iSCurrents

Published by Intel Scientific Computers

Fall/Winter 1987

NEKTON Announced for iPSC®

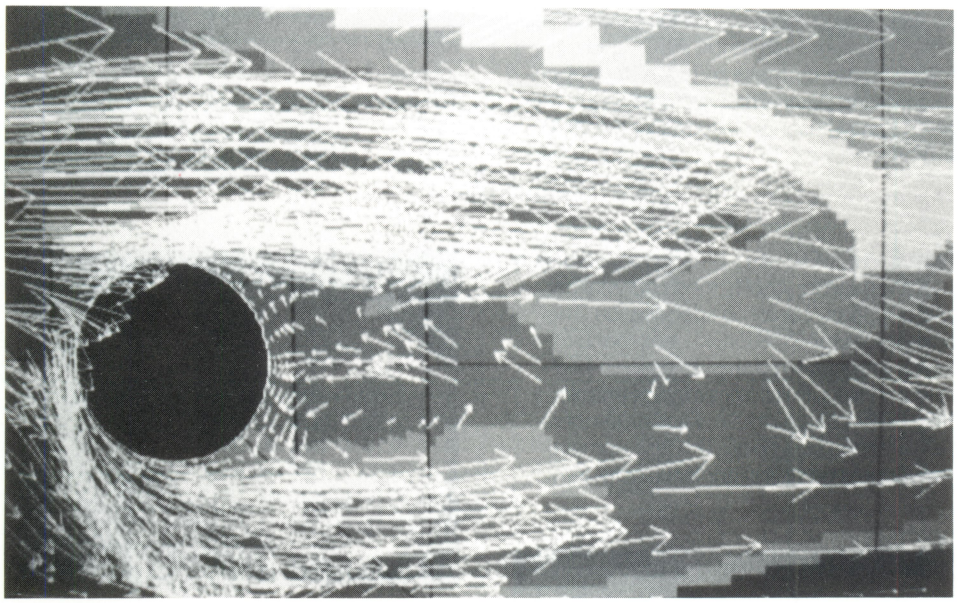

This display shows the velocity and temperature fields in the vicinity of an Eddy-Promoter Heat Exchanger for an unsteady flow, as calculated by the NEKTON fluid dynamics and heat transfer numerical simulation package.

The first full-scale production software package for the iPSC® has just been announced by Nektonics Inc., of Bedford, MA. The *NEKTON* package is a state-of-the-art fluid dynamics and heat transfer simulation code that includes the most recent advances in numerical methods and algorithmic techniques. It features 2D and 3D modeling capabilities, high computational efficiency, and accurate modeling further into transitional regions than competitive packages.

The NEKTON package is already implemented and widely used on a variety of computational platforms, ranging from MicroVAX and Sun workstations to Cray supercomputers. The package has been used to study thermal management of electronic components, crystal growth, contaminant dispersion, and many other fluid flow and heat transfer processes.

Now, NEKTONICS has announced a beta release of NEKTON for the iPSC-VX Vector Concurrent Supercomputer, with production shipments to begin early in 1988. "The iPSC-VX architecture is ideally suited to the computational methods used by NEKTON," said Dr. Anthony Patera, chairman and co-founder of Nektonics. "The price-performance advantage of this system makes it possible to put a dedicated, supercomputer-class simulation engine into the hands of a small engineering team."

The first public announcement of the iPSC-VX version of NEKTON came at the August meeting of the American Society of Mechanical Engineers in New York City.

The performance of the iPSC-VX / NEKTON system makes it a highly cost-competitive alternative to traditional, sequential approaches. "NEKTON on the iPSC-VX system can dramatically improve engineering productivity," said Elliot Swan, Market Development Manager for Intel Scientific Computers. A NEKTON simulation that takes *four hours* on a VAX 8600 system takes just *three minutes* to complete on a comparably priced, 32-node iPSC-VX. "When you can do a half-day of simulation in the time it takes to get a cup of coffee, there will be significant changes in the way engineers do their work," added Swan.

NEKTON performance on the iPSC-VX is roughly equivalent to a CRAY X-MP costing 10 times as much. "With this level of computing power available *locally* and *dedicated to an engineering team,* simulations can be more frequent, more finely tuned, and consequently of higher quality," continued Swan. "And with the iPSC-VX as a shared engineering resource, simulations formerly done on

continued on page 2

In this Issue of iSC

iSCurrents

Volume 2 Fall/Winter 1987

iSCurrents is published quarterly (Spring, Summer, Fall, Winter) by:

Intel Scientific Computers
15201 N.W. Greenbrier Parkway
Beaverton, OR 97006
503-629-7629

Nektonics, Inc.:

Planning for Concurrency

The primary goal of Nektonics, Inc. is to provide the engineering and scientific communities with *leading edge floware* — proven, cost-effective tools for solving a wide range of problems in fluid dynamics and heat transfer. NEKTON is Nektonics' first commercially available package. The company is already working on major enhancements for release early in 1988 (see related article), as well as customer support and training strategies.

Nektonics is one of an emerging class of software developers who recognize the importance and benefits of parallel computing. Because of the price/performance of parallel computers such as the iPSC® and iPSC®/2 systems, these software developers are targeting parallel architectures as the preferred computational platforms for the 1990s and beyond.

Nektonics, Inc. is headquartered in Bedford, MA, with offices in Cambridge, MA and Princeton, NJ. The company was founded in 1985 by a group of engineering and mathematics faculty from M.I.T. and Princeton University. The principal co-founders are Dr. Anthony Patera (M.I.T.) and Dr. Steven Orszag (Princeton), both leading researchers in computational fluid dynamics.

Today, Drs. Patera and Orszag serve as Chairman and Vice-Chairman of the Board of Nektonics, respectively, and continue to provide technical consultation as needed. Dr. Patera's research expertise is in heat transfer, spectral element methods, and parallel processing. Dr. Orszag is the founding principal investigator of the NSF-funded John von Neumann Supercomputer Center; his research expertise is in turbulence modeling, spectral methods, and supercomputing.

NEKTON Announced for iPSC®

continued from page 1

workstations can now be done on the hypercube, with greater speed and accuracy."

NEKTON is a new code, written in ANSI standard FORTRAN 77. While the computational approach is based on fundamental research supported by ONR, AFOSR, DARPA, and NSF, subsequent code development was funded by Nektonics.

The base system was written for a virtual message-passing parallel processor with nested vectorization, making it an excellent fit for the iPSC-VX. Nektonics, working closely with iSC, has now tailored the code to take full advantage of the features of the iPSC. The customization included development of a specialized matrix multiply operation that achieves 8 to 15 MFLOPS (million floating point operations per second) per processor.

The NEKTON code solves the full incompressible, unsteady Navier-Stokes and energy equations. It is particularly well suited for detailed calculations of flows in the laminar and transitional regimes, for both 2D and 3D problems. The code is built around the *spectral element method* (see related article) and provides more accurate results than were previously attainable for transitional regions.

The package includes graphics oriented pre- and post-processors that run on a workstation attached to the iPSC. The pre-processor, PRENEK, uses a menu/mouse-driven user interface to simplify mesh generation and boundary condition specification. The post-processor, POSTNEK, uses menu-driven interactive graphics to assist in interpreting results. Full-color graphics enable clear presentation of flow, velocity, temperature, and pressure fields. The flow domain may be viewed at any angle; velocity, temperature, and pressure can be examined on arbitrary planes in the flow.

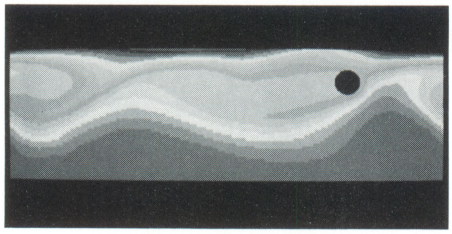

Isotherms for an Eddy-Promoter in an unsteady flow.

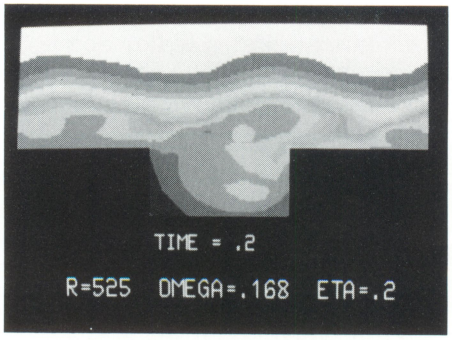

Isotherms for pulsed flow in a grooved channel.

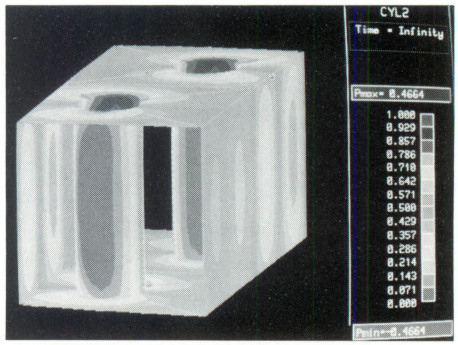

Stokes flow past a staggered cylinder array in a three-dimensional duct (calculations performed on an iPSC-VX hypercube).

Two major enhancements are scheduled for release in early 1988: simulation of free surface phenomena, and a turbulence model developed by Dr. Steven Orszag of Princeton. Models of turbulent regions have not been commercially available heretofore. In addition, a release of NEKTON for the Intel iPSC®/2 Concurrent Supercomputer is expected early in 1988.

The NEKTON product license includes support by phone, mail, TELEX, and fax. More extensive support and consultation is available at additional cost. Additional information can be obtained from Nektonics, Inc. PO Box 22, Bedford, MA 01730, (617) 275-4011, or from Intel Scientific Computers.

The Spectral Element Method

A key factor in the accuracy and computational efficiency of the NEKTON package is a sophisticated mathematical technique known as the *spectral element method.* This method is a high-order, weighted residual finite element technique for numerical solution of time-dependent, incompressible fluid flow and heat transfer equations.

The spectral element method combines the geometric flexibility of finite element methods with the rapid convergence of spectral techniques. In a spectral element discretization, the computational domain is broken

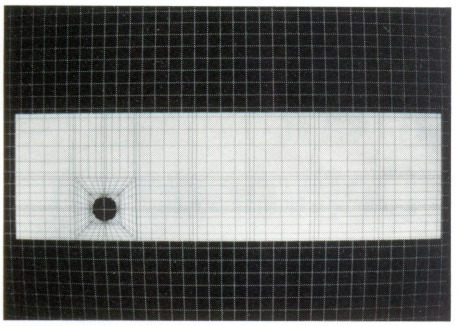

Spectral element mesh for an Eddy-Promoter heat exchanger.

up into relatively few, large, spectral elements. Variables are approximated by Nth order tensor-product polynomial expansions, and convergence is achieved by increasing the order of the polynomial while keeping the number of elements fixed.

The spectral element method gives high accuracy with relatively few grid points by using high-order interpolation and by automatically clustering points near domain boundaries. Elements may be placed in regions of rapid variation to improve accuracy.

Solutions are achieved efficiently due to the relatively few degrees of freedom required, the use of tensor-product sum factorization, and the use of pre-conditioned iterative inversion procedures.

The iPSC®-VX release of NEKTON uses an innovative algorithm that maximizes the computation-to-communication ratio. Inasmuch, speedups are nearly linear with the number of nodes. With the Direct-Connect™ Routing scheme, the iPSC®/2 should provide even greater speed and efficiency.

Intel Demos Firsts in Parallel AI at IJCAI

Intel Scientific Computers demonstrated a number of firsts in parallel AI at the IJCAI (International Joint Conference on Artificial Intelligence) in Milan, Italy during the week of August 21 to 28. "It was an excellent conference," commented David Billstrom, Product Marketing Manager of iSC, "and we were pleased to show the AI community a number of first-ever parallel tools and demos." Intel teamed up with Artificial Intelligence Limited, an iSC distributor in England, to bring 14 software demonstrations to the show, in three overall categories.

For LISP programmers, Intel showed two new **programmer-oriented user interfaces: HYDRA** on a XEROX 1186 workstation, and **VIEWS** on a Sun-3 color workstation. Each iPSC node was viewed through a window using the standard windowing environments on the Sun and XEROX workstations. These were firsts in parallel AI, allowing programmers access to the CCLISP programming environment via familiar workstations, and drew praise from CCLISP users. HYDRA was developed by Artificial Intelligence Limited specifically to join the existing XEROX LISP and iPSC environments.

Another highlight of the CCLISP demonstrations was **NIRA (Northern Italian Renaissance Art)**, an art identification program. Using an Intel SugarCube™ with the XEROX 1186/HYDRA workstation interface, the user describes a painting via menus and pointers. NIRA uses a parallel inferencing algorithm on a distributed database to identify the era and artist, asking for additional information when necessary. Because of the small domain size, this program was an excellent fit for the four-node SugarCube, with very good speedups from parallel processing.

The second group of demonstrations centered around **SOPE (Systems Object Programming Environment)** from Advanced Decision Systems of Mountain View, CA. SOPE is a programming environment for object-oriented programming on a parallel computer. The primary demonstration was LOSP, a parallel deduction engine that uses a single proto-logical concept to execute deductive steps in highly parallel fashion. The LOSP program has excellent human interface features and runs on the Sun-3 workstation with an iPSC D4-MX system.

Flat Concurrent Prolog from the Weizmann Institute was shown in a compiled version, the first *compiled* language implementation for the concurrent logic community. The compiled FCP gives substantial speedups over the previous, interpreted versions of FCP. The demo included a concurrently executing polygon graphic modeling package implemented in FCP by Steve Taylor, Ph.D. student at the Weizmann Institute, Israel. The package showed a 14-fold speedup on a 16-node iPSC.

Intel and Artificial Intelligence Limited worked together with the Weizmann Institute in Israel and Advanced Decision Systems in California to bring hardware and software together and to integrate the systems at the show. In spite of 100° temperatures, a severe lightning storm, and heavy rains, all of the systems worked well.

"The show was an overwhelming success for parallel AI in Europe," continued Billstrom. "We were very impressed with the level of interest in parallel AI, and the questions we received showed a high level of understanding of the potential for concurrent computing in AI."

Conference Announcements

Conferences provide a gentle introduction to concurrent computing for the newcomer, and well-known rewards for the initiate. An additional benefit for all is that the following winter conferences are held in southern California:

- **Third Conference on Hypercube Concurrent Computers and Applications (HCCA3).** January 19-20, 1988. Co-sponsored by California Institute of Technology and Jet Propulsion Lab. To be held at the Pasadena Conference Center, Pasadena, CA. For information, contact the conference coordinator: Dr. Andrew Witkowski, Jet Propulsion Lab, MS 138-208, Pasadena, CA 91109, 818-354-2244 ARPAnet: andy@hamlet.caltech.edu.

- **AEROSPACE 88: AIAA 1988 Annual Meeting and International Aerospace Exhibit.** February 9-11, 1988. Sponsored by the American Institute of Aeronautics and Astronautics. To be held at the *L.A. Airport* Hotel, Los Angeles, CA. For information contact: AIAA, 370 L'Enfant Promenade, Washington, DC 20024.

Users Group Seeks Cubelib Submissions

The iPSC Users Group is actively encouraging submissions to Cubelib, the software library, said Victor Jackson, the administrator of Cubelib and an active Users Group member.

Cubelib provides a convenient way for iPSC users to share non-proprietary software, and to gain both feedback and recognition for their development efforts. The library includes tools and techniques to help shorten product development cycles and optimize cube use, as well as complete applications that can serve as *templates* for developing similar applications.

The following packages are currently available through Cubelib:

- Cube_use — utilities for monitoring use of the iPSC
- Eiscube — concurrent routines for symmetric eigenvalue/eigenvector computations
- Lincube — concurrent routines for solving systems of linear equations, including an efficient implementation of dense factor and solve routines
- GESBT — Generic Expert Systems Building Tool, for constructing expert systems in a concurrent environment
- Navier — a concurrent solution to the Navier-Stokes equations using an explicit Lax-Wendroff scheme
- Sparspak — concurrent routines for solving sparse systems of linear equations
- Suprenum — grid-oriented routines for problem partitioning
- Technotes — letters to customers, release notes, and technical notes published by iSC

Users can obtain information about the programs in Cubelib via electronic mail or conventional mail. By email, just send the message "send index"; addresses are:

...!intelisc!cubelib
for uucp users

cubelib@isc.intel.com
for CSNET, Arpanet users

Send requests by conventional mail to the Users Software Librarian as listed below.

User submissions play a vital role in the success of the library, and must conform to certain minimal guidelines regarding generality, ease of use, robustness, and documentation.

For details on submission or current Cubelib contents, contact:

User Software Librarian
Intel Scientific Computers
15201 N.W. Greenbrier Parkway
Beaverton, OR 97005
503-629-7600

The iPSC®/2 Concurrent Supercomputer – Delivering on the Promise of Concurrent Supercomputing

"The iPSC®/2 Concurrent Supercomputer is going to change the way people think about supercomputing," said Charlie Bishop, iSC Marketing Manager, "just as the PC changed how we think about office computing. The performance of the iPSC/2 system initiates the era of *interactive* supercomputing, while the price makes supercomputing truly affordable."

The first-generation iPSC system successfully demonstrated the hypercube architecture as a viable approach to concurrent computing. Incorporating the knowledge and experience gained from the iPSC system, the 80386-based iPSC/2 system incorporates design innovations that substantially broaden the range of applications that are appropriate to concurrent machines. "We are seeing the 'first fruits' of these applications in high performance software packages such as the NEKTON fluid dynamics code," stated Bishop. "The iPSC/2 is the first true production code hypercube."

The iPSC/2 system consists of three major components:

- The *cube* — a system of 32 to 128 computational nodes, each built around Intel's 32-bit 80386 as a node processor, coupled with an 80387 floating point accelerator. Memory capacity is 1, 4, 8, or 16 MBytes per node.

Two numerical performance enhancements are available: the *SX Scalar Extension* triples the scalar performance of the basic iPSC/2 system, and the *VX Vector Extension* increases vector performance over that of the iPSC by an order of magnitude.

- The *Direct-Connect™ Routing Network* — connects all of the iPSC/2 processing nodes together. Although nodes are physically connected as a hypercube, Direct-Connect Routing offers performance that functionally connects every node to every other node.
- The *Concurrent Workbench™* — a systems programming support package that provides a variety of languages, a simulator, vector development tools, and a concurrent debugger.

The Concurrent Workbench software development environment is standard with

the iPSC/2 system, and is hosted on the System Resource Manager, accessible from a network of workstations such as the SUN-3. The System Resource Manager serves as the administrative console and a gateway to the cube for network users. The System Resource Manager is also an 80386-based system, and features 8 MBytes of memory, a 140 MByte hard disk, and an Ethernet connection.

The iPSC/2 system employs the same modular design and uses the same vector boards used in the first-generation iPSC. A principal objective in the iPSC/2 development effort was to provide an effective migration path for first generation iPSC users. Consequently, even though the hardware and system software have been substantially enhanced, software for the iPSC is upward compatible with the iPSC/2. This allows iPSC users to upgrade to a corresponding iPSC/2 system without losing any part of their software investment.

In addition, the iPSC/2 system offers major advances over the first-generation iPSC system, including:

- Increased node performance (4 to 5 times) with the 80386 node processor.
- Faster node-to-node communications (3 to 10 times) with a substantial reduction in multihop latency made possible with the iPSC/2's *Direct-Connect Routing* (see related article).
- Larger memory capacity, expandable from 1 Mbyte to 16 Mbytes per node.
- Significant improvements in the user development environment, and the benefits of a 32-bit system environment.

Performance

The combination of increased performance for each system node, the Direct-Connect Routing communications scheme, and larger memory capacity enables the iPSC/2 system to give *true supercomputer performance.* For example, a 32-node iPSC/2-VX machine performs a 2D FFT (Fast Fourier Transform) at 154 MFLOPS, 10 times faster than the first-generation iPSC. The 2D FFT represents the worst-case node-to-node communications in hypercube architectures. In addition, a 128-node iPSC/2 system executing the Gabriel Triangle symbolic benchmark for LISP systems tops the performance of the previous record holder, the CRAY 1S, by more than a factor of four.

Productivity and Ease of Use

In addition to the 80386 node processor, three factors contribute to the increased productivity of the iPSC/2 system: mainframe-quality languages and system support, the Concurrent Workbench system support package, and the Direct-Connect Routing communications scheme.

The move from a 16-bit to a 32-bit node architecture gives users access to a new class of production compilers and software tools. "People used to spend most of their time porting code to the limited, 16-bit environment of the original iPSC," stated David Billstrom, Product Marketing Manager. "Now *every node* of the iPSC/2 has the performance of a VAX 8600, *and* with a 32-bit software model. Users can spend their time solving problems, not porting code."

Another major factor in productivity and ease of use is the iPSC/2 system's Concurrent Workbench. Now, all of the software tools developed on the first iPSC system are gathered together in an integrated environment, and this environment appears to run on the user's Unix*-based workstation. Hosted on the System Resource Manager, the Concurrent Workbench allows the user to access, program, and control the iPSC/2 system in the more familiar and productive windowed environment of today's popular workstations.

The programming of first generation hypercubes was characterized by an emphasis on problem decomposition to maximize "nearest neighbor" communications. Direct-Connect Routing enhances productivity by effectively removing these "nearest neighbor" constraints.

With Direct-Connect Routing, it is no longer necessary to precisely match problem and machine topologies to achieve high computational efficiencies (see related article). You can now program the hypercube as an ensemble of processors with an arbitrary communications network in which each node communicates efficiently with all other nodes.

Production Software

This fall, Intel and Nektonics, Inc. jointly announced the availability of the NEKTON fluid dynamics and heat transfer simulation code for the iPSC/2. NEKTON is the first production code for any large-scale parallel computer (see article, p. 1).

In addition, Intel plans to announce four other production-quality codes in the next few months: a molecular modeling package, an extruded materials modeling package, a VLSI device simulation code, and an event-based simulation code. These applications will bring the supercomputer performance of the iPSC/2 Concurrent Supercomputer to a new class of users, those who use computing resources as a standard productivity tool but do not have access to a traditional supercomputer.

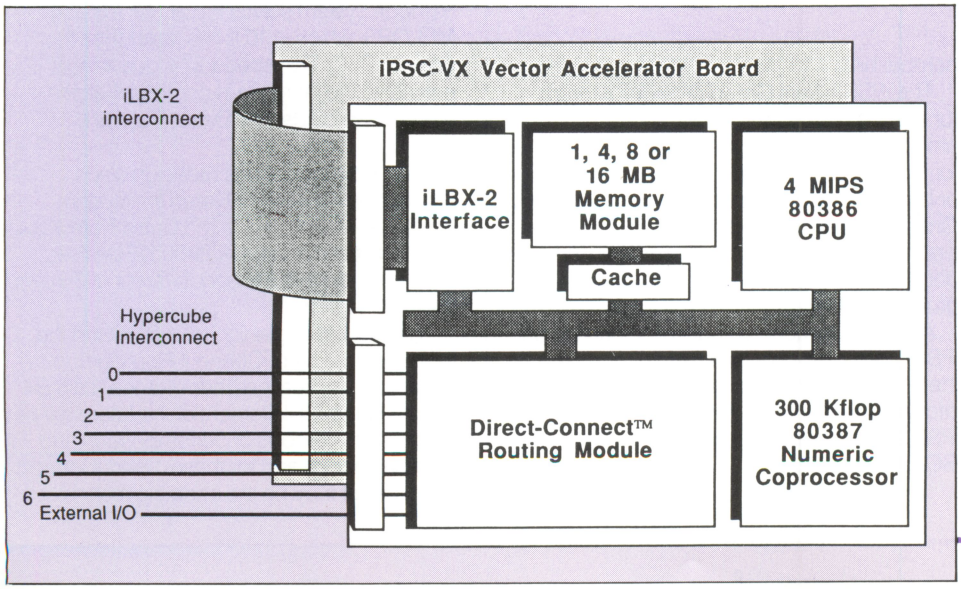

Direct-Connect™ Routing Solves Node Communications Challenge

With the introduction of the iPSC®/2 Concurrent Supercomputer, Intel Scientific Computers has solved one of the major bottlenecks in hypercube computing, the problem of passing messages to distant nodes quickly and without degrading node processor performance.

"The benefits this brings are much more than just speed," stated David Billstrom, iSC Product Marketing Manager. "With Direct-Connect™ Routing, message passing times are essentially *uniform,* as though every node were *connected directly* to every other node. This feature is unique to the iPSC/2 system, and will change the way programmers use the hypercube architecture. It heralds a major change in hypercube technology."

As an example, a large system such as the 128-node iPSC/2 d7 system requires only about 10% longer to send a message between the most distant nodes than

continued on page 6

Direct-Connect™ Routing Solves Node Communications Challenge

continued from page 5

between "nearest neighbor" nodes. As a result, computational efficiency is essentially independent of the problem domain-to-machine topology mapping. This greatly simplifies problem decomposition and frees the programmer to concentrate on resolving other programming issues.

The Direct-Connect Routing scheme uses improvements in both hardware and software to shorten message-passing times. With Direct-Connect Routing, the earlier "store and forward" method used in the first-generation hypercubes is replaced by a hardware switching system, the *Direct-Connect Module,* on each node.

Each Direct-Connect Module (DCM) functions much like a telephone exchange with eight full-duplex channels. Seven of these channels are used for internode message passing, while the eighth channel provides connection to the System Resource Manager or I/O facilities. Because DCM channels are independent, up to seven messages can be passed simultaneously over each node, avoiding bottlenecks when several nodes must communicate at the same time.

The basic concepts for Direct-Connect Routing were developed at California Institute of Technology by Dr. Charles Seitz and Dr. Bill Daley, with research cosponsored by Intel Corporation and DARPA. In addition to uniform communications, DCR provides significantly improved node-to-node communications by reducing latency and increasing communication bandwidth. Latency — the time to send a zero-length message between nearest neighbors — is about 3 times faster than the original iPSC and an order of magnitude faster than many "real time" operating systems.

Experience has shown that the optimum performance from hypercube systems requires both small latency for short messages *and* high communications bandwidth for long messages. The Direct-Connect Routing scheme minimizes the time required to deliver short messages (less than 100 bytes), while allowing messages greater than 100 bytes to move at the full bandwidth of the network. Both types of messages benefit from this scheme, which first establishes a direct connection between the source and destination nodes, then streams the message through at essentially hardware data rates (see illustration and table).

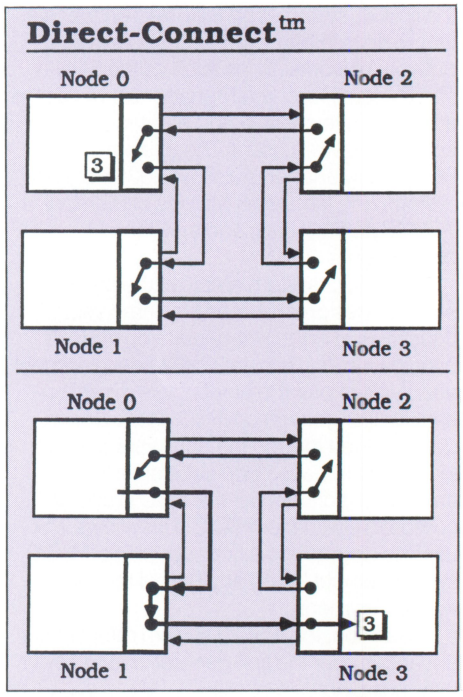

Direct-Connect™ Routing uses a special algorithm for messages longer than 100 bytes. This algorithm first sends a header message to the destination node (Node 3). This header sets gates in each DCM module on the intermediate nodes, clearing a data path for the message. Once the destination node acknowledges receipt of the header, the message is streamed through at essentially hardware data rates, without "packetization," to the destination node.

Hop Length	Store & Forward Latency (μsec)	Direct-Connect Routing Latency	DCR Percentage	Latency Speedup
1	950	350	37%	2.7
2	1,300	352	27%	3.7
3	1,650	354	21%	4.7
4	2,000	356	18%	5.6
5	2,350	358	15%	6.6
6	2,700	360	13%	7.5
7	3,050	362	12%	8.4

This improvement is demonstrated by the iPSC/2 system's performance on the 2-dimensional Fast Fourier Transform (FFT). The FFT has a reputation for "worst case" communications on hypercubes, because all nodes need to trade intermediate results with all other nodes at roughly the same time. An iPSC/2 32-node system solves the 2D FFT at 154 MFLOPS, which is 10 times faster than the original iPSC. This speedup is due almost entirely to the increase in communications efficiency provided by Direct-Connect Routing.

The communications redesign also improves data rates between the System Resource Manager and the Cube, from about 80 KBytes/sec to a peak 2800 KBytes/sec, giving a system I/O speedup factor of 35 over the iPSC system.

"We are just beginning to understand the full benefit of Direct-Connect Routing and its impact on application performance," said Billstrom. "But the excitement of the future will be automated methods of decomposition and load balancing, and for these, Direct-Connect Routing just may be the enabling technology."

Users Group to Hold iPSC®/2 Workshop

The iPSC Users Group has scheduled a one-day users meeting and special iPSC®/2 system workshop for Monday, January 18, 1988 in Pasadena, CA. This places the workshop just before HCCA3, the third annual hypercube conference. Workshop arrangements will be convenient for HCCA3 attendees.

The workshop will include personalized, in-depth presentations, tutorials, and "inside information" to help users take full advantage of the new features of the iPSC/2 Concurrent Supercomputer. Topics for the workshop include new product features, staged enhancements to the iPSC/2 system, concurrent I/O solutions, benefits of the new user interface (with expanded front-end support), and iPSC compatibilty.

The iPSC Users Group provides a forum for sharing techniques, ideas, and tools among iPSC users, as well as an opportunity for input concerning new products and enhancements. Intel is working with the Users Group to stage the workshop and tutorial session. The Users Group, however, is an independent organization, with officers elected by group members.

Anyone with access (logon authority) to an iPSC system can join the Users Group, and there are no membership fees. For further information about the Users Group or the iPSC/2 workshop, contact Victor Jackson of iSC, 503-629-7704.

iSC Training Refocused for iPSC®/2 System

The iSC training program in programming concurrent computers is being updated and refocused to support the new iPSC®/2 Concurrent Supercomputer, said Victor Jackson, Senior Training Specialist for iSC. "The basic class has been modularized to provide individual emphases on numeric and vector computing," continued Jackson. "The class provides a solid introduction to the new iPSC/2 system, but users of the first-generation iPSC system will continue to find the class an excellent resource."

The basic sequence now begins with a three-day session, *Introduction to Numerical Concurrent Computing* that includes hardware and software overviews and basic concurrent programming techniques. This is followed by a two-day session, *Introduction to Vector Concurrent Computing,* that covers vector programming fundamentals for the iPSC/2 VX and use of the VAST-2 Fortran Vectorizer.

The iSC Training Center also offers *Introduction to Symbolic Concurrent Computing,* featuring the fundamentals of programming in Concurrent Common LISP. This workshop is offered alternate months. "We are very enthusiastic about the symbolic computing class," said Jackson. "With the price/performance and memory capacity of the iPSC/2 systems, we expect to see growing interest from the AI community." While familiarity with LISP programming is recommended, no background in concurrency is needed for the class.

Although you can come into either of the workshops "cold," a little advance preparation can make class time more productive, added Jackson.

"If you haven't programmed recently, it helps to review a bit, just to get the wheels turning again. Also, come with a specific problem to solve."

Fall Training Schedule

Introduction to Numeric Concurrent Computing
February 22-24
March 28-30
April 25-27

Introduction to Vector Concurrent Computing
February 25, 26
March 31, April 1
April 28, 29

Introduction to Symbolic Concurrent Computing
February 17-19
April 20-22

Concurrent Workbench™: Easing the Software Development Task

With 30 years of sequential codes for scientific and industrial applications, it's no surprise that software development is a major concern in concurrent computing. One of the prime objectives in the iPSC®/2 system design was to make software development as easy and productive as possible.

The result is the *Concurrent Workbench™,* a Unix*-based development environment that allows the user to access, program, and control the iPSC/2 system from his or her own workstation. This is a significant change from the earlier, first-generation systems, which required all development to be done on a XENIX-based Cube Manager.

Most Concurrent Workbench tools running on the System Resource Manager appear to run on the user's workstation. This means that the user can write, edit, and debug programs for the iPSC/2 system in an environment that is the most familiar, comfortable, and productive. The SUN-3 workstation from SUN Microsystems is the first workstation to be supported; other popular workstations will be supported in the near future.

The Concurrent Workbench includes multiuser, network access to the iPSC/2 system. It supports *cube sharing,* which allows several users to partition the cube into functionally distinct subcubes and increases the versatility and productivity of the iPSC/2 system.

Based on AT&T Unix* V.3 on the System Resource Manager, the Concurrent Workbench provides a full set of software development tools. These tools include the iPSC/2 Simulator software package, C compiler, optimized FORTRAN 77 with extensions, and *DECON,* a high-level symbolic process debugger. DECON supports simultaneous debugging of node processes and internode communications for all nodes of the cube, and represents one of the most sophisticated debugging facilities available for concurrent systems.

Other Concurrent Workbench tools are available at additional cost. For vector processing applications, the *VAST-2 FORTRAN Vectorizer* discovers sections of Fortran code that can be vectorized and translates these sections into vector commands that execute on the iPSC/2 VX machine. VAST-2 has been highly ranked among a group of supercomputer vectorizers by a national laboratory, and comes from Pacific Sierra Research, a leading supplier of vectorizing software for supercomputers.

For AI applications, a full implementation of *Concurrent Common LISP (CCLISP)* from Gold Hill Computers, Inc. is also available. LISP is the most popular AI language in the United States, and CCLISP is the first commercial LISP for a parallel computer. A single iPSC/2 node running CCLISP gives performance roughly equivalent to a Symbolics/LISP system. With the iPSC/2's memory structure, up to 128 nodes can be configured to run CCLISP, making the iPSC/2 one of the fastest AI systems available.

With the Concurrent Workbench™, the user can access, program, and control the cube from a convenient, windowed workstation environment. The SUN-3 workstation is the first to be supported; support for other popular workstations will be announced soon.

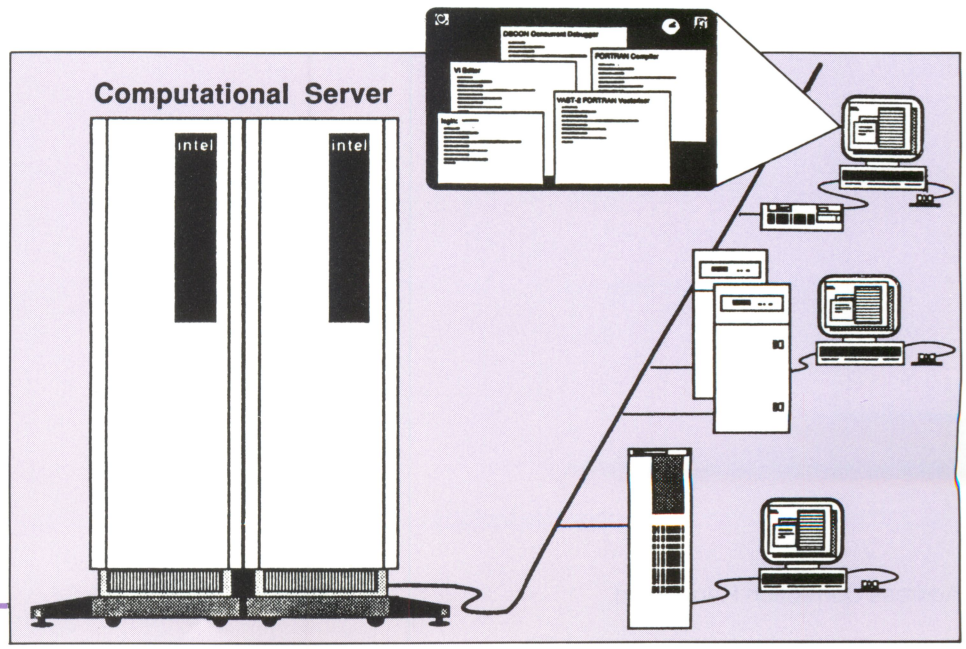

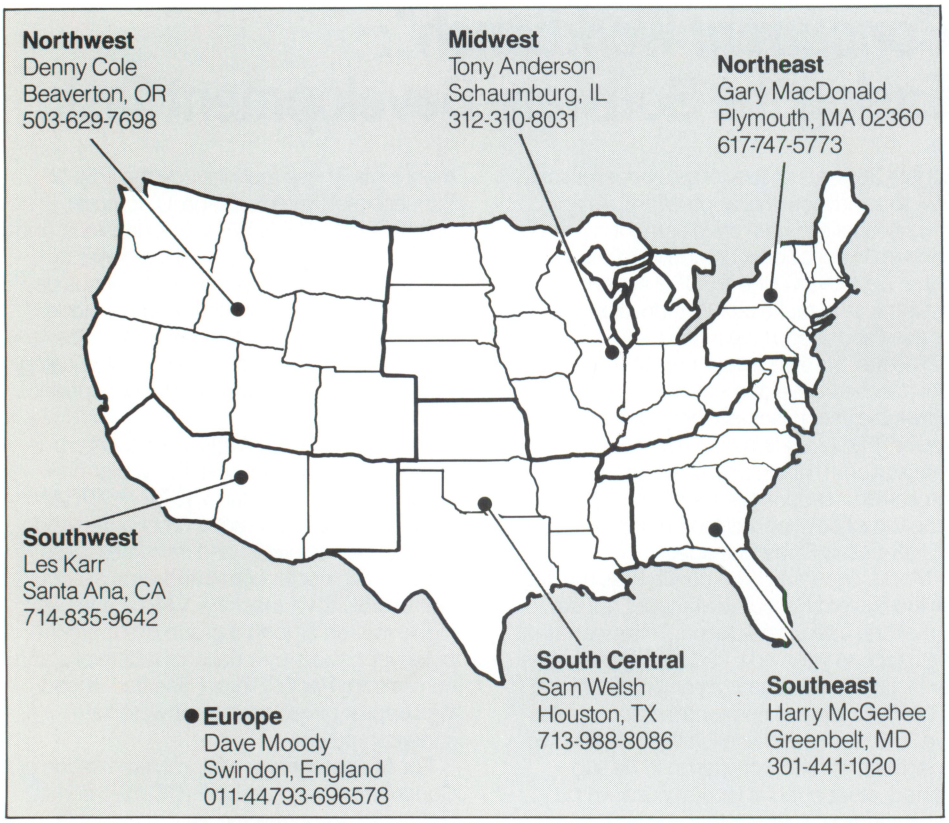

For More Information...

"We work more like partners than vendors," says Charlie Bishop, Marketing Manager for Intel Scientific Computers. "Concurrent computing is still new to a lot of folks. So a lot of our work is simply technical consultation, helping people to understand the new technology, its benefits, and how to use it."

Each iSC Representative has several years of training and experience in the field, and each one provides a range of services, from informational seminars to post-sales support.

For more information about the iPSC®, iPSC®/2, or Sugarcube™ concurrent computing systems, contact one of these iSC Representatives.

Do you know anyone....

...who might be interested in receiving *iSCurrents*? If so, please contact:

iSCurrents
Intel Scientific Computers
15201 N.W. Greenbrier Parkway
Beaverton, OR 97006

Intel Scientific Computers
15201 N.W. Greenbrier Parkway
Beaverton, OR 97006

Address Correction Requested