intel

iPSC 860

SPECIFICS

The iPSC®/860 is an ultra-high-performance computer system designed to operate in business, manufacturing, and research environments. The system's distributed-memory multicomputer architecture, and use of standard microcomputer components produces a scalable system that offers both high performance and high economy. Peak-rated at 7.6 billion double-precision operations per second and configurable with up to 8 GBytes of system memory, the iPSC/860 exceeds the floating point performance of any conventional vector supercomputer. Priced at less than $450 per MFLOPS, the iPSC/860 is also ten times more cost-effective than traditional supercomputers. Life-cycle costs are low as well, since iPSC/860 systems are air-cooled, compact and reliable.

The iPSC/860 system comes with a comprehensive software suite for developing parallel applications and adapting existing programs to take advantage of Intel's parallel supercomputer architecture. The iPSC/860's UNIX workstation-based cross-development environment includes performance-optimizing compilers for industry-standard implementations of the FORTRAN and C languages. Supplementing the compilers are a parallel-oriented CASE development environment, debugger, performance analyzer, and an iPSC/860 simulator. High-performance mathematical libraries and matrix solvers are also available.

For easy integration into networked computing environments, the iPSC/860 can connect to other computer systems through one or more high-performance Ethernet channels. Support for network standards, such as the X-Window System, the Network File System (NFS), and the FTP and TCP/IP protocols, foster integration into VAX/VMS, UNIX, and IBM environments.

LEVERAGING STANDARDS FOR LOW COST SUPERCOMPUTER PERFORMANCE

The underlying technologies that implement the iPSC/860 architecture account for its exceptional performance and cost-effectiveness. The iPSC/860 is based on standard commercial microprocessors, memories, and disks, and standards-based software and I/O interfaces. Intel's ten-year investment in parallel systems technologies transforms these capabilities into an integrated multi-function supercomputer.

The iPSC/860's use of commodity CMOS memory components and 5 1/4'' disks leverages the enormous mass market R&D investment in PC and workstation technology. This ensures maximum cost-effectiveness and also puts future product developments on a predictable course for further improvements in capacity and speed. As the inventor and leading manufacturer of high-performance microprocessors, Intel is now actively exploiting this technology for technical computing applications.

The iPSC/860 exceeds the floating point performance of conventional supercomputers at one-tenth their cost.

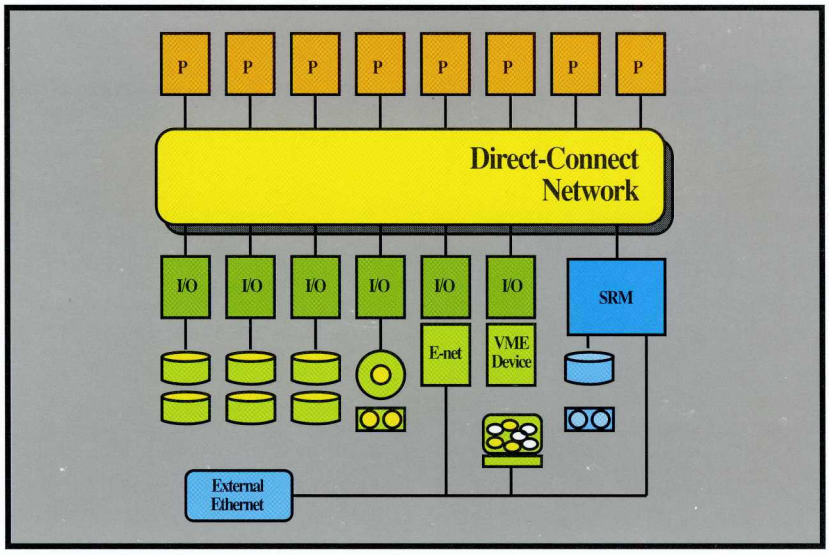

Figure 1: *The iPSC/860's distributed memory, message-passing architecture provides exceptional parallel efficiency and performance.*

Along with its use of standard commercial microprocessor components, the iPSC/860 implements a wide range of industry standards, thus increasing interoperability and further contributing to lower system costs. iPSC/860 standards range from the VME and SCSI I/O interface standards through software standards such as the UNIX operating system, TCP/IP networking services, and NFS file access services.

iPSC/860 HARDWARE OVERVIEW

The iPSC/860 is based on an advanced, large-scale parallel multicomputer architecture using independent computing elements that collectively process information in parallel. The computing elements, called nodes, are each a fully functional computer system containing a CPU, local memory, and I/O facilities.

Communication among nodes is a critical performance element for parallel processing computers. iPSC/860 nodes communicate via Intel's proprietary, high-speed Direct-Connect™ internal network, passing messages that contain instructions or data among the nodes. This internal network is inherently scalable and allows communications capacity to grow proportionately with the number of processors.

This simple but powerful architecture, with its distributed processors, on-board memory and message-passing communications, produces several significant benefits:

- performance: because the internal network can be expanded to accommodate a large number of processors;
- economy: because individual nodes can be constructed from low-cost CMOS components; and
- flexibility: because of the the system's inherent modularity, scalability, and ability to support varied programming models.

The iPSC/860 provides exceptional parallel efficiency up to the maximum number of processors since resource bottlenecks are minimized. By contrast, parallel computers that use shared memory architectures must contend with a single scarce resource: memory bandwidth. This contention limits the effective use of parallel processing to eight or fewer high-performance CPUs. Facing this limitation, manufacturers must resort to extremely expensive components based on ECL or gallium arsenide semiconductor technologies to achieve high system performance. These tradeoffs severely limit the cost-effectiveness of such systems.

Intel's multicomputer architecture produces a system that is inherently upgradable. Customers can make a modest investment in supercomputing, then scale up their system as their needs grow. iPSC/860 systems are scalable from 8 processors to 128 processors, with peak performance ranging from 480 MFLOPS for an 8-processor system to 7.6 GFLOPS for a system with 128 processors. Memory is scalable as well. Each processing node can have 8 to 64 MBytes, providing total system memory ranging from 64 MBytes to 8 GBytes.

The iPSC/860's distributed memory, message-passing architecture also provides the flexibility to mix node types in a single system and to dedicate nodes to specific tasks such as processing or I/O.

Finally, Intel's supercomputer architecture supports a Multiple Instruction Multiple Data (MIMD) computational model that stands in contrast to the Single Instruction Multiple Data (SIMD) model used by several of today's massively parallel supercomputers. The SIMD model severely restricts programmers by requiring that processors execute identical programs in lockstep fashion. The iPSC/860's MIMD model, on the other hand, gives programmers the flexibility to tackle problems with complex boundary conditions, for example, or problems that have many subprograms of differing computational complexity.

NUMERIC PROCESSING NODES

The iPSC/860 performs its numeric computations on specialized processing nodes. Each numeric node has is own CPU, kernel operating system, local memory, and interface to the Direct-Connect internal network.

The CPU is based on a 40 MHz Intel i860 microprocessor, the world's first supercomputer-on-a-chip. Each i860 processor contains a 64-bit RISC processor, a floating-point adder, a floating-point multiplier, and 8 and 16 Kbyte caches for instructions and data. Peak-rated at 60 MFLOPS double precision and 80 MFLOPS single precision, a single i860 microprocessor actually exceeds the scalar processing performance of a Cray 1 supercomputer. Intel's single-chip, million-transistor implementation of the i860 ensures computational

860 nodes and to the System Resource Manager disk files. Concurrent I/O disk files can be accessed directly via the File Transfer Protocol (FTP).

For interactive processing, the iPSC/860 offers the graphical interface of the industry-standard X-Window System, UNIX sockets, and the iPSC/860's Remote Host facilities. X Window System client services allow iPSC/860 nodes to communicate with distributed workstations over the network. The UNIX Sockets interface can be used to create custom process-to-process communications. Remote Host is a network communications facility that extends the iPSC/860's System Resource Manager user and application interface onto Sun workstations. Sun-to-iPSC/860 communications are made transparent to the user, allowing users and developers to avoid the distraction of remote logins. Remote Host allows applications running on the workstation to communicate with the iPSC/860 via Intel's message passing, as if the workstation were another node in the system.

The iPSC/860's capabilities as a network server include batch-style processing using the Network Queuing System (NQS), an industry-standard batch manager for UNIX client/server environments. With NQS, iPSC/860 users can submit jobs in the form of UNIX shell scripts or command files. Multiple job queues and multiple priorities enable a mix of small tasks to run in small partitions of the system while large tasks are dedicated to a partition of many numeric nodes. NQS also lets users direct input and output to remote printers, file servers and workstations on the network.

LANGUAGES AND SOFTWARE DEVELOPMENT TOOLS

For software developers, Intel provides a rich environment of cross-development tools that run on UNIX workstations via network connection to the iPSC/860. FORGE® and CAST™, state-of-the-art CASE tools for parallel programming, are available as aids to porting and managing parallel FORTRAN applications. FORGE offers powerful facilities for analyzing the structure and data flow of FORTRAN programs, while CAST provides interactive structure editors and FORTRAN code generators that transform applications into parallel programs ready for compilation and execution. Graphical performance analysis tools are available for performance tuning of parallel applications.

For outstanding run-time performance, iPSC/860 C and iPSC/860 FORTRAN offer sophisticated optimization techniques that ensure effective pipeline utilization and efficient cache management. To simplify porting and programmer training, the compilers conform to industry standards. iPSC/860 C conforms to the proposed ANSI C standard (X3JII), while iPSC/860 FORTRAN is a superset of FORTRAN77 that includes VAX/VMS 4.0 and IBM/VS FORTRAN extensions. The compilers also have language-compatible libraries for iPSC/860 message-passing and file system services, as well as both IEEE-compliant floating point run-time libraries. Ada for the iPSC/860 is under development.

LIBRARIES AND SOLVERS

The iPSC/860 ProSolver family of matrix solvers has been specifically designed and tuned for the iPSC/860. The family includes the ProSolver-DES direct equation solver, the ProSolver-IES iterative equation solver, and the ProSolver-SES skyline equation solver. ProSolver matrix solvers can be used as the starting point for new applications

iPSC 860 System Specifications

Number of CPUs	8	16	32	64	128
Peak Performance (GFLOPS)					
Double Precision	.48	.96	1.9	3.8	7.6
Single Precision	.64	1.3	2.6	5.1	10
Memory Capacity (MBytes)					
Standard	64	128	256	512	1,024
Maximum	12	1,024	2,048	4,096	8,192
Storage (GBytes)					
Standard	.65	1.3	1.9	4.5	9.1
Maximum	9.1	20	40	82	165
Internal Network (MBytes/sec.) (Bisection Bandwidth)	22.4	44.8	89.6	179.2	358.4
External I/O					
Available Channels	7	15	31	63	127
Maximum I/O (MBytes/sec.)	19.6	42.0	86.8	176.4	355.6
Physical Size (feet)					
Width	1.8	1.8	3.6	3.6	5.3
Height 5.0, Depth 2.1					
Operating Temperature (°C)	10-30	10-30	10-30	10-30	10-30

density and exceptional performance.

Complementing the i860 microprocessor, numeric nodes also contain from 8 to 64 MBytes of on-board memory. This combination of high-performance CPU with large local memory means each node can support a wide selection of programming models, algorithms, and tools. This approach also simplifies the program development process, allowing programmers to run and test programs on a single node before parallellizing the program. Other parallel supercomputer systems lack this basic facility, forcing developers to effect a complete and flawless parallel implementation in a single step.

INPUT/OUTPUT NODES

In keeping with its architectural philosophy, the iPSC/860 uses specialized I/O nodes to provide a scalable I/O facility that can be configured to meet specific application requirements. Based on the Intel 386 microprocessor, I/O nodes can control peripherals, external interfaces and local area networks.

For example, each I/O node can have an industry-standard 4 MBytes/sec SCSI interface. Using 760-MByte 5.25-in disks, disk arrays can be configured with as much as 165 GBytes of storage. Intel's Concurrent File System (CFS) software combines high-performance block interleaving with standard UNIX file interfaces to simplify data management for disk arrays and make parallel data access to large numbers of disks transparent to users.

The iPSC/860's I/O nodes also allow use of other mass storage options such as 8mm cartridge tape for file backup and nine-track reel-to-reel tape drives for data interchange. In addition to the SCSI interface, I/O nodes can provide VME support for various data acquisition boards and custom interfaces.

For networking to other computing resources, the iPSC/860 offers multiple TCP/IP Ethernet channels. Dedicated network nodes perform protocol processing and data buffering; the network nodes also support a UNIX sockets software interface.

DIRECT-CONNECT INTERNAL NETWORK

The iPSC/860's Direct-Connect internal network provides high-speed data pathways between all nodes of an iPSC/860 system and supports literally hundreds of simultaneous processor-to-processor communications with uniform performance between all nodes. Each channel is a bi-directional pathway that delivers data at 5.6 MBytes/second. Automatic switching hardware frees the programmer from the details of message routing, making Direct-Connect easy to use.

The number of available channels increases with system size, maintaining a balance between communications performance and processing performance even as the system grows in size. This expandability is a key distinction between iPSC/860 systems and traditional, bus-based parallel architectures. Because of Intel's Direct-Connect technology, iPSC/860 systems can scale to 128 compute nodes and 127 I/O nodes, while conventional multiprocessor systems are limited to 8 processors.

Along with providing an internal pathway between processors, the Direct-Connect data highway also provides channels from individual processors to the iPSC/860's System Resource Manager (SRM).

SYSTEM RESOURCE MANAGER

The System Resource Manager oversees system allocation and system diagnostics services. The SRM is responsible for allocating and deallocating portions of the iPSC/860 system to various users. As the iPSC/860's diagnostic console, the SRM also contains a comprehensive set of system test routines that assist in isolating faults.

In addition to these primary functions, the SRM also supports tasks such as file storage, network gateway services, and a UNIX System V software development environment. For maximum system throughput, these functions are ordinarily performed by I/O nodes, while the software development tools run on Sun workstations connected to a shared Ethernet network.

PACKAGING AND ENVIRONMENTAL ISSUES

The iPSC/860 uses a modular, air-cooled design that keeps expansion simple and facility costs low. Even the largest standard system configuration, with 128 numeric nodes and a complementary I/O facility, is maintained in a normal computer-room environment and occupies less than 12 square feet of floor space. iPSC/860 systems run on standard AC power at 208 volts and less than 30 Amps per cabinet. And since iPSC/860 systems are air cooled, there's no need for the specialized plumbing and cooling that can drive up the operating costs of traditional supercomputers.

iPSC/860 SOFTWARE OVERVIEW

Intel offers a comprehensive software system optimized for the iPSC/860's multicomputer architecture and designed for efficient usage in a distributed computing environment.

For software development, the iPSC/860's workstation-based cross-development tools allow programmers to remain in their familiar program development environment and keep the supercomputer's processing power dedicated to user applications. Intel offers mathematical libraries and parallel solvers that reduce programming effort and parallel CASE tools and a parallel debugger that aids in obtaining high run-time performance.

System management functions allow users to share the computer. Typically, users prepare a program or data set at a local workstation, then obtain some or all of the supercomputer's processing nodes for testing or production runs. For example, one user might be granted half of a system's available processors to run a particular job, while four others run concurrently, each on one eighth of the system's processing nodes. Each user's program operates independently and concurrently. Users can use the system interactively or take advantage of the system's UNIX batch facility for deferred execution. Programs and data can reside on the iPSC/860's

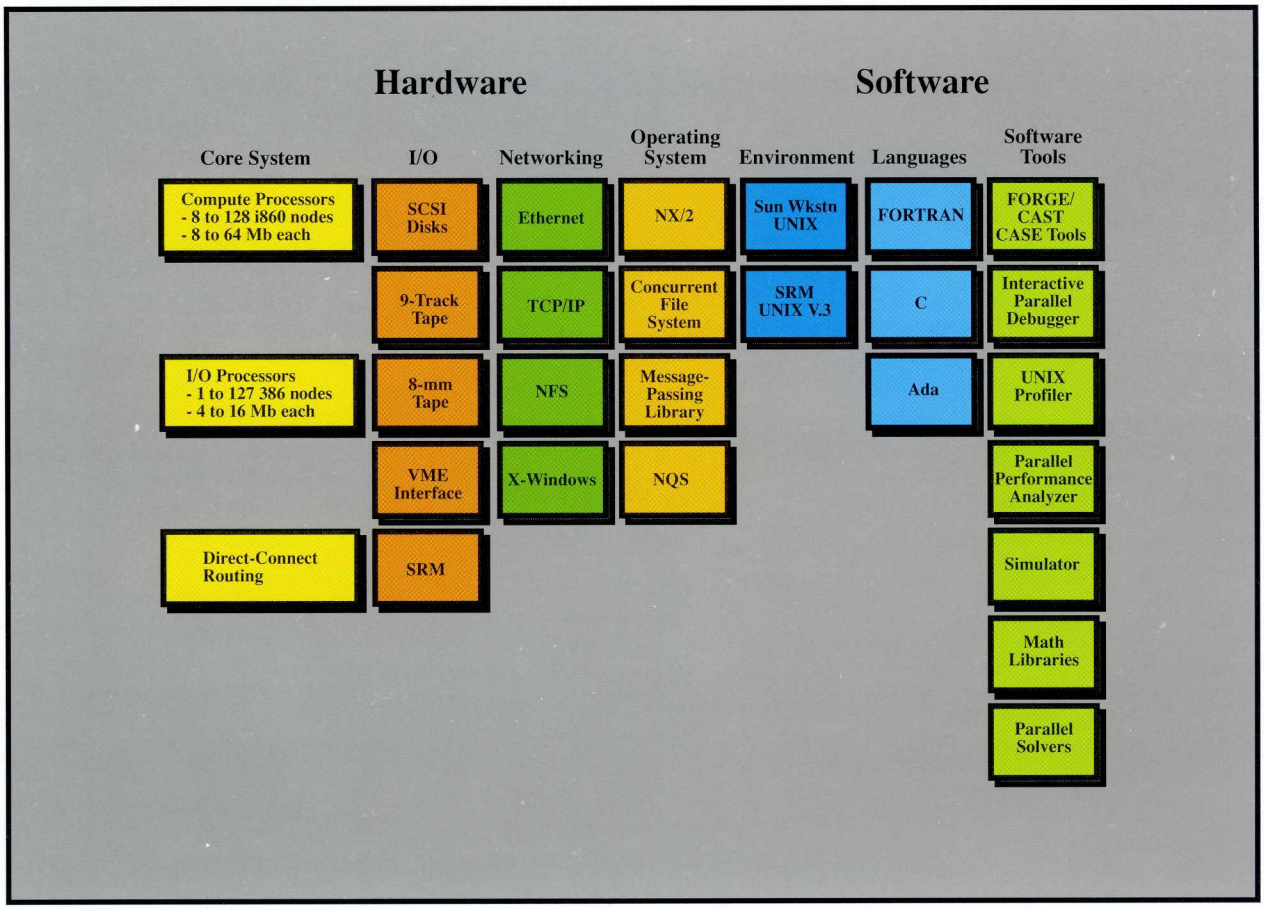

Figure 2: *The iPSC/860's system components create an efficient, easy-to-use multifunction supercomputer designed for performance, economy, and flexibility.*

System Resource Manager, on its Concurrent File System (CFS), on the user's workstation, or at any NFS file server on the network.

iPSC/860 OPERATING SYSTEM

iPSC/860 system software emphasizes performance, UNIX standard application interfaces, and ease of use. Based on a light-weight kernel operating system, called NX/2, the iPSC/860 maintains a consistent run-time environment while offering fast and flexible message passing services. NX/2 resides on each of the compute nodes and I/O nodes in the system and provides synchronous, asynchronous and interrupt-driven communication services with automatic flow control and buffering as well as support for global operations. By maintaining compatibility with UNIX I/O services, NX/2 also simplifies the effort required to port new applications to the iPSC/860.

FILE SERVICES

The iPSC/860's Concurrent File System (CFS) provides a UNIX-compatible file system while hiding the concurrency of multiple I/O nodes and disks. Limited in size only by the aggregate space of all the disks in the system, The CFS can support file sizes of many GBytes without partitioning.

The Concurrent File System executes reads and writes to multiple disks in parallel, for greater performance. The CFS manages parallel I/O through a method that is similar to disk striping but more efficient because it does not require users to synchronize reads or writes. The file system automatically distributes the blocks of individual files across all available disks, using algorithms for reading and writing that allow several independent applications to run simultaneously. Programmers thus gain the performance of parallel disk access while keeping the simplicity of a single, non-partitioned file.

Filing services also support nine-track and 8mm tape drives.

NETWORK SERVICES

The iPSC/860 provides a rich set of features for using the supercomputer in a networked environment. The system uses Ethernet and TCP/IP software as the base for an extensive set of standard networking capabilities. Dedicated I/O nodes with hardware-assisted TCP/IP protocols provide a high-performance, industry-standard connection to a wide variety of systems. Through the use of gateways, TCP/IP can be extended to connect to non-UNIX systems such as DEC VAX/VMS and IBM MVS computers. Multiple I/O nodes can expand network connections and throughput while keeping networking tasks from interfering with applications running on the iPSC/860.

In addition to TCP/IP networking, a full complement of standards-based applications-level software is available for file access, interactive processing, and batch processing on the network. The Network File System (NFS) is provided to iPSC/

development, or they can be incorporated into existing codes as part of the parallel conversion process. The result is quick time to market and maximum applications performance.

Intel also offers the widely used BLAS and NAG math libraries.

INTEL SUPPORT SERVICES

To help customers fully exploit the power of their iPSC/860 supercomputer, Intel provides extensive technical support and consultation, including end-user training, applications consulting, benchmarking, systems integration support, and porting services. Intel also offers a comprehensive system maintenance program supported with trained technicians in each of Intel's 120 offices worldwide.

Intel Corporation
Supercomputer Systems Division
15201 NW Greenbrier Pkwy
Beaverton, OR 97006
(503) 629-7600