iSCurrents

Published by, Intel Scientific Computers

Fall

iPSC Calculates Peptide Conformations

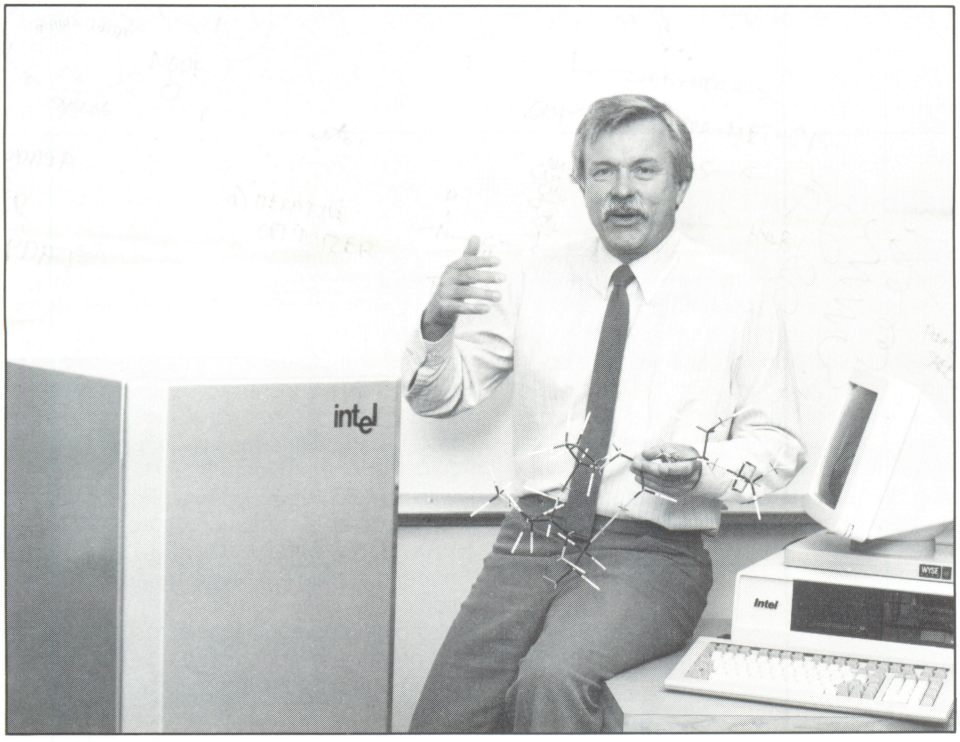

Al Larrabee is shown here beside the Oregon Graduate Center's iPSC d5 concurrent machine, which he used to calculate peptide conformations. For this task, Larrabee developed ECEPP86, a concurrent Fortran program that iteratively calculates minimum energy levels in peptide molecules.

Al Larrabee, Ph.D. in biochemistry (MIT, 1962), has just completed work calculating minimum energy conformations of various peptides, using the Intel iPSC system at the Oregon Graduate Center in Beaverton, Oregon.

For this study, Larrabee developed a concurrent version of the 5100-line Fortran program, ECEPP83 (Empirical Conformational Energy Program for Peptides-1983 Edition, courtesy of Polygen Corporation, Waltham, MA). Larrabee's program, called ECEPP86, performed 17 times faster on the iPSC d5 32-node system than its counterpart on the VAX 11/780.

Because the algorithms for computing peptide conformations have a high degree of inherent concurrency, Larrabee said that his study provides reliable estimates of the speedups possible with larger iPSC configurations. ECEPP86 will soon be available as public domain software from the iPSC Users' Group Software Library.

Larrabee's interest in peptides stems from the computational challenges involved, as well as the critical roles peptides play as hormones and hormone-releasing factors in cell interactions. Each peptide is a small protein, a sequence of amino acids folded and wrapped into a characteristic three-dimensional *conformation*. This conformation determines exactly how the protein interacts with receptors, membranes, and other cell structures.

Methods for determining protein structures include spectroscopy and X-ray crystallography. However, since many peptides cannot be crystallized, their conformations (specified as dihedral angles for non-rigid molecular bonds) must be determined by computational methods.

These computational methods are based on the principle that each molecule assumes a conformation *that minimizes its global energy level*. The energy of each bond within the molecule depends on the bond angle. As the molecule folds over itself in space, however, second order interactions occur, so that the global minimum energy conformation cannot be found simply by minimizing individual bond energies.

(Continued on page 2)

In this Issue of iSC

iSCurrents

Volume 1 Fall, 1986

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

Intel Scientific Computers
15201 NW Greenbrier Parkway
Beaverton, OR 97006
(503) 629-7629

iPSC Users' Group to Meet at Hypercube Conference

The iPSC Users' Group will hold its second meeting on Tuesday, September 30, at 7:00 PM in the United Room of the Airport Hilton Inn, Knoxville, Tennessee, at the Oak Ridge Hypercube II Conference.

The Users' Group was founded to provide a forum for sharing information, techniques, and tools among users of the iPSC and related products from Intel Scientific Computers.

"The need was there," said Bill McLain (who organized the group), "and with so many hypercube users within iSC, we were the natural ones to organize it. But the group belongs to the users, not to iSC. Once officers are elected, the group will become self-governing. However, iSC will continue to provide support as requested by the group."

Organized in April, 1986, the Users' Group already has over 100 members from 30 different organizations. It has published a membership roster to promote information exchange, and a users' software library is under development, with details to be announced at the Hypercube II conference.

Anyone with access (authority to log on) to an iPSC is eligible to join (no membership fees). Members receive an information packet that includes the charter, membership lists, and other pertinent information.

iPSC Calculates Peptide Conformations (Con't.)

The approach taken by ECEPP86 is to use many *conformers*, or initial guesses for the dihedral bond angles. From each conformer, Larrabee's program computes a neighboring conformation that locally minimizes the peptide's total energy. If enough conformers are tested, the smallest of these local minima will be the global minimum energy of the peptide (Figure 1).

For each conformer, the program performs multiple *minimization cycles*. During each cycle, the non-rigid bonds are rotated, one by one, to minimize total molecular energy. Then, to account for second-order effects, this process is repeated until additional cycles give no further energy reduction.

Even for small proteins, this is a formidable computing task. "Exhaustive testing of a small peptide (6 or 7 amino acids), with perhaps thousands of conformers and thousands of minimization cycles for each conformer, would require several days of CPU time, even on a mainframe," said Larrabee. "One mainframe test (using a Xerox Sigma 9) took 17 days of CPU time to complete."

Previous benchmarks show that the iPSC finds conformations 3 to 5 times faster than mainframes such as the IBM 370/168, Xerox Sigma 9, and Prime 350 with an FPS AP-120B Array Processor.

Larrabee's study compared a DEC VAX 11/780 with the Intel iPSC d5 (32-node) machine running a single minimization cycle using many conformers. By distributing computations across the iPSC's 32 nodes, Larrabee found that the computation could be speeded up by a factor of 17.

"This was an ideal benchmark study," said Larrabee, "because it is almost perfectly parallel. Each node minimizes energy for a different set of conformers, so the speedup increases as the number of nodes increases. A 64-node iPSC system will be 34 times faster than the VAX 11/780."

Larrabee plans to extend his work to a concurrent vector environment in the near future. "There are many processes in ECEPP86 that can be vectorized," he said. "I intend to explore these on the Intel iPSC-VX."

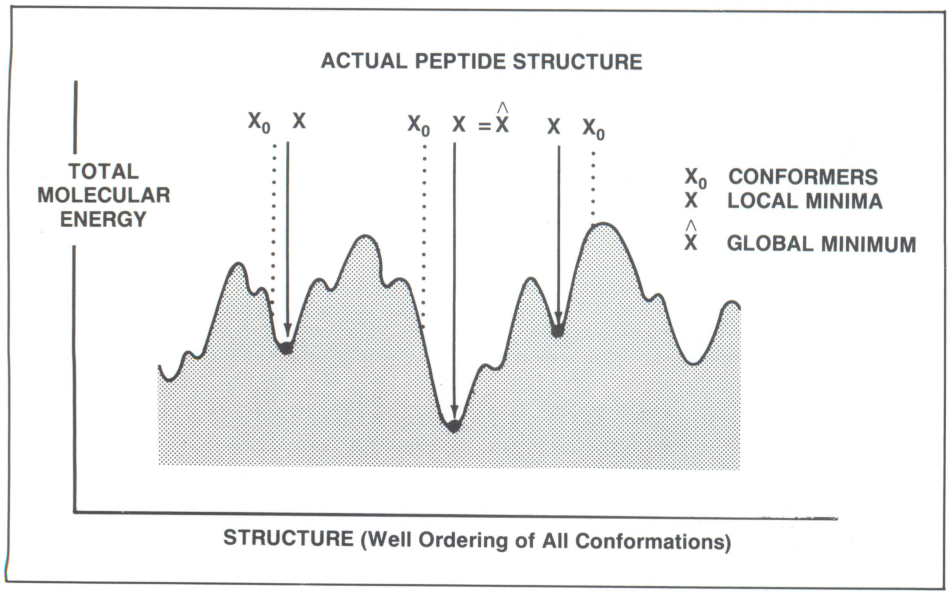

For each conformer, ECEPP86 calculates a neighboring conformation with (local) minimum energy. If the conformers are well chosen, the global minimum energy will lie among the discrete set of local minima, and the corresponding conformation will be the peptide's actual conformation.

Concurrent LINPACK, EISPACK Routines Available

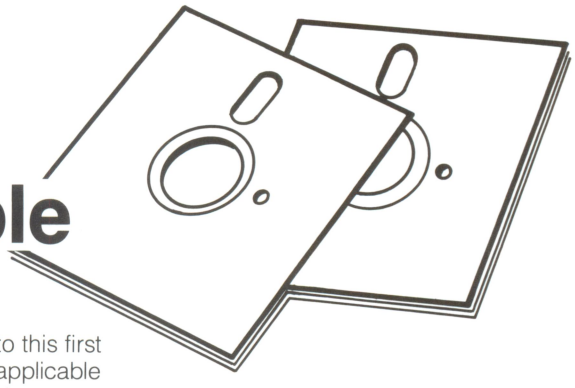

Key modules of LINPACK and EISPACK, industry-standard matrix computation packages, are now available in concurrent versions for the Intel iPSC system, announced Cleve Moler, manager of the Applications Research Group at Intel Scientific Computers. These routines are available through the iPSC Users' Group Software Library.

Work on the LINPACK (linear equations) and EISPACK (eigenvalue/eigenvector) routines is being carried out by Cleve Moler, Bill Hughey, and David Scott, members of the applications group.

Routines now available include factor and solve routines for general square and symmetric positive definite matrices, least squares routines for rectangular matrices, and eigenvalue/eigenvector routines for symmetric matrices. Because these are the most frequently used routines, they were the first to be implemented in concurrent versions.

The programming approach to this first collection of routines is generally applicable to the remainder of LINPACK and EISPACK, said Moler. Complex versions of the above routines are now under development, as well as eigenvalue/eigenvector routines for nonsymmetric matrices and singular-value decomposition routines for rectangular matrices. These will be available over the next six months.

In addition, this approach was intended from the beginning to be applied to the iPSC-VX concurrent vector machine. LINPACK and EISPACK routines were used as debugging tools during VX development. Moler expects to have a detailed performance analysis of LINPACK on the iPSC-VX system by the time of the Oak Ridge Hypercube II Conference in September.

Performance of LINPACK factor and solve routines has become a standard computing benchmark. The iPSC d5 (32-node) system solves a 100×100 double-precision matrix in about 4 seconds, with overall execution time between 4 and 5 seconds with 8- to 32-node configurations.

Efficiency of concurrent methods increases with matrix size. For example, a full dense matrix of order 1300 can be solved on a 64-node system in about 15 minutes. This problem requires 16 MB for matrix storage and represents about 70% utilization of the iPSC's floating point capacity.

Moler will discuss a general approach to concurrent matrix computation, including an analysis of performance measurement and modeling, in a paper titled "Matrix Computation on Distributed Memory Multiprocessors." This paper will appear in the Proceedings of the Hypercube Conference, as well as in a Technical Report available from Intel. A second report will discuss use of the software.

Moler is internationally known for his work in numerical computation, especially for development of LINPACK (with J. J. Dongarra, J. R. Bunch, and G. W. Stewart) and EISPACK (with B. T. Smith, J. M. Boyle, J. J. Dongarra, B. S. Garbow, Y. Ikebe, and V.C. Klema). The *LINPACK Users' Guide*, coauthored by Moler, has become the bestselling SIAM publication ever, with over 10,000 copies in print.

Oak Ridge Hosts Hypercube Conference

Oak Ridge National Labs will host the Second Conference on Hypercube Multiprocessors, September 29-31 in Knoxville, Tennessee. The conference, co-sponsored by Oak Ridge and SIAM, features eight invited speakers, numerous contributed talks, and exhibits by several vendors.

Invited speakers are:

Brian Beckman, JPL
George Cybenko, Tufts
Geoffrey Fox, Cal Tech
William Gropp, Yale
John Hayes, Michigan
Michael Heath, Oak Ridge
Don Heller, Shell Research Corp.
Andrew White, Los Alamos

Each invited speaker will discuss hypercube activity in his organization, then focus on one area of particular interest. Topics include hypercube research (architecture, algorithms, etc.), as well as a range of applications such as seismic processing, computational fluid dynamics, and others.

"This year will see much more emphasis on applications," said Mike Heath, conference chairman for both years. "Last year commercial hypercubes were brand new. People talked about what they were planning to do. This year, people will talk about what they've done — reality, not dreams."

The proceedings of last year's conference are being published by SIAM, with Heath as editor, and should be available at the September meeting.

Do you know anyone....

...who might be interested in receiving this newsletter? If so, please send his or her name and address to us so it can be added to our mailing list. Send to: *iSCurrents*, Intel Scientific Computers, 15201 NW Greenbrier Parkway, Beaverton, OR 97006.

iPSC Training Class Enhanced

"The iSC workshop *Programming Concurrent Computers* is being enhanced to reflect changing needs and interests of our customers," said Victor Jackson, Senior Training Specialist. Beginning in November, the five-day workshop will consist of three days devoted to an introduction to the concurrent architecture and programming of the iPSC, and a two-day focus on advanced topics. "This will give us the flexibility to tailor each workshop to individual interests," said Jackson.

One focus for the final two days will be concurrent vector programming, emphasizing architecture and programming techniques for the iPSC-VX vector machine. Another will cover advanced concurrent programming on the iPSC. Attendees will choose which focus to explore during these last two days.

The workshop will continue its strong applications focus, with afternoons devoted to hands-on programming. Although lab exercises are provided, "we encourage people to bring their own problems," said Jackson, "perhaps a core of very important code, or a simplified version of their application. Often, people leave the class with part of their application already running concurrently."

Although most attendees are iPSC users, the class is open to everyone. "This is an excellent way to get an overview of iSC products or just explore the whole concept of concurrent computing," said Jackson. Class dates for the remainder of this year are:

September 15-19
October 13-17
November 17-21
December 15-19

Separate classes in microcode development for the iPSC-VX and advanced concurrent applications programming are also under development and scheduled to be offered in the first half of 1987.

For details, contact Intel Scientific Computers, Training Center, 15201 NW Greenbrier Parkway, Beaverton, OR 97006, (503) 629-7629

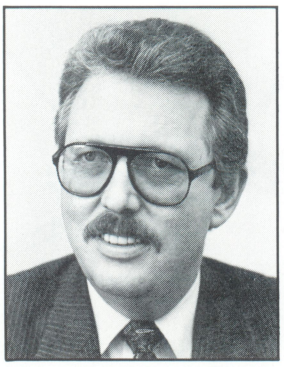

Robert Rockwell, new General Manager of iSC.

iSC Names Rockwell as New General Manager

Intel has named Robert Rockwell as the new General Manager of iSC. "My responsibility is to build iSC into a strategically important business for Intel," said Rockwell, who has over 20 years experience in computer sales and marketing at DEC, Okidata, and most recently, Axiom. "My presence here is a reflection of Intel's commitment and what I believe to be Intel's potential in this market."

Rockwell continued, "With the hypercube architecture and with Intel's strong technology base, resources, and commitment to the systems arena, iSC is ideally positioned to become the dominant player in the high-performance market." Market analysts predict that by 1990, the high-performance market will be between $1 billion and $2 billion.

The fundamental success factor, Rockwell believes, will be rapid development of concurrent applications software. With the Applications Research Group, formed to promote software development both by Intel and by third-party vendors, iSC is playing an active role in this effort.

As General Manager, Rockwell replaces Justin Rattner, who has been named to the new position of Director of Technology. Rattner, the "father" of iSC, will continue to drive product development, to explore the evolution of concurrent architectures and to apply new techologies to these architectures.

SIMULATOR Package Enhanced

In Release R2.1 (August 1986), the SIMULATOR concurrent software development package has been enhanced to support the *dynamic loading* capability of the current iPSC operating system. Dynamic loading allows the node operating system software and application software to be downloaded to the nodes independently. This reduces the number of compile, link, bind and load steps, giving faster turnaround times during software development.

First shipped as part of the iPSC system, the SIMULATOR package was unbundled and offered as a separate package in February, 1986. This enabled customers to begin developing concurrent algorithms and prototyping code immediately. In addition, for those with a developing interest in concurrency, it provided a stand-alone exploratory tool for a variety of popular UNIX hosts at a minimal price.

The SIMULATOR also continues to be shipped as part of the iPSC system, to allow offline code development on a workstation and reduce the computing traffic on the Cube Manager. The package runs under any Xenix or Unix BSD 4.2 operating system, including the IBM PC/AT, Intel System 310, Sun and Apollo workstations, and host systems such as DEC VAX, Sequent, and Pyramid.

In the United States, the SIMULATOR package is available from Intel for $495 (with educational discount, $297); contact Intel Scientific Computers (503-629-7629). In Europe, the package is available through Intel International; contact Dave Moody in Swindon, England at 44-743-696578.

Next in iSCurrents

Our next edition will feature "Europe" as our theme. We will be showing you some facinating concurrent applications developed by our European colleagues who are using the iPSC. The issue will also cover such topics as European support and training.