The CONVEX C Series

The first family of affordable supercomputers.

CONVEX

The leader in affordable supercomputing

Today's competitive markets require nothing less.

Supercomputers from CONVEX Computer Corporation are helping scientists and engineers all over the world make breakthroughs in product design.

Researchers in computational chemistry are using CONVEX systems as a dry lab, simulating reactions between proteins to uncover the nature of recombinant DNA.

Automotive and aeronautical engineers are using CONVEX systems for computational fluid dynamic studies, simulating wind tunnel tests and obtaining diagnostic information without building expensive prototypes.

Electrical engineers are using CONVEX supercomputers to develop new semiconductors; financial institutions, to do financial modeling; government agencies, to break cryptic codes; and petroleum corporations, to model oil and gas reservoirs and to analyze seismic data.

The technical community is choosing CONVEX systems for their ability to solve computationally demanding problems and for their proven reliability, usability, and affordability.

Scientists, engineers, and senior executives appreciate CONVEX supercomputers for their integrated scalar, vector, and parallel technology. But these users would never invest in a supercomputer for the technology alone. While they understand the potential of new technology, they also demand computer systems that work in a highly efficient, production-oriented environment.

CONVEX systems deliver the processing power for the complex simulations that can give corporations and research facilities the edge over their competition. And CONVEX combines this superior hardware with multiuser, interactive, production-oriented software to make supercomputing accessible and easy to use.

The C Series: The first family of affordable supercomputers.

The C Series family of systems puts supercomputing within the reach of engineers, scientists, and financial experts. It is the first to offer a smooth growth path from a single-processor scalar/vector machine to the high-end, four-processor, tightly coupled parallel supercomputer. Each system delivers an outstanding price/performance ratio.

Those who move from superminicomputers to CONVEX experience a quantum leap. They are suddenly freed from the impediments of minicomputers. They no longer have to scale down and simplify their research.

Instead, they become more creative and undertake more ambitious problems. If the application is mechanical computer-aided design, a C Series system can enable scientists and engineers to go from two- to three-dimensional modeling; if the application is metallurgy, they can go from linear to nonlinear simulations. If the area is cryptology or signal processing, they can discern patterns that were indecipherable before.

With a CONVEX system, users work with more realistic models. They gain new insight; they envision new possibilities and new ways of testing them. In short, they free their imagination from the mental handcuffs imposed by the limited capabilities of superminicomputers.

C Series systems combine high performance with accessibility, making them the ideal supercomputers for large production environments.

C Series hardware breakthroughs.

The C Series machines constitute a major breakthrough for large production-oriented facilities in which many jobs must share the same computing resources. Each machine in the C Series addresses the needs of the scientific and engineering communities with features such as the following.

- 64-bit integrated scalar/vector processors execute concurrently, guaranteeing fast turnaround for a variety of applications.
- Large physical memory, expandable up to 1 or 2 Gbytes (depending on the system), allows applications to fit in physical memory, reducing time to solution for large simulations.
- Large virtual memory permits easy processing of larger user applications.
- A multiuser, interactive operating system makes CONVEX systems accessible in large computing environments.
- An intelligent input/output subsystem moves high volumes of data between peripherals and main memory at speeds of up to 200 Mbytes per second. Data transfers occur at high speeds, leaving the central processing units free for user applications. Standard interfaces for both Multibus® and VMEbus are supported. In addition, a High-Speed Parallel interface (HSP) is supported for custom applications that require extremely fast transfer rates.
- Cost of ownership is extremely low because all C Series systems are air-cooled. There is no need to invest in exotic cooling systems that are expensive to install and maintain. Dense, compact packaging helps to conserve floor space in the computer room, laboratory, mobile research facility, or wherever the system is used.

In addition to the features that all C Series systems share, the parallel processing models have achieved even higher performance. CONVEX has taken parallel technology from the realm of research and experiment to the challenging environment of industrial production. The C Series systems consist of several components that make them the best parallel supercomputers for production environments.

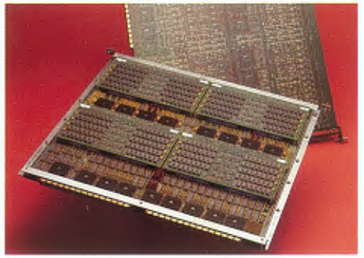

The non-blocking memory crossbar subsystem delivers outstanding performance for large-scale simulations.

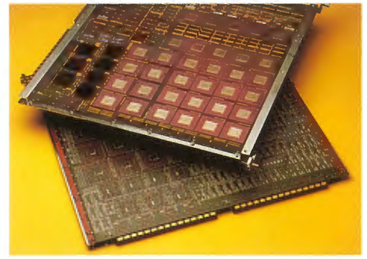

High-speed ECL ASICs and high-density CMOS ASICs boost the performance of C Series supercomputers.

- Automatic Self-Allocating Processors allow each processor in the system to be available continuously. No CPU is idle as long as there is processing to be performed. This is vastly superior to the commonly used static allocation approach. Static allocation assigns processors in a parallel complex to execute a single job. While that job is executing, no other process can access those statically allocated processors. This leaves processors idle during serial portions of the program and wastes expensive CPU cycles.
- CONVEX compilers recognize constructs within programs that lend themselves to vectorization and parallelization, and generate the proper machine level vector and parallel instructions. The CONVEX parallel processors use the parallel instructions to carry out two functions of parallelism:
  - informing other processors that they may assist in the execution of a program, and
  - funneling the portions of code that executed on different processors into a serial stream.

Both of these functions are handled in CONVEX high-speed hardware. As a result, the overhead to split a portion of a job to run on other processors is extremely low, and the operating system is relieved of the time-consuming chore of managing the parallelism. Using hardware rather than the operating system represents a major breakthrough in parallel technology. In addition, the CONVEX method of parallelism is completely automatic and transparent to the user. Users are free from the tedium of deciding which portions of a program to run in parallel: the system does all the work.

- A non-blocking memory crossbar ensures each CPU access to the full bandwidth it requires. This memory can serve up to four memory ports simultaneously at 200 Mbytes per second, providing a total bandwidth of 800 Mbytes per second. Since each CPU in the system has its own 200-Mbyte-per-second path to memory, all CPUs can access memory simultaneously. Memory access is not a bottleneck, so the CPUs can operate at their full potential. This makes the C Series systems the perfect choice for multiuser, multiapplication environments.

• Each CPU is a supercomputer—ECL technology in each individual CPU dramatically raises the baseline performance of the system. This is the speed at which scalar, or serial, or unoptimized codes execute. CONVEX has raised the baseline performance for individual processors by making each one a supercomputer. Higher baseline performance means that all types of codes run faster, not just the codes that can be parallelized or vectorized. Raising the average speed at which all applications run benefits multiuser, multiapplication facilities where high throughput is essential.

Add to this high-performance scalar processor vector and parallel capabilities, and you have the best of all possible worlds. A few strong CPUs operating together is much more efficient than a multitude of weak CPUs. The CONVEX approach to parallel processing is both fast and efficient, offering businesses the chance to beat the competition and save money at the same time.

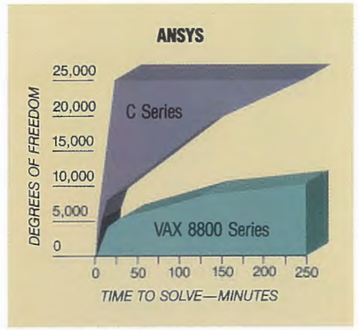

C Series systems' performance on ANSYS® delivers tremendous time savings and allows more complex calculations.

The CONVEX C Series.

C120
The C120, CONVEX's entry-level supercomputer, offers integrated vector and scalar processing to achieve up to 40 megaflops and 11 MIPs of processing power. This range approaches large-scale supercomputer performance. The processor unit is a high-performance, 64-bit, integrated scalar/vector supercomputer. The processor is coupled with main system memory via its own 80-Mbyte-per-second memory bus. Memory in the C120 can be expanded up to 1 Gbyte, which makes processing even the largest applications easy. Both the highly pipelined scalar and vector units can execute concurrently. These units, with an instruction dispatch and address translation unit, form a pipeline that can execute up to eight instructions simultaneously.

The C120 is air-cooled and densely packaged to fit conveniently into a variety of environments. Its cost of ownership is surprisingly low for the computing power it delivers.

C201-C202
The C201 and C202 are midrange, expandable processors that provide a low-cost entry into the CONVEX line of parallel processing systems. The base system (C201) can be expanded from one to two processors, and upgraded to other members of the C Series. When configured with two processors (C202), the system takes advantage of CONVEX's Automatic Self-Allocating Processors parallel technology. The C201 achieves performance of 23 MIPs and 36 megaflops, and the C202 achieves 46 MIPs and 72 megaflops.

The processor unit is a high-performance, 64-bit integrated scalar/vector supercomputer. The processor is coupled with main system memory via its own 145-Mbyte-per-second memory bus. A non-blocking, dual-ported memory subsystem expandable up to two Gbytes allows applications to fit in physical memory, reducing the time to solution for large simulations.

Like the C120, the C201 and C202 are air-cooled and compact to save on operational costs.

C210
The C210 is a high-performance uniprocessor that extends the performance growth path of the C Series. The processor unit is a high-performance, 64-bit scalar/vector supercomputer that is tightly coupled with main system memory via its own 200-Mbyte-per-second bus. The highly pipelined scalar and vector units execute concurrently and can execute up to eight instructions simultaneously.

The C210's increased scalar performance ensures consistently high performance for all types of applications.

The C210 delivers performance of up to 50 megaflops and 26 MIPs. Its clock speed is two and a half times faster than the C120's, and its memory system is expandable up to 2 Gbytes. The non-blocking crossbar prevents the memory system from becoming a bottleneck. Large virtual memory enables the C210 to process even the most complex user applications with ease.

Cost of ownership is surprisingly low. Compact packaging, air cooling, and low power requirements make it very affordable to place and maintain.

C220
The C220 is a dual processor, 100 megaflops, 50 MIPs scalar/vector/parallel supercomputer. It uses Automatic Self-Allocating Processor technology to run multiple jobs quickly and efficiently while boosting overall performance for a single job. Single jobs can be split to run concurrently on both processors, then revert to one when scalar code is being processed, then back to two, and so on. This allocation takes place automatically in the hardware, with no programmer intervention or operating system overhead.

Compact and air-cooled, the C220 has a remarkably low cost of ownership. It has large physical memory (up to 2 Gbytes) and large virtual memory to handle even the largest production jobs.

C230
The C230 is a three-processor system capable of performance levels of up to 150 megaflops and over 75 MIPs. Like other machines in the C Series, the C230 achieves supercomputing performance by fully integrating scalar, vector, and parallel processing. It uses the Automatic Self-Allocating Processor technology to run multiple jobs quickly and efficiently while boosting overall performance for a single job.

Densely packaged and air-cooled, the C230 has a remarkably low cost of ownership in contrast to traditional supercomputers with similar performance levels. The C230 has large physical memory (up to 2 Gbytes) to handle even the most complex simulations. The non-blocking crossbar prevents memory system bottlenecks, and large virtual memory allows processing of large user applications.

C240
The C240, with four processors, is the most powerful system in the C Series. It achieves up to 200 megaflops of processing power and over 100 MIPs.

As with other machines in the C Series, the C240 uses scalar, vector, and parallel processing to achieve its supercomputing power. It uses the Automatic Self-Allocating Processor technology to run multiple jobs quickly and efficiently while boosting overall performance for a single job. With the C240, CONVEX achieves and exceeds traditional supercomputer performance in a machine designed for production environments.

The C240 has large physical memory (up to 2 Gbytes) to store the vast amounts of data required for detailed simulations. It uses the non-blocking crossbar technology for high memory bandwidth and has large virtual memory to process large user applications.

Dense packaging and air cooling maintain the low cost of ownership, even with this high-end C Series system.

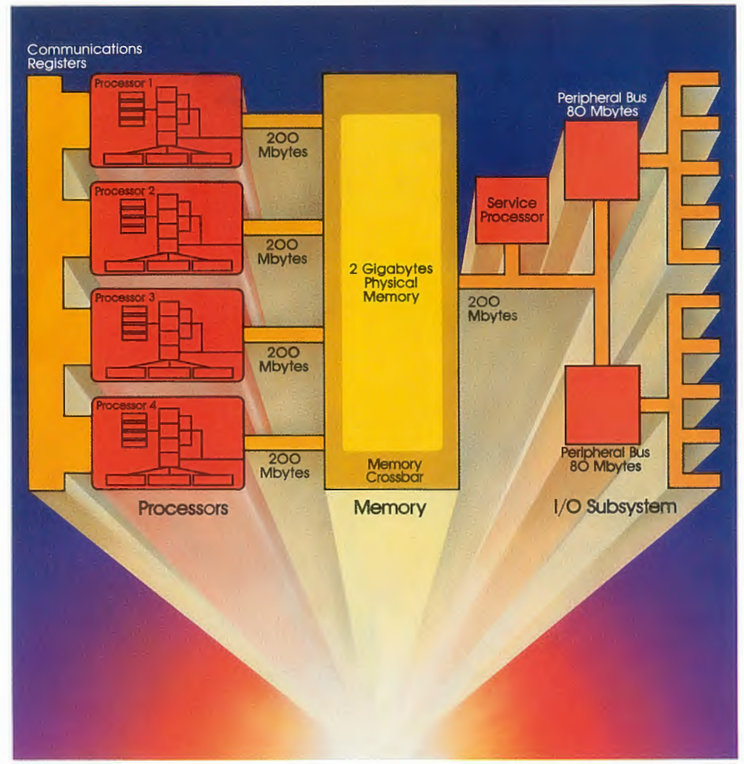

The CONVEX C201-C240 architecture combines supercomputer class CPUs, high I/O performance, and an 800-Mbyte-per-second memory bandwidth to provide a balanced solution to production supercomputing.

CONVEX software: A series of advances.

A production supercomputer begins with the architecture. It also requires a full complement of software: advanced compilers that benefit from the system architecture; an operating system suited for the environment; programming, system management, and networking utilities for accessibility; and a collection of the most popular industry-standard third-party packages to make the system immediately productive. CONVEX systems provide all this and add parallel processing capabilities to deliver results even faster.

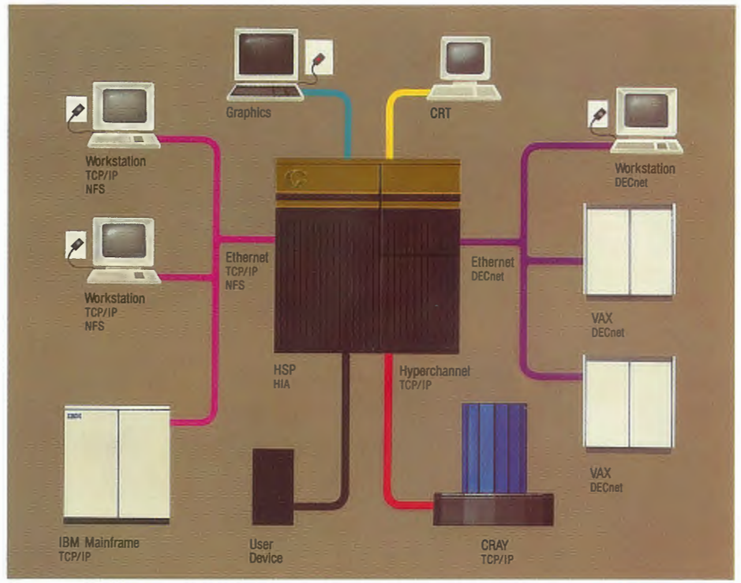

CONVEX networking products, including COVUEnet, allow CONVEX systems to perform well in heterogeneous environments.

Leading-edge compiler technology.
CONVEX already offers the world's finest vectorizing, optimizing compilers. Now, CONVEX has advanced its compiler technology by adding automatic parallel processing capabilities.

CONVEX compilers generate an internal dependency analysis that gives a global view of variable dependencies and flow of control within a program. Using dependency analysis, the compilers automatically determine opportunities for vectorization, parallelization, or both. There is no need for programmer intervention to take advantage of vector or parallel processing; the compilers do the work.

The last step of compilation produces highly efficient object code that exploits the machine architecture, including parallel execution of instructions and parallel operations on data structures. CONVEX compilers determine which instructions may overlap during execution on an individual processor and schedule the resulting machine instructions accordingly. The result is the maximum possible throughput for a given set of instructions.

CONVEX compilers also make it easy to transport codes from other systems. The compilers not only accept industry-standard languages (FORTRAN, C, and Ada) but are compatible with the most popular dialects of each of those languages.

For CONVEX FORTRAN, that means compatibility with FORTRAN 77 and VAX™ FORTRAN. For CONVEX C, that means compatibility with the portable C compiler (PCC). For Ada, that means an Ada that is validated by the Ada Joint Program Office. No unusual syntax or extensions to any of these languages is required for vectorization or parallelization.

CONVEX UNIX extended for production supercomputing.
Traditionally, supercomputers offered only a minimal operating system environment with few user services. CONVEX realizes that users require more control and flexibility in their work than these impoverished operating systems allow. The CONVEX UNIX® operating system is rich with user services without sacrificing any of the system's supercomputing power. CONVEX UNIX is based on the industry-standard UNIX 4.2BSD and is extended to support large, sophisticated supercomputing environments.

Extensions range from performance boosters that deliver consistently high performance to resource management utilities that make the system easier to control and manage.

In addition, CONVEX systems come with the UNIX utilities necessary for software development: debuggers for both source and object code, profilers, text editors, and libraries of mathematical routines.

CONVEX also offers several networking options with all C Series supercomputers so that users can connect CONVEX systems to each other and to others in their environment. Among the industry-standard network interfaces supported by CONVEX are the following.

- Ethernet™ and Terminal Connect Protocol/Internet Protocol (TCP/IP) allow connections to other UNIX-based systems, non-UNIX systems, and workstations.
- CONVEX Network File System (NFS™) allows users to share a common base of files across different systems. This provides a time savings because there's no need to copy files from one system to another to get work done.
- CONVEX systems can be integrated into DECnet™ networks with the COVUEnet software package.
- HYPERchannel™ provides a link between CONVEX systems and large-scale supercomputers and mainframes that support TCP/IP.

Third-party applications you depend on.
CONVEX offers more than 200 premier third-party application programs in these areas:

- structural analysis,
- computational fluid dynamics,
- petroleum,
- chemistry,
- simulation/operations research, and
- graphics.

In addition, CONVEX has a robust collection of application libraries, such as VECLIB™, that maximize program performance with a minimum of programming effort.

COVUE makes affordable supercomputing even more accessible.
CONVEX computers are more than high-performance machines. They are more than state-of-the-art technology. They are also highly accessible.

For VAX users, CONVEX offers the COVUE™ (CONVEX-to-VAX User Environment) family of software products to ease the transition from the VMS environment to the UNIX environment.

COVUEshell is a VMS-like shell that emulates most of the commonly used DCL commands. Many first-time CONVEX users have found COVUEshell so easy to use that they have logged in and run an application within 5 minutes.

COVUEnet makes it easy to network a CONVEX C Series machine into a VAX environment. With COVUEnet, your CONVEX system becomes a powerful node on a DECnet Phase IV network.

COVUEbatch allows a CONVEX system to be used as a computational server. Users can submit jobs on the VAX and receive output on the VAX, using the CONVEX to run computationally intensive applications and the VAX for time sharing.

Built-in reliability and worldwide support.

CONVEX systems are well known for their reliability. The production facility in Richardson, Texas employs state-of-the-art manufacturing and testing techniques and a comprehensive quality assurance program that detects and corrects problems before they leave the plant.

CONVEX systems average 2600 hours of continuous uptime in all types of environments. And commitment to our customers doesn't end with the quality we build into each CONVEX system. It extends from our factory, to our field service offices throughout the world, to our customers' sites.

CONVEX has structured a support services team to provide the excellent service our customers have come to expect from us. This team includes support for hardware and software with the following services.

- Field support experts worldwide perform installations and maintenance.
- A dial-in hotline to the Technical Assistance Center (TAC) provides advice and quick fixes to keep your systems up and running.
- Problem reporting service via electronic mail automatically tracks problems until they are resolved.
- CONVEX maintains a multimillion-dollar inventory of spare parts placed strategically around the world to ensure coverage of every customer site. This inventory is constantly evaluated by CONVEX Logistics Management to provide fast, cost-effective support.
- Sophisticated remote diagnostics capabilities provide quick problem resolution from CONVEX's home office.
- The CONVEX training center provides classroom and laboratory classes to help you get the most out of your CONVEX system.

With a price/performance ratio no other manufacturer can touch and accessibility for large production computing environments, it's no wonder that so many scientists and engineers require the standard of affordable computing excellence: The CONVEX C Series.

The CONVEX hotline puts customers in touch with product specialists in the Technical Assistance Center (TAC).

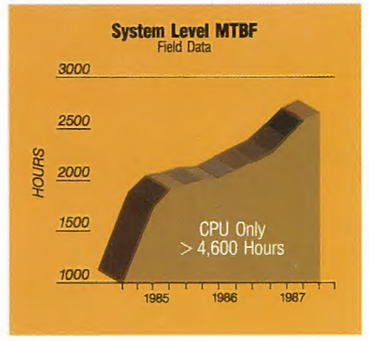

Current mean time between failure (MTBF) for CONVEX systems is more than 2600 hours. This figure is derived from actual field data, not laboratory testing.

Corporate Headquarters:
CONVEX Computer Corporation
P.O. Box 833851
701 Plano Road
Richardson, TX 75083-3851
Phone: 214 952-0200
FAX: 214 952-0550
Telex: 5106002063

European Headquarters:
CONVEX Computer Corporation
Randalls Research Park
Randalls Way
Leatherhead
Surrey, UK KT22 7TS
Phone: 011-44-372-386696
FAX: 011-44-372-375877
Telex: 914827

Asia/Pacific Headquarters:
CONVEX Computer Corporation
1 Scotts Road
#25-06 Shaw Centre
Singapore 0922
Phone: 65-733-4355
FAX: 65-733-4354

080-001006-001
Printed in U.S.A.