assembler.

Performance: Each PEM is rated at 10 MIPS, and speeds in excess of 7 MFLOPS have been achieved on one PEM for linear algebra kernels coded in HEP assembler language. It is rare to exceed 3 MFLOPS for purely Fortran code on one PEM.

Cost: The cost of a 1-PEM configuration is around $3 million.

Status: Company filed Chapter 11 in 1985. No systems operational. HEP2 plans uncertain.

Hitachi S-810

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

415-872-1902

Vector Register Architecture

The Hitachi 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).

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 megawords.

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 pipe to/from memory, each capable of loads/stores at a rate of two word (64 bits) per cycle.

The scalar speed of the Hitachi S-810 may be slower than either the CRAY X-MP or Fujitsu VP-200.

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.

IBM 3090/VF

IBM
Neighborhood Rd
Kingston, New York 12401

In Europe:
David Marshall
IBM Warwick Engineering, Science and Industrial Centre
PO Box 31
Birmingham Road
Warwick CV34 5JL
England
0926-32525 Telex 311601

Vector Register Parallel Shared Memory Architecture

The IBM 3090 is the top end system available from IBM. It uses the System/370 Extended Architecture for scalar operations.

18.5 nsec cycle time.

3090 Model 150 is a uni-processor with 32 MB or 64 MB of central memory.

3090 Model 180 is a uni-processor with 32 MB or 64 MB of central memory and 64 up to 256 MB extended storage.

3090 Model 200 is a dyadic processor with 64 MB of central storage and up to 256 MB of expanded storage.

3090 Model 400 is a four-way processor with 128 MB of central storage and up to 512 MB of expanded storage.

For the 3090 each processor has a high-speed-cache of 64 KB. The cache is system controlled.

Vector Facility (VF):

Optional feature to the 3090.

Pipelined vector processor with vector registers.

Each VF has 8 vector floating point registers of 128 64-bit elements.

171 vector instructions are added for the VF.

32-bit operands in the VF are treated as 64-bit operands. Fixed stride addressing on vectors is allowed as well as indirect addressing or mask control.

Each VF has has a theoretical peak performance of 108 MFLOPS.

Models 150 and 180 can have 1 VF added.
Model 200 can have one or two VFs added.
Model 400 can have one, two, three, or four VFs added.

System Software:
MVS/XA
VM/XA
VM/SP High Performance Option

Languages:
Assembler H Version 2
VS Fortran 2 including Library Program Multitasking Facility and
Interactive Debug.
Engineering and Scientific Subroutine Library.

The Fortran compiler will automatically vectorize existing codes.

Power consumption:
7.8 KWatts

Closed water/air cooled.

171 Sq. Ft.

Cost:
3090 Model 200 rough cost is $5M, VF option is 10 per cent per processor additional cost.

International Parallel Machines Inc. (IP-1)

Robin Chang

International Parallel Machines, Inc.
700 Pleasant Street
New Bedford, Massachusetts 02740

617-990-2977

Parallel Architecture

Sales:
Walter Stuart Pye
V.P. Marketing
6767 Forest Hill Ave.
Suite 305
Richmond, VA 23225 U.S.A.
804/272-5678 Telex 888648

Technical:
Dr. Robin Chang
President
700 Pleasant St.
Top Floor
New Bedford, MA 02740
617/990-2977 Telex 888648

Parallel Architecture

Proprietary CPUs (9 used in base system) (IP-1-9)
Local and global-shared memory
NxN crossbar interconnection switch
32-bit physical memory addressing, expandable to 48 bits; 64-bit data paths
80M to 430M main memory
170M to 3G disk space
double-precision IEEE standard
9CPU system, 133 MFLOPS double precision
72 MIPS (9 CPU configuration)
52 I/O ports

Configurations:
- Stand-alone
- VAX front-end
- IBM MVS front-end
- IBM MVS front-end
- various VME/Unix workstation front-ends
- Symbolic processing workstation front-end (Prolog or Lisp)

Can add:
- 1/2-inch tape drives
- multiple disk drives running in parallel
- plotters and printers
- close-coupling high speed communication interface to other CPUs
- TCP/IP, HyperChannel
- more CPUs up to 33 for 1987 delivery, up to 1025 CPUs for 1Q 1989 delivery

Software
- UNIV System V.3, up 64 users, real-time version available
- C with IP parallel math routines called from library
- Fortran 77-to-C converter
- Fortran 77 (VAX compatible)
- IP-1 virtual machine package for software developers, IBM-AT and VAX hosts, with debugging facilities, nominal charge

Application Software Available:
- Database management, printed circuit board layout, oil reservoir simulation, seismic data analysis, will port serious applications depending on market potential

Performance:
- 9-CPU peak, 144 MFLOPS double precision IEEE
- 33-CPU peak, 600 MFLOPS double precision IEEE

Status:

First machine delivered October, 1985

Oil reservoir simulation beta sites in progress

multiple OEM contracts

Cost: $22K to $1M+, plus possible application porting charges

Scientific, aerospace, engineering, military and university users

Intel's Personal Supercomputers (iPSC)

Intel Scientific Computers
15201 NW Greenbriar PW
Beaverton, Oregon 97006

503-629-7600
General Manager: Robert Rockwell
Applications Manager: Cleve Moler
Marketing Manager: Charlie Bishop
Marketing and Customer Support: Ellen Bailey

In Europe:
David Moody
Intel Scientific Computers
Intel International Limited
Pipers Way
Swindon SN3 1RJ ENGLAND

Hypercube Architecture

Developed from Caltech work on Cosmic Cube.

The cube manager, or intermediate host, is a 286/310 workstation with 2-4 Mbytes of memory, a 140-Mbyte Winchester disk, a 320-Kbyte floppy, a proprietary ethernet connection to the hypercube itself, and a TCP/IP ethernet connection to remote hosts. The manager runs Xenix.

The hypercube has 32, 64, or 128 nodes, termed the iPSC/d5, d6, or d7. Each node consists of an 80286 CPU, an 80287 floating point coprocessor, and 0.5 megabytes of memory. The 80287 has IEEE arithmetic with 32-, 64-, and 80-bit formats and a speed of about 30-50 kiloflops. Each node also has 8 bi-directional communication channels rated at 10 Mbits/sec per channel. One of the channels is used for communication with the cube manager and the other 5, 6, or 7 are used for communication with other nodes in the cube.

The basic system may be modified by replacing node boards with memory expansion boards or higher speed floating point vector boards. A memory board increases the node memory from 0.5 to 4.5 megabytes. The resulting systems are known as the iPSC-MX/d4, MX/d5 and MX/d6. Software available from Gold Hill called CCLISP, for Concurrent Common LISP, provides communicating LISP environments for each node of the MX systems.

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 megabyte of 250 nsec data memory, and 16 kilobytes 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 Megaflops on a single node, while a dot product involving short precision vectors in the small, fast memory can approach 20 Megaflops. Peak floating point rates of the VX systems, obtained by multiplying the peak rate of a single node by the number of nodes, reach 424 megaflops for long precision and 1280 megaflops for short precision on a 64 node, iPSC-VX/d6. VAST II, a Fortran vectorizer from Pacific Sierra Research, is expected to be available in the summer of ’87.

Software:
Manager operating system: Microsoft Xenix 3.0
Node executive: Intel NX
Languages: Fortran, C, LISP, FCP (Flat Concurrent Prolog), ASM286,
Ada under development.
Tools: CCLISP, VAST II, Debugger, Crystalline Operations System (Caltech),
Cosmic Environment (Caltech), NETCUBE (Oak Ridge)

Physical characteristics of one 32-node cabinet:
16 x 16 x 19 inches;
footprint 26 x 26 inches;
180 lb.

Cost and performance summary:

System	Nodes	Memory	MFLOPS	Price

iPSC/d5	32	16 MBytes	2	$155K
iPSC/d6	64	32 MBytes	4	$280K
iPSC/d7	128	64 MBytes	6	$525K
iPSC-MX/d4	16	72 MBytes	2	$176K
iPSC-MX/d5	32	144 MBytes	4	$306K
iPSC-MX/d6	64	288 MBytes	6	$556K
iPSC-VX/d4	16	24 MBytes	106	$250K
iPSC-VX/d5	32	48 MBytes	212	$450K
iPSC-VX/d6	64	96 MBytes	424	$850K

Loral Dataflo

Loral Instrumentation
8401 Aero Drive
San Diego, California 92123

619-560-5888

Parallel Dataflow Architecture

The Loral DATAFLO system is a parallel processor that can be incrementally expanded from approximately 10 processors to approximately 256 processors. Each processor is composed of two National Semiconductor NS32016 microprocessors. One processor is dedicated to token (data) management and store and the other is dedicated to application execution. The application processor has a National Floating Point Unit associated with it. The applications processors each have 128 K of local static RAM that is used for application execution.

In general, communication between processors is via messages (dataflow tokens). Communication is handled on a 32-bit time mutliplex bus. This bus is used to broadcast dataflow tokens that have 16 bits of tag and 16 bits of data. A large dataflow system is composed of multiple chassis, with at most 14 dataflow processors programmed to pass dataflow tokens between chassis. Since these interfaces pass only those tokens that they are programmed to pass, bus saturation within a chassis is minimized. Shared memory can be added to the system in 2-Mbyte increments by replacing a dataflow processor with a shared memory board. Shared memory can be accessed by any processor in the chassis via a device bus that is separate from the dataflow bus.

A program is composed of two components, a data graph description and a set of graph node implementations written in some standard language like C or Fortran. Applications development and monitoring system activity is accomplished through a dedicated UNIX-based processor occupying a position in one of the clusters.

The "grain" size for the system is approximately the size of a procedure, around 60 to 100 lines of source code.

A wide variety of real-time I/O and data storage controllers may be included in the dataflow environment through an extension of the dataflow bus.

Price: $67K to $2M

Meiko

Meiko Incorporated
6201 Ascot Drive
Oakland, CA 94611
(415) 530 3055   Telex 797748

In Europe:
Meiko Limited
Whitefriars
Lewins Mead
Bristol BS1 2NT
England
(0272) 277409 Telex 449731   Fax (0272) 277082

Parallel MIMD architecture

Founded in 1977.
First shown in July 1985 at SIGRAPH in San Francisco.

Contact: Roy Bottomley and Miles Chesney (England)

Parallel MIMD architecture

The founders of Meiko were the managers of the design group responsible for the transputer and its peripherals. Thus the whole design philosophy of the Meiko system units is based around the INMOS Transputer. These are available in three flavors:

T414-15 15MHz 32-bit 7.5 MIPS
T414-20 20MHz 32-bit 10 MIPS
T800-20 20MHz 32-bit 1.2 Mflops sustained (peak of 3 Mflops)

Connection topology is user configured, either
(i) hardwired by means of wire wrap, patch links or PCBs plugged onto the backplane, or
(ii) by electronic configuration. The connectivity is defined by the program. A distributed electronic switch implements this connectivity on the computing surface.

Each unit contains a transputer processor with eight unidirectional 10Mbit/sec autonomous message channels. These communication channels can be used for high-speed direct memory access or for low latency message passing to or from other computing elements.

Communication between units is by explicit I/O or message passing.

Message passing is a single instruction in which the appropriate process scheduling is achieved in an efficient microcode sequence.

The units are:

Local host with 3Mbytes 15Mbytes/sec error-checked RAM and 128Kbytes of 10 Mbyte/sec EPROM. IEEE 488 and dual RS232 I/O interfaces. At least one local host is required in any system (computing surface).

Computing element. The only memory is that of the transputer, namely, 256Kbytes of 15Mbytes/sec error-checked RAM.

Mass store with 8Mbytes of 15Mbyte/sec error-checked RAM and 2Mbytes/sec DMA controlled SCSI disk and peripheral interface. The third level of this memory hierarchy is 2048 bytes of single cycle static RAM for frequently accessed local variables.

Display which has 128Kbytes of private SRAM, 1.5 Mbytes dual-ported display store. 70 MHz pixel rates and 200Mbytes/sec pixel highway. CCIR/RS-343-compatible video with programmable sync generator supports interlace and non-interlace.

The units are held in slots in the Computing Surface. The local host, mass store, and display each require one slot, but the Compute Board contains 4 computing elements and occupies only one slot.

The units are grouped as Computing Surface Modules which can themselves be combined to form the Computing Surface. Two standard modules are the 10-slot M10 and the 40-slot M40.

The Computing Surface contains an infrastructure to facilitate debugging. The Computing Surface can be used stand alone or as an attached resource to a VAX, SUN, IBM PC, or Prime.

Software includes :
Occam II compiler
and the sequential language compilers ...
C
Fortran 66
Fortran 77
Pascal

BCPL

Current applications include molecular modelling, naval simulators, computational fluid dynamics, lattice gauge theory, quantum chromodynamics, ray tracing, and solution of partial differential equations.

A single M40 module with computing elements employing the T800 transputer is capable of a sustained performance of 187 Mflops.

Price depends on system which is ordered. Prices for a fully operational system start at around $13K (the M10). The M40 Computing Surface Module with 39 computing elements and a local host costs around 250K pounds ($417K). This configuration has 157-way parallelism, total MIPS rating of 1175, and 42 Mbytes of RAM.

Computing Surfaces containing over 300 processors, spread across several modules have been demonstrated. The delivery of a 1024 processor, 1 Gbyte, 1 Gflop, 3M pounds Computing Surface is expected during the middle of 1987.

First deliveries were in March 1986 and since then over 2 dozen machines have been shipped.

MIPS Computer Systems, Inc.

John Hennessey
MIPS Computer Systems, Inc.
930 Aqures Ave.
Sunnyvale, CA 94086

408-720-1700

RISC Technology

This is a new organization (2 years old), with about 95 people, including the founders John Hennessey, John Moussouris, and Skip Stritter.

Architecture:

Family of Products:
- component kits
- boards
- development systems

Family of CPU boards.
3 - 5 - 8 MIPS (VAX 1.0 MIPS)
Custom floating point 3 MFLOPS
IEEE arithmetic

Software: UNIX (C, IEEE Pascal, Fortran 77)

Cost: $4,000 for the OEM board

Status: products are shipping now

Goodyear MPP

Goodyear Aerospace Corporation
1210 Massillon Road
Akron, Ohio 44315

Ken E. Batcher
216-796-4511

Parallel 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 neighbor 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 data 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.

Language: Parallel Pascal

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

Multiflow

Donald E. Eckdahl
Joseph A. Fisher
Multiflow Computer, Inc.
175 N. Main St.
Branford, CT 06405

203-488-6090

VLIW (Very Long Instruction Word) Architecture

Performance: Vector/parallelism capabilities by different techniques

Software: IEEE standard arithmetic and UNIX

Applications: Scientific engineering market

Languages: Fortran, C, F77 VAX extensions

Cost: under $1 million

Myrias 4000 System

Martin Walker
Myrias Research Corporation
200 - 10328 - 81st Avenue
Edmonton AB T6E 1X2
Canada

(403) 432 1616 Telex 037 - 42759

Martin Walker - R&D Program Manager
UUCP:ihnp4!alberta!myrias!maw

Parallel Architecture, hierarchically Managed Local Memory

Main design goal of the architecture is scalability of memory capacity and performance.

Each processing element (PE) contains one Motorola MC68000 (10 MHz) and 512 Kbytes of 150 nsec DRAM with DMA interface to a board level bus; a multiple processing element (MPE) board contains 8 PEs, a supervisory PE and an interface to a printed wire backplane; 16 MPEs fit in one card cage. Card cages have eight parallel ports for communication with other card cages, or with external devices; they can be interconnected in a fractal network of arbitrary size; physical packaging in Krates of eight cages (1024 PEs; 512 Mbytes of memory).

The architecture supports the Myrias memory model: independent parallel tasks execute in distinct memory spaces; spaces are merged upon task completion; these memory spaces are not tied to particular PEs. Virtual memory (32 bit addressing) and the hierarchical clustering of PEs provide a distributed cache system.

Architecture is implemented as a virtual machine on which all user software (applications, compilers, editors and optimizers) runs. The virtual machine provides user transparent virtual memory (paging and scheduling) and run time support to user processes. The virtual machine can run on many different hardware substrates; hardware failures are circumvented by the machine's control mechanism.

Configuration:

- off the shelf components
- two sided printed circuit boards
- two kinds of board
- maintenance by on site board swap
- stand alone system
- standard network interface (eg. VME) or to suit customer

Arithmetic: 32-, 64-, and 128 bit floating point; 8-, 16-, 32-, and arbitrary precision fixed point; IEEE 754 option.

Software:

- UNIX System V and BSD 4.2 operating system (user visible)/Myrias 4000 (user transparent)
- upwards compatible with existing serial computers

Languages available: Myrias Parallel Fortran (Fortran 77 with parallel DO loops, recursion and dynamic array dimensions); Myrias Parallel C (ANSI C with parallel DO loops).

Fortran characteristics:

- single instruction provides access to parallelism (parallel DO loops)
- upwards compatible with Fortran 77
- will run conforming programs
- will have parallel debugging aids
- recursive parallel programming methods allow straightforward implementation of optimal divide and conquer algorithms which can minimize computational complexity

Applications: physical modeling (neutron transport, magnetic fusion, drug design, chemical engineering, quantum chemistry, aerodynamics and hydrodynamics, seismic processing and hydrocarbon recovery, geophysics, meteorology, and structural design); data processing (image processing and generation, searching and sorting); VLSI design; algebraic manipulation. Will provide (recursive parallel) mathematical library.

Performance: proportional to size of configuration; achieved through scalable architecture and algorithmic reduction of computational complexity.

Status: prototype 1986. Cost: price proportional to performance; more than $1M.

AS/91X0

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

Integrated Vector Processor

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

AS/9140/50 are uni-processors with 48 MB of central memory.

AS/9160 is a uni-processor with 64 MB of central memory.

AS/9170/80 is a dyadic processors with 64 MB 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.

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

46 vector instructions are added for the VPF.

32-bit operands in the VPF are treated as 64-bit operands. Fixed stride addressing on vectors is allowed as well as indirect addressing or mask control.

Based on the Hitachi S-9 plus IAP.

System 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.

Closed water/air cooled.

Cost: rough cost is $3M

NCUBE

Sales Office:
700 E. Baseline Rd., Suite D1
Tempe, AZ 85283

Headquarters:
1815 NW 169th Place
Suite 2030
Beaverton, OR 97006

John Palmer (602)839-7545

Hypercube Architecture

Node Processor
  Custom VLSI
  11 Interrupt driven DMA channels at 2 Megabytes/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 (e.g. 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: 128 Kbytes SECDED about 110 KB available for application

Processor Board (16"x22") contains 64 nodes + 8 MBytes SECDED memory
Host Board (16"x22") contains
  Intel 80286/80287 with 4 Megabytes SECDED memory
  1 ESMD Disk Interface for up to 4 disks (160,330,500 Megabyte)
  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 Megabytes/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; 16 local NCUBE nodes; 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.

Configurations

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

NCUBE/seven: 16 to 128 Nodes; 14" wide by 29" by 29"; 110v; office environment; 4 slot backplane: 2 for I/O options, 2 for Processor Boards; 160 or 330 MB disk 16 MB 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.

routing; message typing; process debugging support

Fortran 77 and C are available.

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

Price: NCUBE/ ten or seven: $40K(cabinets+peripherals)+$60K/Host Boards+$100K/Processor Boards (University discount available) NCUBE/four: $10K/board (4 nodes) + $4K O.S. license.

Schedule: Betasites working with I/O systems since February, 1985

Product announcement November 18, 1985, SIAM meeting on Parallel Processing

First complete system shipments in December, 1985

Approximately 30 systems sold and installed.

NEC SX-1E, SX-1 and SX-2

Mr. S. Adams
NEC Information Systems
1414 Massachusetts Ave.
Boxborough, Massachusetts 01719

617-264-8800

In Europe:
Garry Foley
Manager - Marketing
Communications Systems Division
NEC Business Systems (Europe) Ltd.
NEC House
164-166 Drummond Street
London NW1 3HP
01-388-6100 Telex 261914 NEC LDN
Fax : (01) 387 4867 (GIII)
(01) 388 5704 (GIII)

Vector Register Architecture

The SX system has two processors, the Central Processor (CP) and the Arithmetic Processor (AP) sharing the main memory. 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 user programs executing. Although SX runs in standalone mode, NEC supports its ACOS series mainframes and also IBM mainframe connections.

	SX-1E	SX-1	SX-2
Cycle time	7 ns	7 ns	6 ns
Number pipes	4 v-pipe	8 v-pipe	16 v-pipe
Length regs	20K v-reg	40K v-reg	80K v-reg

Architecture:

AP 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.
- 1 Kbit bipolar memory with 3.5 nanosecond cache memory access time.
- 256 Megabyte memory (512-way interleaving) with 2 Gigabyte extended memory.
- 64K bit static MOS memory chip with 40 nanosecond access time, giving a memory-to-register rate of 11 Gbytes per second.
- Register-to-register machine with 40 (80 on the SX-2) Kbytes of vector registers.
- register-to-register with far more (and more flexible) vector functional units.

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 Architecure

- The extension of the NEC mainframe computer.
- virtual storage support.

Software:
- does not run the IBM instruction set (unlike other Japanese computers)
- 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).
- uses its own operating system

Languages: Fortran 77, ALGOL, PL/I, BASIC, Pascal, C, LISP, PROLOG and COBOL. In vector mode only Fortran is supported.

Performance: Maximum rating of the SX-1E is 285 MFLOPS and of the SX-1 is 570 Megaflops and of the SX-2 is 1.3 Gigaflops. Peak performance for the SX-2 will be in the 1.3 Gigaflop range. 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 U.S. was July 1986. The NEC machine is available for benchmarking. NEC has sold seven of its supercomputers in Japan and in the USA.

Cost:

SX-1E: \$8-9 million
SX-1: \$10-12 million
SX-2: \$14-16 million

NUMERIX MARS-432

Numerix Corporation
20 Ossipee Road
Newton, MA 02164
(617) 964 2500

In Europe:
Numerix UK Limited
Ambassador House
181 Farnham Road
Slough SL1 4XP ENGLAND
(0753) 29411
Attn : Martin C Allen .. Director of Sales and Marketing

Company formed in 1980 as co-operative exercise between
Analog Devices Inc and Standard Oil (Indiana).

Pipelined Array Processor

32-bit floating-point array processor. Clock cycle time is 100 nsec. There are two pipelined adders and one pipelined multiply that can each deliver one result per cycle. Simultaneously two data reads or one write can be performed. Computational power of 30 Megaflops (32-bit arithmetic). Access of memory from arithmetic pipes is via a cross-bar switch. Data memory of 64 Megabyte of directly addressable memory. Program memory of 4K words cache and virtual memory space (64 word pages) of up to 64K words. A path exists between the memories so that programs can be stored in the data memory.

Communication with the host is through an interface box and a 5 MHz 32-bit data bus with control through a second bus (the CBUS). DMA transfers at I/O bus rates of 20 Megabytes/second. Interfaces currently exist to DEC machines, ELXSI (Embos), Apollo (Aegis), and Sperry (OS1100) systems.

Software includes:
Fortran development system
Microcode development system
AP run-time executive support package
Application libraries including mathematics, signal processing, and
geophysical processing.
1024-point complex FFT in 1.7 msec.

Dimensions are 19"w x 21"h x 24"d Weight 180 lbs.

Customers include :
ERIM (Michigan), Honeywell, Naval Research Lab, Kodak, Pratt & Whitney, Naval Weapons Center (USA)
and Rolls-Royce, GESMA, Ensign Geophysics, Queen Mary College, and BGT (Europe).

PS 2000 (Russian supercomputer)
U.S.S.R.

Parallel Architecture (SIMD)

Today in the Soviet Union there is assembly line production of PS-2000 computers with a capability of up to 200 million ops. All these processors (the number of which varies with modifications of the machine) do the same operation at the same time or are in the wait mode.

The PS-2000 complex is classified as SIMD (single instruction stream-multiple data). The complex includes an SM-2 and the PS-2000 processor. The latter consists of 8-64 processor elements, each with its own memory of 4K-16K 24-bit words. All processor elements are under common control. The complex was 'first commissioned' in 1980. Unspecified type of addition speed is 0.3 microsec, with a memory access or cycle time (source gives both in heading without saying which the number applies to) of 0.64 microsec.

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 12K or 48K-byte 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 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.

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

The PS-2000 can have 8, 16, 32, or 64 processors, and these 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 16K 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 U.S.-made 30Mflop Numerix 432 fills half of a 22' rack.

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

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.
Cost: "retails at 800,000 rubles".

SAXPY Computer Corporation

SAXPY-1M

B. Friedlander
Director of Advanced Technology
SAXPY Computer Corporation
255 San Geronimo Way
Sunnyvale, California 94086
408-732-6700

Reconfigurable systolic architecture

The machine has 5 basic components:
- System Control Unit (DEC Micro VAX II)
- Matrix Processing Unit (Systolic processor)
- Vector Processing Unit (Numerix MARS 432)
- System Memory (64 MB to 2 GB)
- SAXPY Interconnect (320 MB/sec transfer rate)

Stand-alone computer
- With capability to connect to VAX family of equipment
- Additional interface -
  - High speed mass storage subsystem (HMS) (100 MB/sec)
    - Connection to disks, tapes, VME, hyperchannel
  - Network Input/Output (NIO) - VAX Cluster Interface

The matrix processing unit is a linear array of 32 systolically connected processors.
MPU to system memory trasnfer rate is 62.5 MWords/sec.
64 nsec cycle time.
32-bit floating point arithmetic.
Peak performance 1000 MFLOPS.

Software on the System Control Unit:
- VMS Operating system Fortran 77 Pascal Ada C Matrix Math Subroutine Libraries. Access to the MPU is through subroutine calls.

Size: 95.2" wide x 78.2" high x 40.4" deep Power: 15 KWatt Air cooled

Cost: $2 million base price

SCS-40

Scientific Computer Systems Corporation
25195 S.W. Parkway Ave.
Wilsonville, OR 97070

503-682-7223

President: Bob Schuhmann
Technical: Carl Haberland

In Europe:
Pierre Hassid
Scientific Computer Systems Corp.
5 Villa Alexandrine
92100 Boulogne Billancourt
France
+33-1-48.25.73.47

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:
- 8-, 16-, 32-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.)
- designed as to interface to a front end, either VAX 11/780 or VAX 11/750. (Interfaces planned for CRAY X-MP, IBM 4300 series, and NSC 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:
- software licensing agreement with CRAY.
- 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

Performance:
- peak of 44 MFLOPS 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 $500,000. 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.

Sequent Balance 21000

Ron Parsons
Sequent Computer Systems, Inc.
15450 SW Koll Parkway
Beaverton, Oregon 97006-6063

503-626-5700
800-854-0428
Telex 296559

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

Chicago Office
Karl von Spreckelsen
District Manager
200 Tri-State International Drive
Suite 110
Lincolnshire, IL 60015-1480
312-940-9299

In Europe:
SEQUENT UK
Chris Arnold
Compass Peripheral Systems
Bridge House
Faraday Road
Newbury
Berkshire RG13 2DH ENGLAND
(0635) 33933   Telex 846301

Incorporated in January 1983
(old name of company was Sequel)

Parallel Bus Architecture

Machine 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 32-bit wide pipelined packet bus supporting multiple, overlapped memory and I/O transactions with a sustained data transfer rate of up to 27 Mbyte/sec.

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 256K-RAM components. Processor and memory boards can go in any slot on the SB21000 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.

Software: The operating system, called DYNIX, is a version of Berkeley 4.2bsd UNIX, enhanced for application-transparent multiprocessing and user-controlled parallel processing. Among the enhancements are a completely reentrant kernel, 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.

Industry-standard I/O, interfaces:
  MULTIBUS - has terminal multiplexor with controllers
  Ethernet - at 10 Mbits/sec. Connnection to PC as virtual disk through Ethernet.
  SCSI - at 2.5 Mbyte/sec. Offers 5-1/4 in. disk drives (72 Mbytes formatted) and streamer tape drives with adaptor boards for the SCSI bus.

Peripherals include a 1/2" tape drive and a 396-Mbyte disk drive asynchronous

The packaged system includes a 26-slot SB21000 bus backplane and an 21-slot MULTIBUS backplane; can take up to fifteen dual processor boards.

Other features:
  Table height packaging.
  Dimensions 30.5" x 23.25" x 28.625" (HWD)
  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

Software: supports 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

Parallel programming library callable from any language.
Extension to Fortran to allow shared common blocks.

Performance:
Fully populated machine seen as 21 times a VAX 11/780 in power.
Designed as a high throughput system, with support for parallel processing at user level.

Status: Shipments began 12/84, and Sequent currently has manufactured more than 140 systems (as of Nov 86).

Cost:
$286,000 for the complete machine with all software, 10 processors, 8-Kbyte cache/processor, 16-Mbyte memory, and four Fujitsu Swallow 264 MByte disks (total of 1056 MBytes); $140,000 for a 4-processor system; and $62,000 for a small 2-processor Balance 8000 system.

Silicon Graphics Inc

Forest Baskett
Silicon Graphics
2011 Stierlin Rd.
Mountain View, CA94045
415-960-1980

Very High Performance Workstation

Goals:

Heavy emphasis on interactive graphics for large computational problems.

Markets: CAD/CAM/CAE Molecular Modeling Image Processing Scientific/Engineering Research and Development

Unisys Integrated Scientific Processor System ISP 1100/90

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

Vector Parallel Architecture

The ISP operates under the control of one basic Integrated Scientific Processing system consists of a Unisys 1100/90 CPU with one I/O Unit, the ISP, and a 4 M-word Scientific Processor Storage unit.

The peak performance of a single ISP is a 133 MFLOPS in single precision (36 bit word) and 67 MFLOPS in double precision (72 bit word). Two ISPs may be connected to a single Unisys 1100/90 host system.

The high speed memory that supports the ISP is capable of transferring data to an ISP at 133 M-words/sec. The sustained performance is 20 to 30 MFLOPS in double precision and may double for single precision.

First delivery was June 1986.

30 nsec clock

16 MWords memory

Peak performance single precision (36 bits) 133 MFLOPS

Peak performance double precision (72) 67 MFLOPS

Cache based 4K words (36 bit words) - Scalar processor part only/vector processor can address into cache

register to register architecture; vector register set is 16 x 64 words.

Vector processor also has scalar imbedded processor.

Hetrogeneous processing system - up to four scalar processors (IP) and two vector processors (ISP).

Vectorizing compiler UFTN

ST-100

Star Technologies Inc.
515 Shaw Road
Sterling, Virginia 22170

703-689-4400

Technical: Phil Cannon

In Europe:
Stephen D Bean
Star Technologies Inc.
Rosemount House
Rosemount Avenue
West Byfleet
Surrey KT14 6NP ENGLAND
09323 5281   Telex 928764 STAR G

Pipeline Floating Point Architecture

The ST-100 is an array processor, designed to attach to a more general-purpose computer or host via bus.

It has four independent programmable processors. A separate processor is dedicated to each of the following functions: external data flow, internal data flow, arithmetic processing, and synchronization. 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 consists a 320 nsec memory, 8-way interleaved, composed of 64K dynamic RAMs with SECDED. It is expandable to 32 Mbytes in increments of 2048 Kbytes. 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 8192 32-bit words for a total of 192 Kbytes. During each machine cycle, four cache references are permitted: three by the arithmetic processor and one by the storage/move processor. Information flow from host to main memory to cache to functional unit to cache to memory to host.

Other features:
40 nsec clock cycle,
Bipolar VLSI circuits with 1200 gates.
32-bit floating point arithmetic, pipelined functional units.
2 adders, 2 multipliers, and a 480 nsec divide/square root functional unit.
Ambient air cooled
Size 56" x 33" x 67"

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 7 device adapters; 12.5 Mbyte/sec data transfer rate.

Software:
Fortran-like control language (APCL)
Macro assembler
Simulator/debugger and Linker
Library Maintenance Program
Applications Library available.

Fortran compiler implemented using KAP Pre-compiler from
Kuck and Associates.

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

Cost: $265,000 base price.

Stellar

Wallace E. Smith
VP Sales
Stellar
100 Wells Ave.
Newton, MA 02159
617-964-1000

Very High Performance Workstation

Company founded by John Poduska (from Apollo).

Goals:

Heavy emphasis on interactive graphics for large computational problems.

Price $75K - $125K

Availability: 2nd half 1987

Markets:
CAD/CAM/CAE
Molecular Modeling
Image Processing
Scientific/Engineering Research and Development

Vitesse Electronics

741 Calle Plano
Camarillo, CA 93011

805-388-3700

Parallel Architecture

Plans are to build a scalar machine with with 1Gbyte memory and a 40-nsec cycle time. The machine will be made of CMOS. It is to support hardware optimization for high run-time performance.

Configuration: First machine is to have up to 8 processors. Connectivity allows for large number of processors, in the thousands. It can be used as a co-processor on a VAX.

Software: 32-64 bit floating point arithmetic supporting the IEEE standard.

Languages: Fortran, Pascal, and C.

Performance: 25 to 150 MFLOPS (uniprocessor range of performance as result of optional hardware boards for each processor).

Status: Started in July 1984, expect to produce a machine by late 1986. It is planned to make a GaAs version in a couple of