Cray Computer Corporation
1110 Bayfield Drive
Colorado Springs, CO 80906

Telephone 719 579 6464
Facsimile 719 540 4256

CRAY 4

High Performance Computing Engines

Outstanding Performance Unprecedented Value

Over the past four or five years the idea has arisen that the only way to achieve significant improvement in the price/performance ratio for powerful computer systems was through the use of off-the-shelf commercial components. Some have even suggested that the days of custom-designed computer systems are past.

We would like to introduce you to a direct contradiction of the above idea . . .

CRAY 4 HIGH PERFORMANCE COMPUTING ENGINES

- Unprecedented Price/Performance Advantage
- 100% Gallium Arsenide Logic
- 1 Nanosecond Clock
- 2 Gigaflops per Processor
- Up to 256 Gigaflops* per System
- Up to 8 Gigawords* of Directly Addressable Fast SRAM
- Unmatched Scalability for Vector-Scalar, Highly Parallel or Mixed
- Memory Bandwidth of 2 Gigawords per Second per CPU
- 100 MBytes per Second per 32-Bit HIPPI Channel**
- IEEE Floating Point
- Extended UNIX Operating System

*These figures are for a four-node cluster using our existing memory part. Memory size will double with a new memory part to be introduced later in 1995.
**Support for 200 megabyte per second, 64-bit HIPPI channels is planned for mid-1995.

Introducing the CRAY-4

The CRAY-4 high performance computing engine is a natural successor to the CRAY-3, offering more than double the performance of the CRAY-3 per processor and as much as ten times the performance for a fully configured system. Even greater levels of performance and scalability can be obtained through *clustering*—the next logical step beyond multiprocessing.

Unmatched Performance and Value

Amazingly, these improvements in performance and scalability have been gained at a price point that is both unprecedented and unmatched by any other computer system in its class. Cray Computer Corporation was able to dramatically reduce the cost of production for the CRAY-4 by building on the tremendous investment in tooling and technology for the CRAY-3, and by simplifying the architecture of the CRAY-4 to significantly reduce the number of hardware components.

The physical size of a CRAY-4 processor has been reduced by almost a factor of four with the use of second generation gallium arsenide integrated circuits with ten times the level of integration of those used in the CRAY-3. The resultant reduction in pathlengths, combined with other improvements, has contributed to more than a doubling of individual processor performance.

Extensive Scalability

The building block of a CRAY-4 system is a 20 module set with four processors, 256 million words of fast Static Random Access Memory and four industry standard High Performance Parallel Interface channels for I/O. This 20 module set (also called a *quartet*) represents the smallest CRAY-4 configuration available. This system serves as a low-cost entry level system with a peak performance of eight gigaflops.

Multiples of this basic building block are offered in larger, single cabinet configurations of up to 32 processors, two gigawords of fast SRAM memory and 32 full-duplex, 32-bit HIPPI channels. Memory capacity can be increased independent of the number of processors to meet specific application requirements. Available memory sizes will double with the introduction of a new memory part later in 1995.

Even larger configurations can be achieved by clustering several single cabinet systems (referred to as *nodes*). In addition to increasing processing power and memory capacity, clustering provides greater potential I/O bandwidth and enhances the resilience of the system to individual node failures.

Natural Architectural Progression

The CRAY-4 system architecture is a natural progression from the architectures employed in the CRAY-1, CRAY-2 and CRAY-3 systems. Therefore, users of those systems will find a high level of compatibility in the CRAY-4. Improvements have been made in several areas. Some of these are explicit performance enhancements while others are simplifications that remove redundant functionality and help to reduce the overall cost of the system.

Commitment to Standards

The introduction of the CRAY-4 high performance computing engine underscores Cray Computer Corporation's commitment to providing standards-compliant, high-end computing solutions. The CRAY-4 is the first high performance computer system to support IEEE floating point numeric formats. This allows the CRAY-4 to act as a true peer in a distributed computing environment with high-end workstations.

Input/Output on the CRAY-4 is accomplished using standard industry channels such as Ethernet and HIPPI. This gives CRAY-4 users maximum flexibility in the selection of peripheral and network vendors.

Innovative Software

The CRAY-4 software environment is a direct extension of that developed for the CRAY-3. It provides complete compatibility with that system as well as a high degree of user level compatibility with systems offered by Cray Research, Inc. Multiprocessing applications may be extended to exploit the newly defined cluster environment. Enhancements are being made to Cray Computer Corporation's CSOS operating system to handle the sophisticated resource management requirements implicit in a distributed system. The compiling software is also being enhanced to support the latest language standards such as Fortran 90 and High Performance Fortran.

System Organization

The CRAY-4 high performance computing engine is based on the proven shared-memory, multiple-vector processor model employed in previous generations of Cray Computer Corporation and Cray Research, Inc. systems.

Nodes and Clusters

The CRAY-4 extends this concept into an additional dimension by introducing the ability to connect multiple CRAY-4 systems together in a cooperative topology or *cluster.*

The building block of a CRAY-4 clustered system is referred to as a *node.* A node contains up to 32 processors and up to two gigawords of very fast SRAM memory housed in a single cabinet. A cluster comprises multiple interconnected nodes.

Clusters may be configured to reflect specific bandwidth, redundancy and cost requirements. For example, a 16-processor CRAY-4 system could be configured in several different ways as shown in the diagram in the next column. This hierarchical approach to multiprocessing provides the potential for higher bandwidths and better hardware resilience than could be achieved in a monolithic multiprocessing design. At the same time the distributed memory and inter-processor communication issues found in many Massively Parallel Processing implementations are minimized. This flexible approach also provides the customer the ability to configure very powerful systems; for example, a four-node CRAY-4 cluster could have 128 processors and eight gigawords of memory offering asymptotic performance approaching 256 gigaflops.

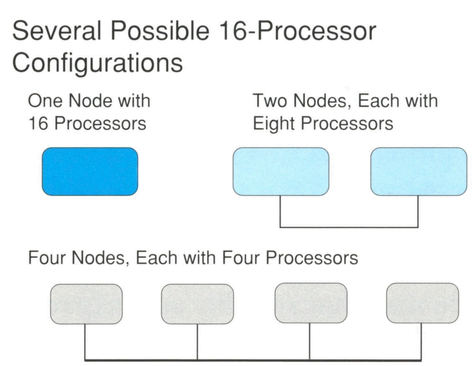

Inter-Nodal Connections

Software applications requiring minimal inter-nodal bandwidth may use one or more HIPPI channel connections. However, many applications will require much higher bandwidths. We are developing a new generation of channel networking technologies with these applications in mind. The initial implementation, planned for the middle of 1995, will provide a full-duplex, two gigabytes per second connection between each node with a maximum of four nodes per cluster. Subsequent implementations will extend both the bandwidth and the maximum allowable nodes, including the capability to interface very high speed OEM devices to the CRAY-4 clustered system.

Input/Output

Each CRAY-4 node supports HIPPI channels for I/O capabilities with optimal bandwidth. These capabilities also provide for easy reconfiguration in the event of an individual node failure. Unlike the CRAY-2 and CRAY-3 systems which used a separate foreground processor to manage I/O activity, the CRAY-4 manages I/O operations via operating system software in the CRAY-4 processor.

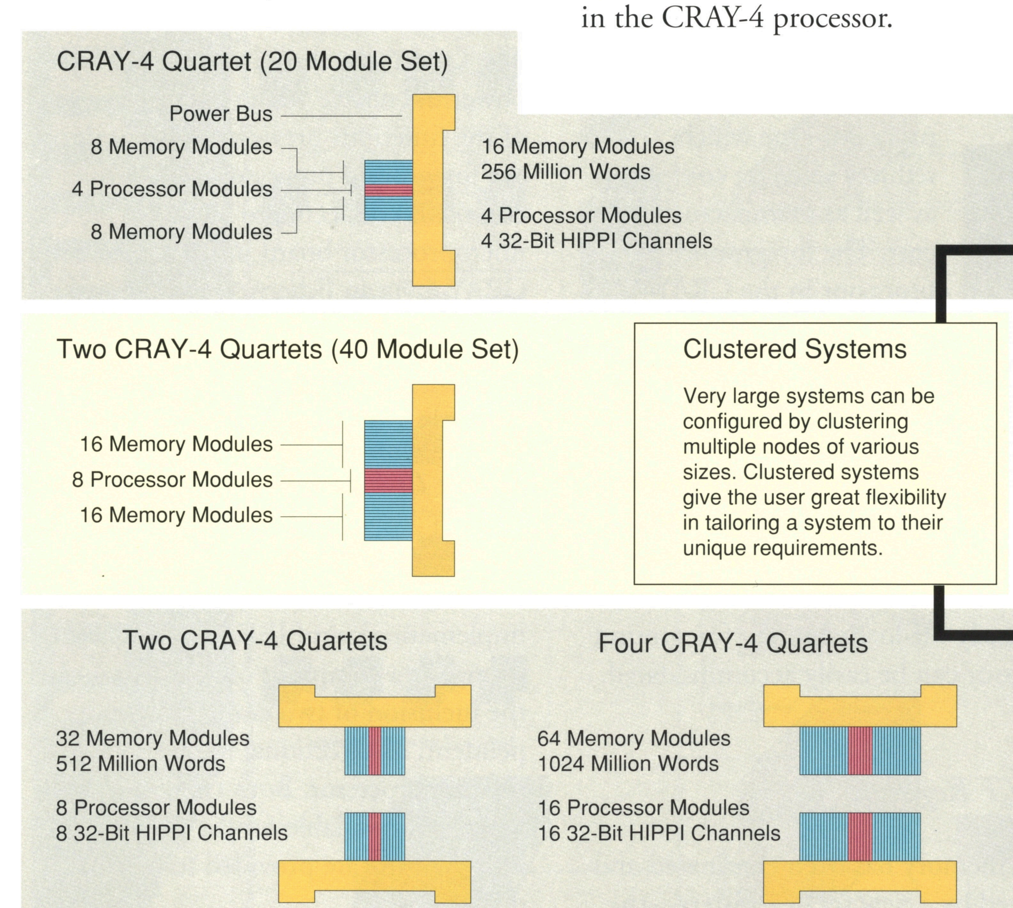

Below: The cabinet for a single quartet node measures just 53 inches wide by 34 inches deep. The cabinet for a two or four quartet node measures 77 inches wide by 43 inches deep. Both cabinets are 48 inches above the floor.

System Architecture

The CRAY-4 architecture is both simplified and enhanced in comparison to the architecture of the CRAY-3. Simplifications include the elimination of the foreground processor and local memory. Enhancements include additional registers, improved memory access, chaining of functional units and IEEE floating point.

New Single-Module Processor

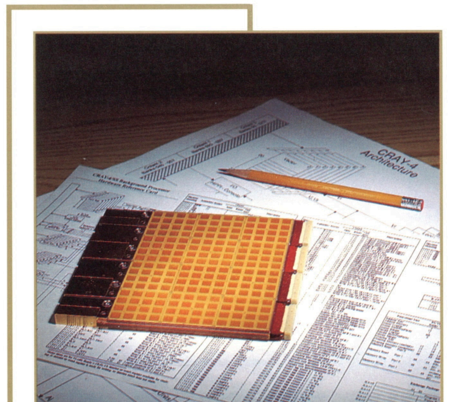

The CRAY-4 employs a single generic processor type which reduces software complexity as well as manufacturing cost. The foreground processor in the CRAY-3 was used to provide overall management of the system—especially control of input/output operations—independent of the operation of the background processors. With the increase in speed of the CRAY-4 processors, the overhead previously off-loaded to the foreground processor can be easily accommodated.

New T Registers

Local memory has been eliminated and replaced by a new set of registers—the temporary (T) registers—which act as a 64-element secondary cache to the scalar (S) registers and the address (A) registers, as well as providing a path to and from the vector registers. All the registers in the CRAY-4, including the address registers, are 64 bits wide.

New More Powerful Console

The CRAY-4 includes a new and more powerful console which assumes some of the functions previously assigned to the foreground processor. The new console is connected to a MIPS microprocessor board inside the CRAY-4 via an Ethernet channel and operates under its own UNIX operating system.

Other Enhancements

Further enhancements provided in the CRAY-4 architecture include the first implementation of IEEE floating point format in a computer of this class and the inclusion of two completely independent, bi-directional memory ports on each processor. Both tailgating, which was available on the CRAY-3, and chaining are provided for vector operations.

Each CRAY-4 Processor Performs Two Billion Floating Point Operations per Second.

CRAY 4 ARCHITECTURE

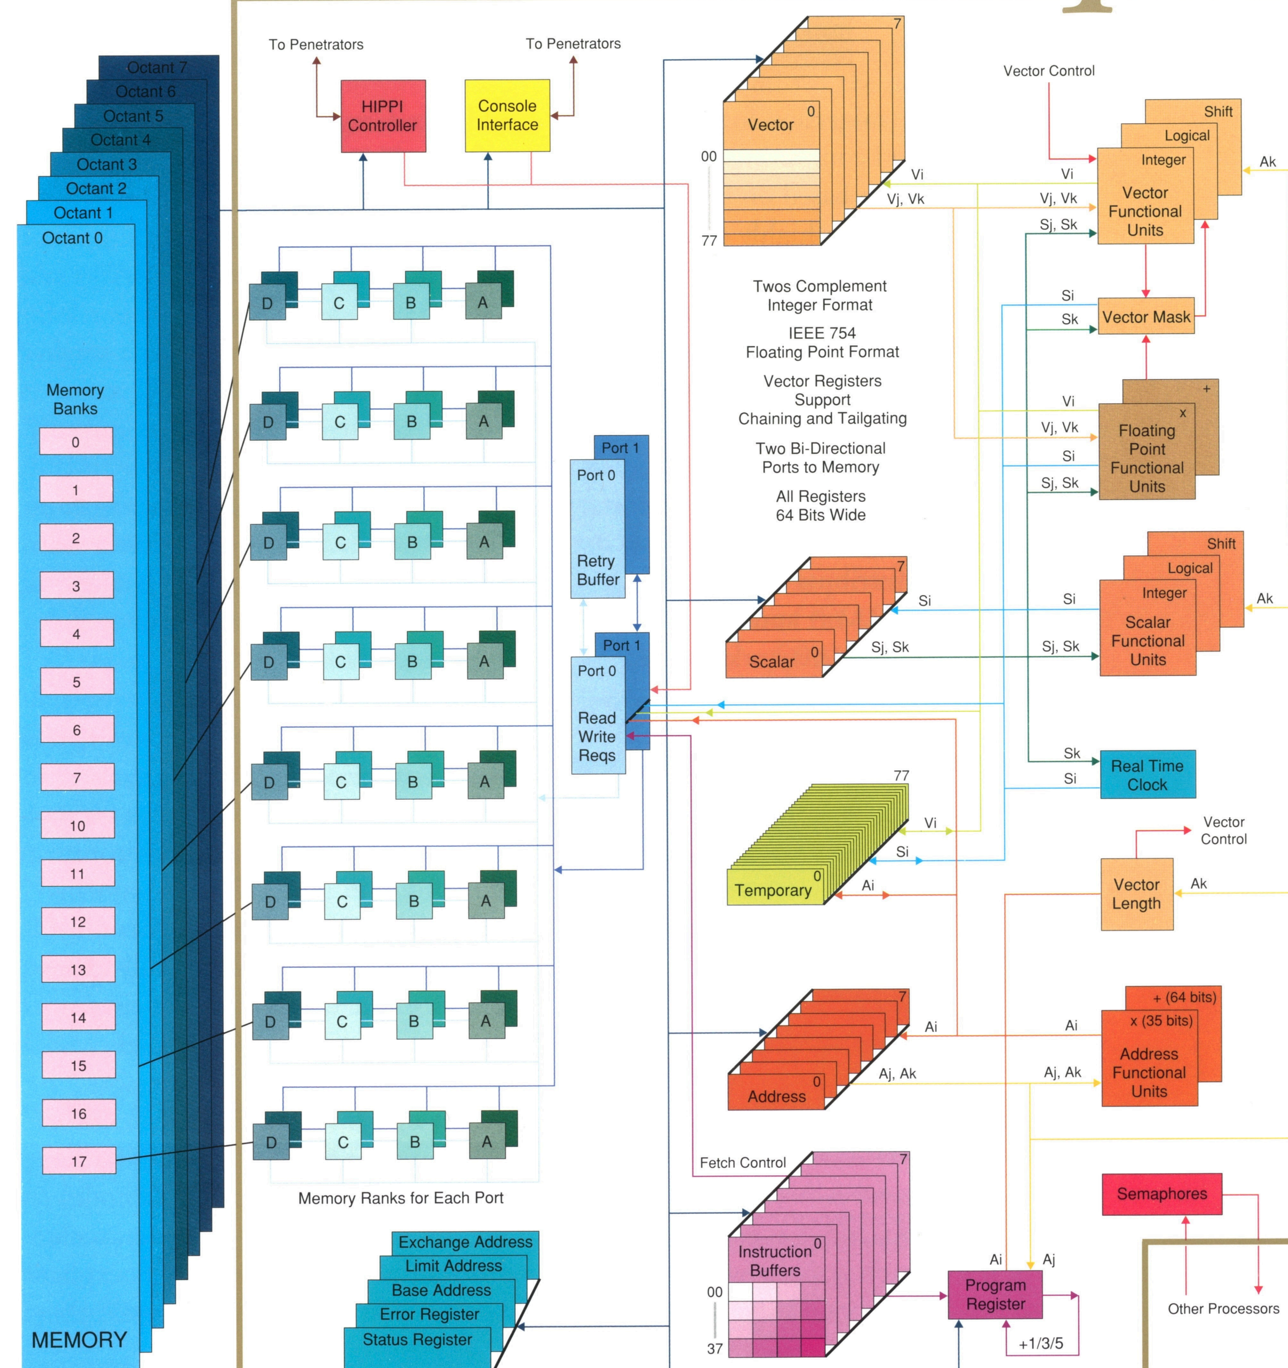

Software Environment

The CRAY-4 software environment provides a simple migration for users of all major high-end computing systems. An operating system based on UNIX is complemented by a highly efficient compiling system supporting both Fortran and C. A range of user interface products focuses on increased productivity through the use of powerful visualization capabilities. Compliance with industry standards has been a high priority contributing to the ability to interchange codes readily with other systems as well as to distribute applications between the CRAY-4 and other systems in a heterogeneous network environment.

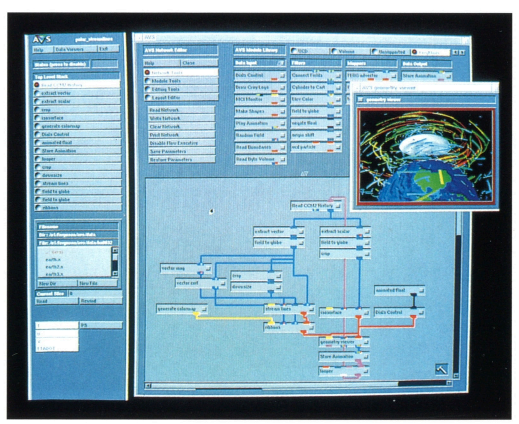

AVS 5.02 has been fully implemented and optimized on the CRAY-3 from the user interface and distributed computing through powerful 3-D graphical rendering capabilities.

Many new capabilities have been added to Cray Computer Corporation's software library to support new features of the CRAY-4 system. Our strategy is to introduce and enhance our support for multi-nodal (clustered) CRAY-4 systems through a series of software releases. The initial release will treat a single CRAY-4 node as an individual vector-parallel system (comparable to current CRAY-3 support) with multiple CRAY-4 nodes treated as a loosely coupled network. Subsequent releases will move progressively towards an environment that fully supports distributed applications and services.

Programming Environment

The Fortran programming environment provided with initial CRAY-4 systems will provide intra-nodal support only; i.e., an application will take full advantage of all the capabilities of a single CRAY-4 node including parallel exploitation of multiple processors within a node. In this initial release automatic support for the distribution of applications across multiple nodes will not be supported. However, such distribution can be achieved by explicit message passing at the user application

level utilizing a message passing library such as PVM or MPI.

The CRAY-4 Fortran programming environment provides ANSI 77 support as well as support for many Fortran 90 features. Cray Computer Corporation also offers an ANSI compliant C compiler and an assembler for the CRAY-4 system.

We view migration to a native Fortran 90/HPF environment as the principal mechanism for supporting fully distributed, data-parallel applications. Subject to demand, we may offer an interim solution based on a Fortran 90/HPF to F77 translator. This would generate data-parallel structures and library calls which would utilize the underlying message passing library for inter-nodal communication.

No support for parallel C is planned at this time. However, if a standard for parallel C is defined Cray Computer Corporation will implement it on the CRAY-4.

We also offer several tools which are essential to the CRAY-4 development environment. These include the bdb interactive debugger, the emacs editor and various user interface products such as the Watson object oriented libraries, X Windows and Motif. The development of multi-nodal support in these products will mirror the development plan for the compiler outlined above.

Operating System

Cray Computer Corporation's extended UNIX operating system (CSOS) has been enhanced to support the new hardware features of the CRAY-4 system. A new, more powerful console interface is provided giving the operations staff considerably greater capability than was present on the CRAY-3. Some of the functionality assigned to the CRAY-3 foreground processor has migrated to the CRAY-4 console.

In initial CRAY-4 configurations each CRAY-4 node will run its own copy of CSOS which will manage the node's resources independent of any other CRAY-4 nodes. This includes the scheduling of multiple processors within a node, both across multiple applications and across multiple tasks within one or more applications.

Basic communication between multiple CRAY-4 nodes is supported via TCP/IP and a message passing library over a HIPPI channel. This enables users to define explicit message passing between applications on multiple nodes. It also facilitates inter-nodal file access via NFS and FTP.

We plan to move the operating system towards a fully distributed, asynchronous environment. This will be achieved through the progressive implementation of a microkernel-based technology with many of the current CSOS capabilities becoming part of a *server* hierarchy.

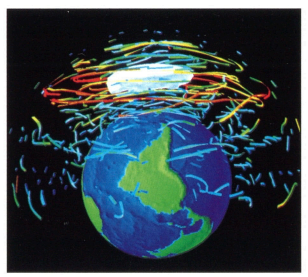

Above: Depletion of Ozone over the South Pole. Ribbons show wind direction and speed by their shape and color. The polar cloud is represented as an isosurface of temperature. Study run on the CRAY-3 by NCAR scientists.

A fully distributed environment will include distributed resource management and file systems. It will also provide failsoft software technology to allow the rapid and automatic reconfiguration of a multi-nodal CRAY-4 system in the event of individual node failure.

Software Applications

A wealth of applications await the user of a CRAY-4 high performance computing engine. Third-party investments in code for the CRAY-2 and CRAY-3 provide a simplified path for bringing codes to the CRAY-4, including codes for global atmospheric studies, large-scale meteorological modeling, reactor safety, computational fluid dynamics, structural analysis, seismic processing, reservoir modeling, biochemical prediction, powerful 3-D imaging and much, much more. Cray Computer Corporation will assist customers and third-party application vendors to insure that the customer's requirements are satisfied.

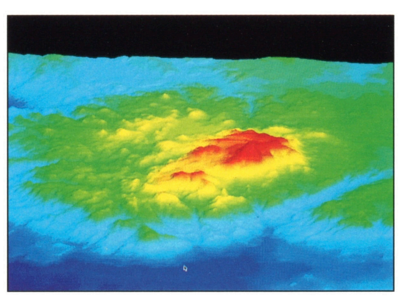

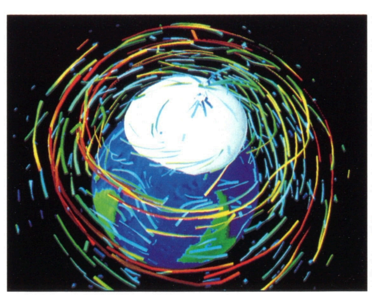

Above Left: Digital Terrain Map of the Pikes Peak Region created with AVS software, a CRAY-3 and data from the U.S. Geological Survey. The color scale is based on elevation. The cursor marks Cray Computer Corporation's headquarters.
Above Middle: The earth's South Pole with polar cloud and stream lines.

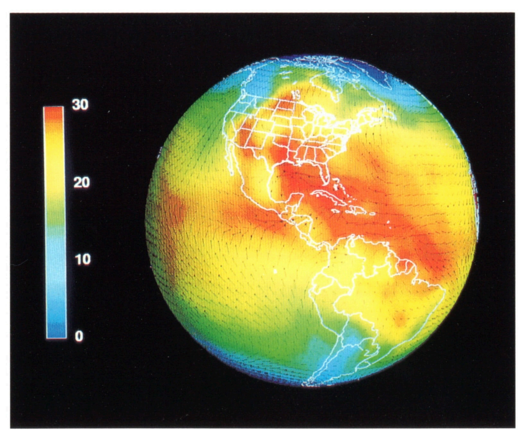

Above: Our earth with white political boundaries and black, surface wind vectors. The color of the earth's surface represents temperature in degrees C. Created from global climate model data generated on a CRAY-3 by NCAR scientists.

CRAY-4 Technology

The CRAY-4 is based on second generation enhancements to the technology developed for the CRAY-3 system. The 500 gate-equivalent gallium arsenide integrated circuits are replaced by 5,000 gate-equivalent GaAs circuits custom designed and manufactured at Cray Computer Corporation's GaAs facility in Colorado Springs. CRAY-4 memory is based on fast SRAM memory parts produced by Toshiba in Japan. Initial CRAY-4 systems will use the same 4 Meg by 1 parts used in the CRAY-3. Later in 1995 we plan to begin use of a 4 Meg by 4 memory part which will essentially double memory size and potential bandwidth.

All logic and memory integrated circuits are mounted on modules slightly larger but otherwise very similar to those used in the CRAY-3. The twisted pair inter-module connections used in the CRAY-3 are replaced by micro-coaxial cables. The mechanical design and cooling technologies used in the CRAY-3 have been refined allowing the CRAY-4 to reside in a single cabinet possessing a smaller footprint than even the CRAY-3. This is in marked contrast to most competitors whose state-of-the-art machines seem to be increasing in size with each new version.

The clock speed of the CRAY-4 has been doubled to one billion cycles per second (one gigahertz) giving a clock period of only one nanosecond. Since no commercial testing system was available to test the die at this speed while still in wafer form, our engineers designed and developed new Modular Acquisition and Generation Circuitry (MAGC) to test the die in wafers.

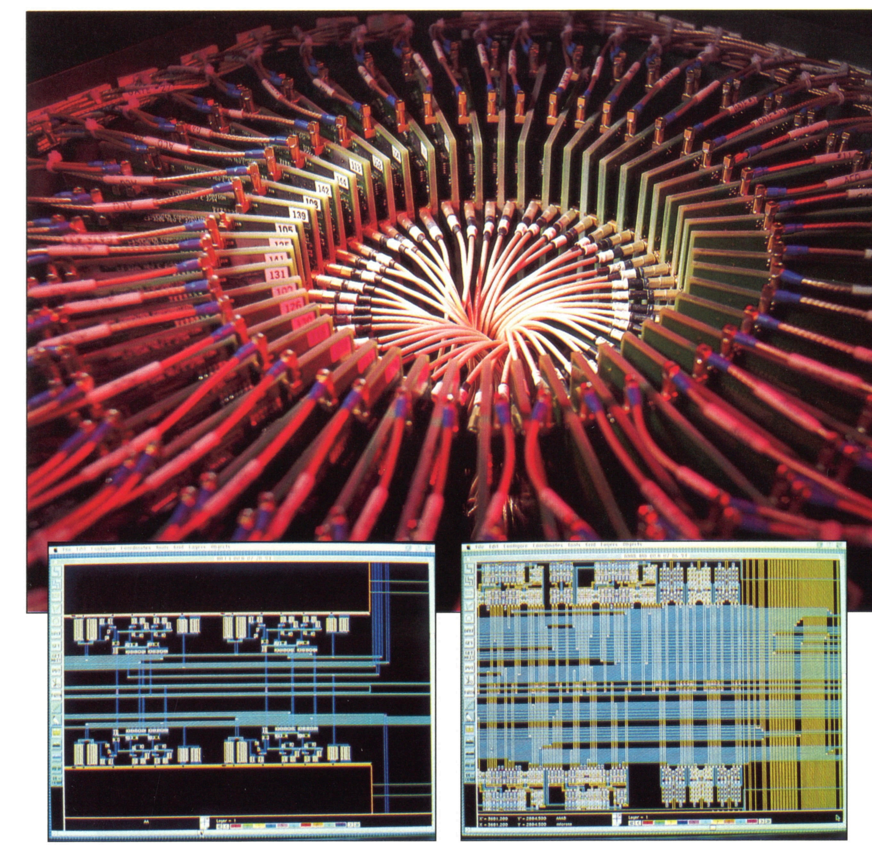

The large photo above shows a portion of the new CRAY-4 die testing system. The smaller photos illustrate the tenfold increase in component and signal density of the CRAY-4 ICs. Each photo is showing an area 400 microns in height for each IC. A CRAY-3 IC is on the left; a CRAY-4 IC is on the right.

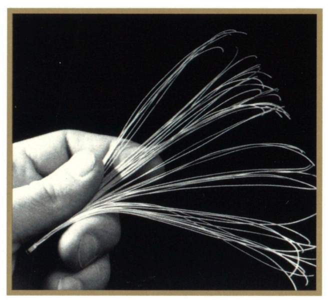

Above: The micro-coaxial cable used for inter-module communication.

Comparing the Technologies

	CRAY-3	CRAY-4
Logic Circuits	GaAs SDFL 500 Gate-Equivalent 3.835 x 3.835 mm	GaAs DCFL 5000 Gate-Equivalent 5.4 x 5.4 mm
Memory Circuits	Silicon CMOS SRAM 4 Meg x 1 25 ns Cycle Time 8 MWords per Module	Silicon CMOS SRAM 4 Meg x 1 (4 Meg x 4) 21 ns Cycle Time 16 MWords per Module
Modules	4.1 x 4.1 x 0.25 Inches Four Modules per CPU 69 Electrical Layers 22000 Z-axis Connections Twisted Pair Interconnect	5.2 x 5.2 x 0.33 Inches One Module per CPU 90 Electrical Layers 36000 Z-axis Connections Micro-Coaxial Interconnect
Cooling	Chilled Water/Fluorinert	Also Air-Cooled Versions
Cabinets	System Tank and C-Pod	System in One Cabinet
Typical Footprint	252 Square Feet	215 Square Feet

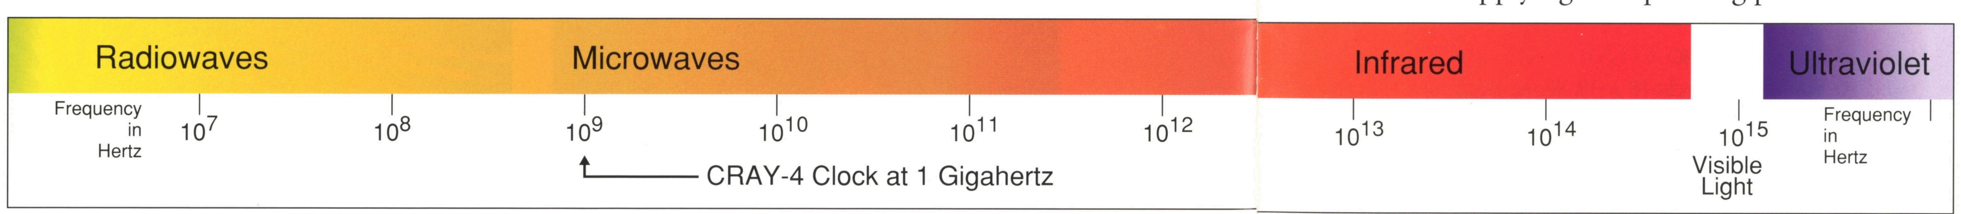

Left to Right: The micro-coaxial cable passing through the eye of a needle.

A CRAY-4 GaAs die with 121 pins for power and signals.

An exploded view of a CRAY-4 module with 90 electrical layers in only 0.33 inches.

The four modules comprising a CRAY-3 background processor next to the one module of a CRAY-4 processor.

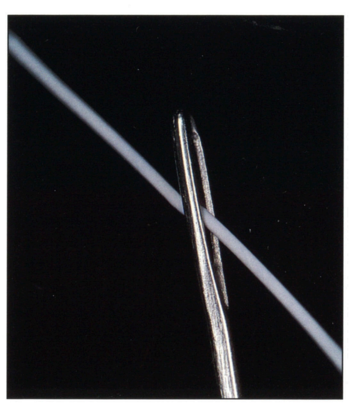

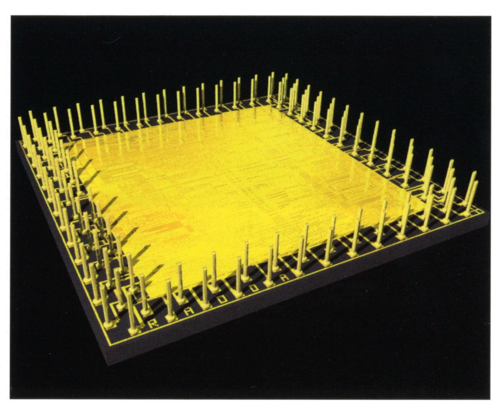

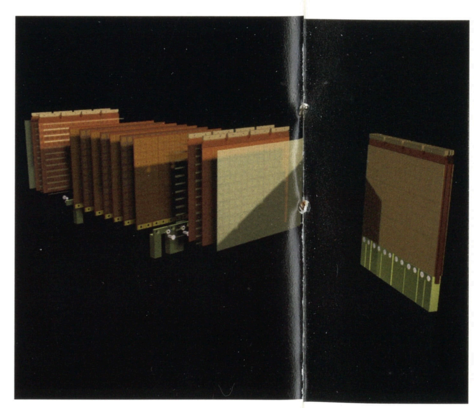

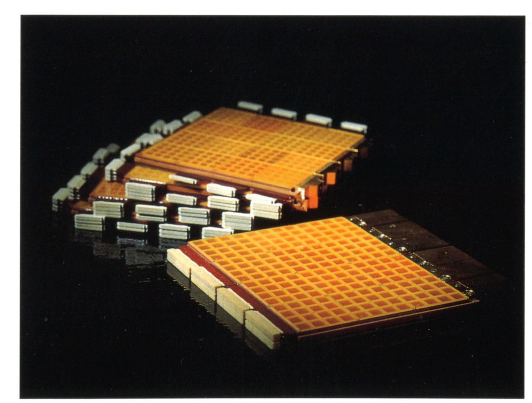

The MAGC system, our GaAs fabrication facility and our precision circuit board manufacturing capabilities are striking examples of the technological investments now beginning to pay dividends. It is clear that still greater potential exists; for example, the ability to increase GaAs die density even further without the need for significant additional investment.

Quartets, Nodes and Clusters

The basic building block of a CRAY-4 high performance computing engine is a quartet. Quartets comprise a 20 module set sharing a common power bus for supplying the operating power and mechanical stability to the modules. This 20 module set includes four processor modules, each with one full-duplex, 32-bit HIPPI channel, and 16 memory modules with 256 million words of SRAM.

A CRAY-4 node normally comprises a single physical cabinet containing one, two, four or eight quartets along with the associated power supplies, cooling equipment, environmental monitoring equipment, HIPPI translator cards, cable bulkheads and MIPS-based microprocessor interface for the external console. In nodes containing more than one quartet, pairs of quartets are combined into 40 module sets (eight processors with 32 memory modules). The largest CRAY-4 node currently planned comprises eight quartets with 32 processors and 128 memory modules (two gigawords of SRAM).

Larger and more powerful systems can be configured by clustering multiple CRAY-4 nodes. The individual nodes can be small or large systems. In some configurations it may be possible to house multiple nodes in a single cabinet. This ability to create clustered systems using various size nodes gives users great flexibility in scaling a CRAY-4 high performance computing engine to their particular requirements.

Specifications

Hardware

PROCESSORS
2 to 32 in a Single Node (Cabinet)
Up to 128 in Four Nodes

MEMORY SIZE
256, 512, 1024 or 2048 Million Words
64-Bit Words with SECDED.
These sizes will double later in 1995.

EXTERNAL I/O CHANNELS
1 HIPPI Channel per Processor

CLOCK
1 Nanosecond

MEMORY TRANSFER RATE
2 Gigawords per Second per Processor

PEAK PERFORMANCE
2 Gigaflops per Processor
(2 Billion Floating Point Operations per Second)

Software

OPERATING SYSTEM
Extended UNIX

COMPILERS
Fortran ANSI 77 with Fortran 90 Extensions
ANSI C

NETWORKS
TCP/IP

USER TOOLS
Editors
Debuggers
X11r5
Motif

Processor Architecture

FUNCTIONAL UNITS
3 Vector
3 Scalar
2 Floating Point (Shared Vector/Scalar)
2 Address
Chaining and Tailgating
IEEE Floating Point Format

REGISTERS
8 Vector
1 Vector Mask
1 Vector Length
8 Scalar
64 Temporary
8 Address
1 Program
1 Base
1 Limit

INSTRUCTION STACK
256 64-Bit Words (8 x 32 Word Buffers)

64-BIT REAL TIME CLOCK

UP TO 64 SEMAPHORE FLAGS

2 BI-DIRECTIONAL PORTS TO MEMORY

Interfaces

CLUSTER
2 Gigabytes per Second (Later in 1995)

CONSOLE
Ethernet

NETWORKS
HIPPI (32- and 64-Bit)

DISK STORAGE
RAID Disks via HIPPI

The Company We Keep

Cray Computer Corporation was established as an independent company on November 15, 1989. The spin-off agreement with Cray Research, Inc. provided for the transfer of currently owned assets, people and initial funding. A transfer of patents and cross licensing of technology allows each company to pursue its own specific projects unencumbered by patent or technology conflicts. Cray Computer Corporation is a public company listed on the NASDAQ Exchange.

The company's mission is to design, manufacture, sell and support high performance, general purpose scientific computers. The company's first product was the CRAY-3 computer system introduced to the marketplace in 1993.

The company employs approximately 360 people at two facilities in Colorado Springs, Colorado. One facility, near Cheyenne Mountain, encompasses corporate offices, design, development, software, testing, and manufacturing areas as well as a complete state-of-the-art GaAs fabrication facility. The other facility, near the Garden of the Gods, manufactures the extremely small, precision circuit boards used in the CRAY-3 and CRAY-4 computer systems.

The CRAY-1 (the world's first supercomputer), the CRAY-2, the CRAY-3 and the CRAY-4 are four extraordinary machines that share one thing in common: a very uncommon gentleman, Seymour R. Cray. Mr. Cray serves as the Chairman and Chief Executive Officer of Cray Computer Corporation.

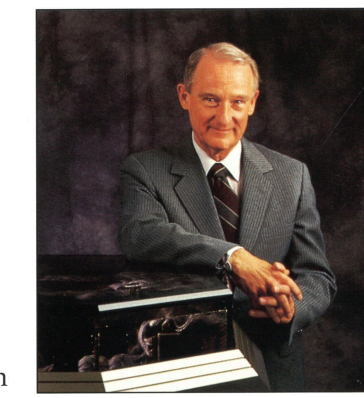

Mr. Cray was the 1994 recipient of the MCI Leadership Award. The award is given annually to an innovator who has demonstrated outstanding leadership and whose work has had a beneficial impact on society.

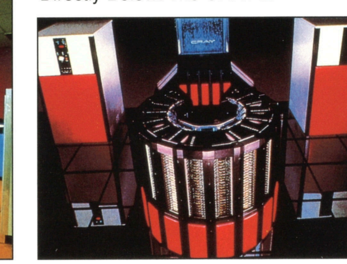

Left: The CRAY-1.
Directly Below: The CRAY-2.

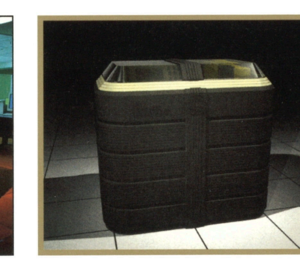

Right: The CRAY-3.
Far Right: The CRAY-4.

The 1976 CRAY-1 has 566 modules per CPU, 12.5 ns clock, 80 megaflops per CPU. The 1985 CRAY-2 has 34 modules per CPU, 4.01 ns clock, 250 megaflops per CPU. The 1993 CRAY-3 has four modules per CPU, 2.08 ns clock, 960 megaflops per CPU. The CRAY-4 has one module per CPU, 1.00 ns clock and two gigaflops per processor (two billion floating point operations per second).

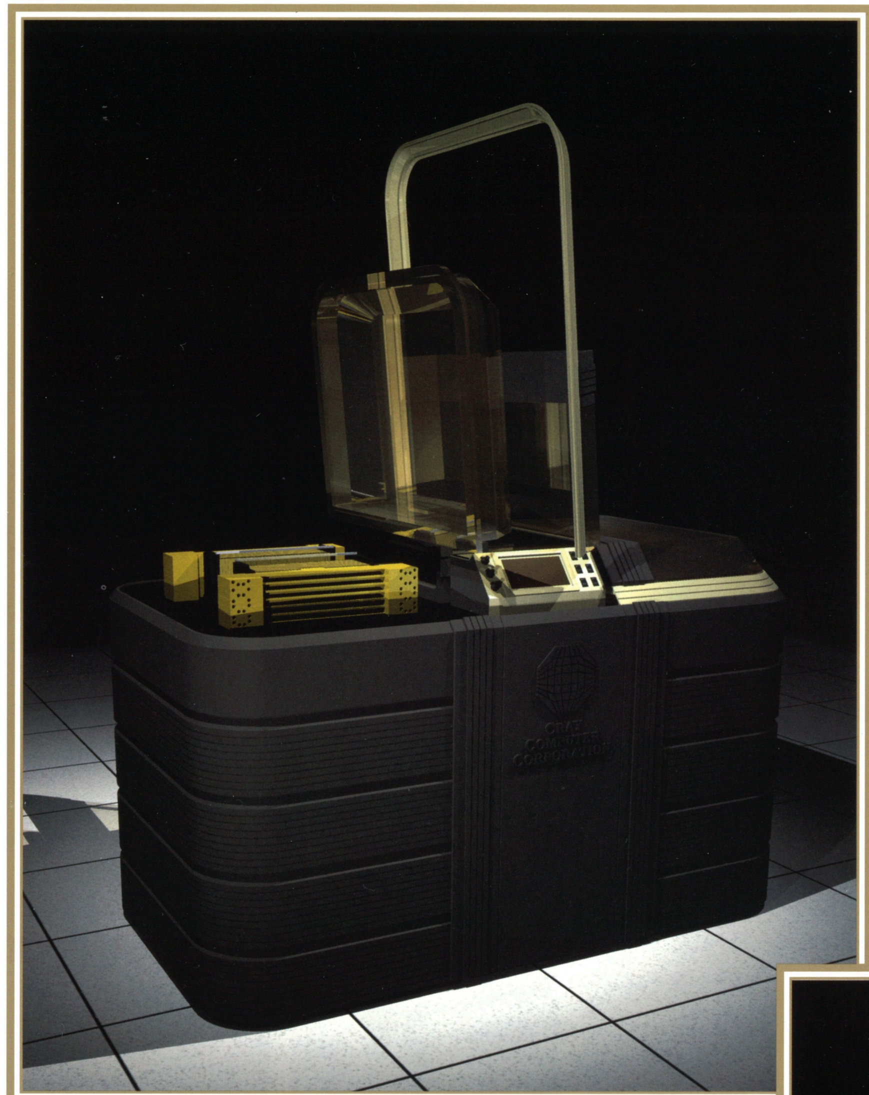

CRAY 4 HIGH PERFORMANCE COMPUTING ENGINES

Outstanding Performance · Unprecedented Value

"Thank heaven for start-up companies or we'd never make any progress. People who get unhappy with structure in companies can move on and start their own, take the big risks and occasionally find the pot of gold. I think that's just wonderful."

Cray Computer Corporation

CRAY-1 is a registered trademark of Cray Research, Inc.
CRAY-2 is a trademark of Cray Research, Inc.
CRAY-3 and CRAY-4 are trademarks of Cray Computer Corporation.
All other companies or product names referenced are believed to be trademarks or registered trademarks of their respective holders.

Editor and Design: Dean Tollefson, Artwork Department.
Photography: Steve Gombosi, Software Department.
Digital Color Renderings: Doug Cowles and Tom Sibert, Mechanical Engineering Department.
Software Application Views: Scott Ferguson, Software Department. Data for views of South Pole ozone depletion study and global temperatures was provided by Byron Boville of the National Center for Atmospheric Research, Boulder, Colorado.
Additional Photography: Photo of Mr. Cray, Jeff Stine Photography. Photos of the CRAY-1 and the CRAY-2, courtesy of Cray Research, Inc.
Lithography: Four-color waterless process, National Printing & Packaging Corporation, Denver.
Digital Scanning: Universal Graphics, Boulder, Colorado.

The product specifications contained in this brochure and the availability of the products are subject to change without notice.

Printed November 1994.