

iSCurrents

Published by Intel Scientific Computers

Spring 1987



Clifford Addison, co-Manager of the Centre for Computer Science at Chr. Michelsen Institute, is shown here with the iPSC System his group is using for seismic analysis.

At Chr. Michelsen Institute in Bergen, Norway, Clifford Addison and a group of researchers are using the Intel iPSC™ System to reduce the time for seismic data analysis from days to hours. Their ultimate goal is a truly interactive seismic processing system.

Determining underground geologic structures from reflected seismic energy patterns is a complicated, multi-stage process requiring expert intervention at critical stages. Until now, processing times of up to several hours per stage have forced experts to function inefficiently in batch mode.

Addison and his group are working to change that. "With processing cycles on the order of minutes instead of hours, an expert can determine critical parameters interactively. Using the Intel iPSC System we have significantly shortened computational times for critical stages that involve expert intervention." With the iPSC System, the group has

iPSC™ System Brings Interactivity to Seismic Data Analysis

reduced the time for these stages from over 2½ hours to 5 minutes, bringing them well within the requirements for interactivity.

Two Critical Stages

Addison and his colleagues have focused their efforts so far on two crucial stages: *velocity analysis* and *normal moveout (NMO)*. "Both velocity analysis and NMO have a high degree of inherent parallelism," said Addison. "Using the iPSC System, we have been able to exploit that parallelism extremely well."

Velocity analysis and NMO are crucial to seismic analysis. Reflection times for waves for the same common depth point or CDP (see article, p.2) vary with the distance between the shot and the receiver. Individual traces or received waveforms for a CDP are not comparable in their raw form. Normal moveout, or NMO, is a process of correcting traces for these differences in travel times so that reflections for a given layer line up. Once this is done, these traces can be stacked – summed – to improve the signal-to-noise ratio.

The mathematics required for NMO is simple, but the calculations are performed many times for each CDP. This stage is critical, since it gives a profile of the subsurface layers at the CDP and is thus the basis for all future conclusions about subsurface structure.

In order to perform NMO on a CDP, it is necessary to estimate the speed at which sound waves have traveled through the subsurface. In velocity analysis, the *stacking velocities* for each layer of the CDP are determined and the velocity function is expressed as a piecewise-linear function of the stacking velocities.

Finding just the right velocity function is an iterative process requiring the experience and judgment of a seismic expert. To aid this process, it is usual to form a correlation plot of reflection times and velocities (Figure 1). Each column of this plot is the stacked result of an NMO applied to the CDP using a constant velocity. Connecting suitable peaks on the plot gives the velocity function. Additional




Figure 1. Correlation Plot. In this plot, estimated velocities chosen by a seismic expert are plotted against reflection times. The contour "peaks" are connected to form the current estimate of the velocity function.

NMOs for this and for neighboring CDPs indicate whether subsurface layers line up properly for a chosen velocity function. After several iterations, the velocity function for this CDP is finalized and the final NMO for this CDP can be done.

Because of the amount of computation required, a full velocity analysis is only performed every few kilometers. Once adjacent velocity analyses are complete, stacking velocities for intermediate CDPs are easily

(Continued on page 2)

In this Issue of iSC

- Concurrent Computing at Chr. Michelsen p. 3
- Strong International Support from iSC p. 3
- iSC Training Adds Concurrent LISP p. 4
- User's Group Software Library p. 4
- Conferences p. 4

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

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

iPSC, iPSC-VX, and iPSC-MX are trademarks of Intel, Corporation. CCLISP and Concurrent Common LISP are trademarks of Gold Hill Computers, Inc.




Figure 2. Stacked and Aligned Traces. In this illustration, each vertical waveform represents the trace for a single CDP after NMO and stacking. At this point a picture of the subsurface structures begins to emerge. ©1985, AAPG, used with permission.

Seismic Data (Con't)

found by interpolation. Then NMO and stacking are performed on these CDPs. When traces for all CDPs have been stacked, a picture of subsurface structure begins to emerge (Figure 2).

The iPSC System Contribution

Addison's group initially sought to determine whether the velocity analysis and NMO for a single CDP could be done quickly enough to allow an expert to select and fine-tune the velocity function interactively. The main bottleneck had been computing a correlation plot, involving many NMO calculations (Table 1). With very minor overhead for data distribution, the total time for a correlation

ing nodes. CDPs for which a full velocity analysis is performed are treated specially. These CDPs are distributed to all the nodes, to balance the computing load and gain speed.

Addison's group will now determine whether this speedup can be sustained to make the entire analysis process interactive. Since a single experiment produces several hundred CDP records, only a portion of these can reside on the iPSC System at a given time. As CDPs are processed, new CDPs must be loaded so that computation can continue.

"The question is whether we can bring in new CDPs fast enough to make the system truly interactive," said Addison. If so, the

Table 1
 Single- vs Multiple-Node Computation Times (Seconds)

Function	Single Node	32 Nodes
Find node, distribute CDP	0.0	4.7
Correlation function	9574.4	299.2
Final NMO	33.9	1.3
Stacking	3.4	0.4
TOTAL	9611.7	305.6

plot and final NMO was reduced from about 160 minutes to 5 minutes, allowing an expert to make initial estimates and determine final velocity values interactively.

The key to success was the efficient distribution of data and calculations over the iPSC nodes. Each node holds the data for several CDP records and is normally responsible for the calculations for these CDPs. Neighboring CDPs are assigned to neighbor-

analysis for an entire experiment that now takes days could be performed in just a few hours. Addison has high hopes that with the iPSC-VX System that is scheduled for delivery this spring, this goal will be realized.

NOTE: Funding for the work described in this article was provided by Statoil and the Royal Norwegian Council for Scientific and Industrial Research (NTFT).

Typical Marine Experiment


In a marine seismic survey (see figure), a series of shots generate seismic energy. At each interface between subsurface layers of different acoustic impedances, some energy is reflected back toward the surface. These reflected waves are sensed by hydrophone receivers at various locations and recorded as digital traces, giving information about the location and nature of each reflecting interface. The sequence of shots and the placement of the receivers are organized so that multiple signals are recorded for the same vertical segment of the subsurface. This vertical segment is called the *common depth point* or *CDP*, and the collection of traces for a given CDP is called the *CDP record*.

In a typical experiment, data is generated along a seismic line from 30 to 50 km in length, with perhaps 40 CDPs per km. Each CDP record includes up to 50 traces, with 4800 bytes of data in each trace. A single CDP record,

then, contains as much as 240 KBytes of data, and a single experiment may generate 400 MBytes of data or more. Clearly, data is plentiful and complex.

Sources of difficulty are also plentiful. A given trace includes reflections for multiple layers and multiple CDPs; this data must be sorted by CDP before analysis can begin. The signal-to-noise ratio of each trace is typically poor, and methods to enhance the signal are required. Reflecting boundaries may not be horizontal, causing discrepancies between apparent and actual boundary locations. Signals from primary events must be distinguished from multiple-bounce reverberations, especially troublesome in marine surveys.

Methods such as signal filtering, deconvolution, and discrete time series analysis enable experts using high-speed computers to successfully analyze and interpret seismic data for both basic research and resource exploration needs.



This figure shows the placement of shot source or CDP. In an actual experiment, there are many distance relationships between sources and receivers.



Concurrent Computing at Chr. Michelsen Institute

Strong International Support from iSC

"Our goal is to help iPSC System users be as productive as possible," says Steve Sterba, Customer Support Manager for Intel Scientific Computers. "We offer a broad spectrum of post-sales support services to ensure users maximum productivity." Customer support services include:

- Initial product support — site preparation assistance, installation (with a half-day orientation session), 90-day on-site hardware warranty, and one-year software warranty.

- On-site service agreements, including an annual agreement to allow budgeting of all hardware service costs.
- Annual software support agreements that provide software updates, TIPS (Technical Information Phone Service), and technical reports.
- Training, documentation, diagnostics, and spare parts.

Strong customer support and a personal approach have been key factors in iSC's success, says Sterba. "By leveraging off of Intel's broad base of existing Customer Support offices, we offer a level of support unexceeded in the industry. European customers receive timely support from iSC specialists based out of local offices in the host countries (such as Sigmund Simonsen in the Skjetten office in Norway). In addition, Dave Moody and the staff at our center in Swindon, England provide the same level of extended hardware and software support that U.S. customers receive from the home office in Beaverton, Oregon."

Intel's worldwide Customer Support infrastructure includes 16 European offices, 80 U.S. offices, and 5 offices in Japan. For details on iSC Customer Support services in Europe, contact one of Intel's Regional Customer Support Centers, or contact Intel Scientific Computers, Beaverton, OR, 503-629-7777.

At Chr. Michelsen Institute near Bergen, Norway, computer scientists are pursuing two computational goals: to support the Institute's many resource-related projects and to become Scandinavia's leading center of expertise in parallel computing. These twin goals are not unrelated.


As western Norway's largest city, Bergen directly faces the North Sea — and 80% of the North Sea oil and gas reserves. Bergen has nicknamed itself the "Undersea City" because of its commitment to offshore research and development. In addition to major research, engineering, and contracting firms, the city also hosts the international Underwater Technology Conference, held semi-annually.

The Chr. Michelsen Institute, located just outside Bergen, was founded in 1930 as Norway's first independent research institute. It now occupies a 9,000 m² facility and employs a staff of 175 people, of whom 130 are attached to the Department of Science and Technology. Of these 130, more than half are graduate scientists and engineers.

CMI established the Centre for Computer Science in November of 1984 to provide computational support for the Institute's many research projects. The Centre also conducts computer research in several areas, including numerical algorithms, advanced computer architecture, image processing, and computer graphics.

Leaders at the Centre soon recognized the importance of concurrent processing in handling the large data bases and computation-intensive processes in offshore research. Their commitment to concurrent processing included the purchase of the first Intel iPSC system to be

(Continued on page 4)



s and receivers around one common depth point CDPs, each one determined by the time and receivers.

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

Institute (con't)

delivered outside the United States. The 64-node iPSC System was delivered in December, 1985, and is the Centre's primary machine for research into concurrent processing.

"We are convinced that distributed-memory concurrent processing is the only practical and cost-effective solution for offshore-related computational needs," stated Clifford Addison, co-Manager of the Centre. "We have found the iPSC System a valuable tool in developing a concurrent processing environment to meet those needs."

Researchers at the Centre are exploring the iPSC hypercube topology and distributed memory architecture as a tool in applications as diverse as seismic processing (see article), cardiac analysis, and computer-generated holograms. But the background theme for most of the work in concurrent processing is the heavy computational requirements of the Institute's seismic and resource projects.

Conference Announcements

The following conferences scheduled for this spring and early summer provide opportunities to share and learn about new developments in supercomputing and concurrent processing, including applications in chemistry, aerospace, and artificial intelligence.

- **International Conference on the Impact of Supercomputers in Chemistry.**
April 12-16, 1987. Sponsored by the University of London, England. Conference to be held at the University of London Computer Center. For information, call Dr. Janet Altmann, University of London Computer Center, 01-405-8400 Ext. 358, London, England.
- **AEROSPACE 87: AIAA 1987 Annual Meeting and International Aerospace Exhibit.**
April 28-30, 1987. Sponsored by the American Institute of Aeronautics and Astronautics. Conference to be held at the Hyatt Regency Crystal City, Arlington, VA. Call AIAA, 212-581-4300, for information.
- **Second International Conference on Supercomputing.**
May 4-8, 1987. Conference to be held at the Santa Clara Convention Center, Santa Clara, CA. Call Drs. Lana or Steven Kartashev, 813-866-2694, for information.
- **Sixth National Conference on Artificial Intelligence.**
July 13-17, 1987. Sponsored by the American Association for Artificial Intelligence. Conference to be held at the Seattle Center, Seattle, WA. Call AAAI, 415-328-3123, for information.

iSC Training Adds Concurrent LISP

The iSC training program now includes classes in CCLISP™ (Concurrent Common LISP™) announced iSC Senior Training Specialist Victor Jackson. The week-long training workshop includes:

- Overview of the iPSC-MX hardware and software architectures
- Introduction to logic programming on the iPSC-MX System
- Message passing schemes in Concurrent Common LISP
- Survey of concurrent AI topics such as cooperating expert systems and distributed knowledge bases

The class is taught on the Beaverton, Oregon campus, in conjunction with Gold Hill Computers, Inc. of Cambridge MA, the developer of CCLISP. Attendees are expected to know LISP programming, but no knowledge of concurrent programming is needed. "Anyone interested in concurrent LISP programming in a distributed memory, message-

passing environment will find the class a valuable introduction," said Jackson. LISP is the most popular language in the United States for AI applications and research.

In addition, the regular Programming Concurrent Computers workshops continue to be offered, also on the Beaverton campus. These five-day classes offer a three-day introduction to concurrent programming on the hypercube architecture, plus a two-day segment of advanced topics such as vector programming.

"Together, these two intensive, hands-on classes give a comprehensive introduction to concurrent numeric and symbolic programming. They also provide a practical way to come up to speed quickly on the iPSC machines," added Jackson. Since their inception in August of 1985, the iPSC training classes have graduated more than 120 people, including attendees from Germany, South Korea, Great Britain, and Norway.

Spring and Summer Training Schedule

**Programming
Concurrent
Computers**

April 6-10
May 4-8
June 8-12
July 6-10
August 10-14

**Concurrent
Common
LISP
Programming**

April 13-17
May 11-15
June 15-19
July 20-24
August 17-21

For additional information, contact Intel Scientific Computers, Training Center, 15201 N.W. Greenbrier Parkway, Beaverton, OR 97006, 503-629-7629.

User's Group Offers "Cubelib" Software

The iPSC Users Group software library, "cubelib", is up and running, says Tony Anderson of iSC, who is responsible for operation of the library. "There are many parallel software tools available or well under development," stated Anderson. "The purpose of cubelib is to give iPSC users a way to share the tools they've developed."

Cubelib now includes packages for matrix and eigenvalue problems, graphics, fluid flow simulation, and peptide energy minimization, as well as utility routines and customer support and release notes. Several other packages will soon be available.

Software librarian Mike Ess gave guidelines on how to use the library:

1. First connect to the library via the appropriate network and address. Access over most networks will be similar to one of the following:

```
Internet or CSNET:  
cubelib@isc.intel.com  
ARPANET:  
cubelib%isc.intel.com@relay.cs.net  
UUCP:  
[your site routing]logvax!intelisc!cubelib
```

2. Obtain an index of currently available software:

```
mail [network address] <'send index'
```

3. Obtain an index of your particular area of interest; for example:

```
mail [network address] <'send index from eiscube'
```

4. Request the software you want; for example:

```
mail [network address] <'send all from eiscube'
```

Software submissions from iPSC users are strongly encouraged. Contact Mike Ess, 503-629-7761, for details.

In Future Issues . . .

. . . of *iSCurrents* you'll read about new developments in iPSC hardware and software, including "Sugarcube" - a low-cost entry into concurrent computing, plus advances in operating systems, compilers, and cube utilization tools. We'll also feature applications of concurrent computing in areas such as AI, modeling and simulation, and image processing.