eliminate the tracking of graduate student comments, trends, and demographic statistics, among other information afforded to individual fields. That information is vital to student perceptions and institutional self-evaluations, and should be maintained, he said.

"We can't forget about oceanography and I don't think they're going to do that," he said. "On the contrary, recognizing that [oceanography] is part of a larger field is a good thing."

Comments on the draft—a preliminary step toward the 2005 report—are due by 1 March. Feedback and discussion about the draft taxonomy can be sent to resdoc@nas.edu.

—JONATHAN LIFLAND, AGU Science Writer

GEOPHYSICISTS

PAGE 31

In Memorium

Melvin Calvin died last year. He was a retired life member and AGU Fellow (Volcanology, Geochemistry, Petrology) who joined in 1961.

Reverend Keith W. Johnson died on 11 September 2002, at age 78. He was a retired life member (Atmospheric Sciences) who joined AGU in 1965.

Hsi-Ping Liu died in 2001. He had been an AGU member (Seismology) since 1974.

Reginald E. Newell died on 29 December 2002, at age 71. He had been an AGU member (Atmospheric Sciences) since 1961.

Garald G. Parker died last year. He was a retired life member (Hydrology) who joined AGU in 1940.

Herbert R. Shaw died last year. He was an AGU Fellow (Volcanology, Geochemistry, Petrology) who joined in 1968.

Nelson W. Spencer died last year. He had been an AGU member (Aeronomy) since 1959.

Neil B. A. Trivett died last year. He had been an AGU member (Atmospheric Chemistry) since 2000.

Honors

John T. Wasson has been selected to receive the J. Lawrence Smith Medal from the National Academy of Sciences "for important studies on the classification, origin, and early history of iron meteorites and chondritic meteorites, and on the mode of formation of chondrules."

Wasson is professor at the Institute of Geophysics and Planetary Physics and Departments of Earth and Space Sciences and Chemistry and Biochemistry, University of California, Los Angeles. He is an AGU Fellow and has been a member (Planetology) since 1961.

FORUM

Supercomputing Moves to Universities and Makes Possible New Ways to Organize Computational Research

PAGES 30, 33

Too cheap to meter? Surely not. But microprocessors have come a long way in delivering inexpensive computing power.

The development is eroding the distinction between computers made from low-cost microprocessors, and more conventional supercomputers operated—for example—at the U.S. National Science Foundation's Supercomputer Centers. It is also shifting the balance in high-performance computing (HPC) away from the NSF's Supercomputer Centers, to powerful, local computer facilities run by academic departments and universities.

These local facilities tend to have user communities with well-defined computing needs, often centered around a few key applications, such as climate modeling or "ab initio" material simulations. In other words, their computing needs are topical. Topical supercomputing is unusually effective. It holds the promise of increasing our modeling capacity substantially, as more and more "grand challenge" calculations are performed locally.

Since genuine supercomputing resources—defined to lie within 10% of the NSF's top supercomputers—can now be purchased for well under $1 million, new and innovative forms of organizing scientific computing should be explored. For example, dedicated support for software development and hardware acquisition should be tailored specifically around individual scientific communities. Likewise, faculty who rely on HPC resources for their research and teaching should be supported by local campus facilities. The resulting "democratization" of numerical modeling should serve large and small institutions alike, and it should go a long way toward achieving a true integration of education and research in the geosciences.

Microprocessors on the Rise

The speed of microprocessors has improved dramatically over the past 10 years. Measured in million floating point operations per second (Mflops), processing power increased by more than two orders of magnitude, from about 1 Mflops on Sun Microsystem's venerable SPARC 1 in the early 1990s, to about 500 Mflops on Intel or AMD processors today. Measured in terms of price-performance—that is, processing speed per dollar—the development is even more breathtaking. On that account, today's microprocessors exceed their earlier cousins by a factor of 10,000. In other words, since the early 1990s, the cost of 1 Mflops plummeted from about $10,000, to less than $1 today.

There is deeper significance to this leap forward. It stems from our ability to harness hundreds of microprocessors at once through a parallel programming technique known as "active message passing." Developed in the late 1980s, the method relies essentially on e-mail to communicate among a group of individual programs. Comparing the technique to e-mail is, of course, an analogy. But it helps us to better understand how groups of programs can tackle demanding computational tasks in a coordinated way across computer networks. The key is that the messages carry information from one program to the other, while providing at the same time an effective way to synchronize the entire group; hence the term "active" message passing. Message passing programs require a considerable development effort. But what makes it worthwhile is that they perform exceedingly well across networks of computers, and parallel efficiencies of better than 90% are not uncommon.

The development bodes well for Earth scientists, who have long yearned for affordable parallel computers to satisfy their insatiable appetite for computing power. Hunger for massive computation is, of course, a consequence of the natural systems we tend to model. These systems are often characterized by vastly different time and length scales, which demands the simultaneous resolution of physical processes over a wide range of temporal and spatial scales. Most of us are familiar, for example, with the oceanographers' need to resolve narrow western boundary currents, such as the Gulf Stream, in global models of the general circulation. Nowadays, hydrologists, space physicists, mineral physicists, geodynamicists, and seismologists are joining their lot with sophisticated groundwater models, simulations of solar wind-[=magnetosphere interaction, first principles calculations of matter under exotic conditions, fluid dynamic simulations of mantle and core convection, or numerical studies of seismic wave propagation, all of which consume vast computing resources.

Responding to growing computing needs seems easy enough: network a group of inexpensive PCs into a so-called "cluster" and run message-passing programs across the system.

There is a term for this innovative approach to parallel computation that has struck a cord with the HPC community: "Beowulf" clustering, named after the original Beowulf cluster project initiated by Jim Fisher at NASA's Goddard Space Flight Center some 10 years ago. A growing number of Beowulfs is now in use. At Princeton's geosciences department, the Instrumentation and Facilities program of NFS's Earth Sciences Division funded a prototype machine for large-scale geophysical modeling in 1998. The cluster with 136 Pentium II processors, fittingly named "Geowulf," rivals far more expensive supercomputers. Its computational speed, for example, matches Silicon Graphic's Origin 2000, a more traditional supercomputer operated at NSF Supercomputer Centers.

Many more geophysical modeling clusters have come online in recent years. In 1999, the earth sciences department of the University of Liverpool in the United Kingdom installed the Networked Earth Sciences SuperComputer (NESSC), which has 260 Pentium III processors. In 2000, the California Institute of Technology's Seismological Laboratory built a cluster with 156 dual-processor nodes, while the departments of geosciences and astrophysics at the University of California at Santa Cruz opened a 132 dual-processor cluster in 2002.

Today, the list of the top 100 supercomputers includes 34 clusters, with two PC clusters ranked among the top 10 most powerful machines (www.topclusters.org). But perhaps more important than this widely publicized advance of clusters into the top tier of supercomputing is their impact at the lower end of HPC. Only a few years ago, many smaller-scale calculations would require the resources traditionally reserved for National Computer Centers. Today, however, with bare-bones PCs selling for around $1000, a 16-32 processor cluster can be purchased for about $20,000–40,000, a price tag well within reach of individual research groups.

Topical Computers: One Size Does Not Fit All

Beowulfs work best when they are targeted at a few key applications; that is, when their usage is topical. The topical approach to supercomputing is unusually effective. It is flexible enough to tailor a machine's processors, network, memory, or disc space specifically to the application at hand. But it also works well because it is relatively unburdened by bureaucratic overhead. Seen in this light, clusters are an attractive supercomputing platform, and one that provides computer resources of surprising magnitude. Simply put, a Beowulf is a computing block grant in line with NSF's largest supercomputer allocations. Taking a 256-processor Beowulf as an example, 256 x 24 x 365 = 2,242,560 CPU hours are available each year.

By comparison, NSF's highest class of supercomputer awards, a "meta-allocation," provides about 100,000–1,000,000 CPU hours per year, or about 5–50% of this amount. While NSF's supercomputer awards must be renewed each year, Beowulf resources are available over the lifetime of the machine, so the simple comparison in terms of raw CPU hours becomes even more striking. More important, the comparison is relevant because today the majority of commercial supercomputers is based on the same off-the-shelf microprocessors that form the core of Beowulfs. Thus, the actual computer power per CPU hour is, in fact, quite similar across all machines.

Topical supercomputing has a long and successful tradition in the Earth sciences. NCAR's and NOAA's computer facilities are topical in nature, with technical staff and hardware dedicated to support atmospheric and oceanographic researchers across the country. What is new, however, is that this highly effective approach to supercomputing can now be exploited by far smaller scientific communities, as well as by academic departments and universities seeking to enable ambitious computer modeling efforts at their institution.

Organizing Next-generation Community Modeling Infrastructures

Computational geodynamicists comprise a small scientific community that has drawn on HPC resources for many years. To explore how this community would take advantage of the new opportunities provided by topical supercomputing, Mark Richards and Peter Olson organized a 2-day workshop in Lake Tahoe in July 2002. Funded by the NSF's geophysics program and scheduled immediately after the 2002 "Study of the Earth's Deep Interior" meeting, the workshop brought together more than 30 Earth and computational scientists in fields ranging from planetary science, seismology, and geodynamics, to numerical analysis, visualization, and computer science. The workshop's main focus was to discuss the basic framework necessary to develop and maintain a topical computational infrastructure for geodynamicists that could serve the entire community.

The case for organized computing frameworks is strong in geophysics. While PC clusters are cheap today, the development cost for novel parallel programs that can utilize this hardware efficiently has remained stubbornly high. In fact, state-of-the-art geophysical modeling software, like modern convection codes and seismic wave propagation models, would require more development time than all but the most ambitious Ph.D. research projects would allow. This makes agency support for dedicated software development and maintenance by senior numerical analysts essential for future progress. Consequently, the workshop participants envisioned software support by senior numerical analysts and computer scientists for the development of unified programming standards, common model interfaces, community modeling codes, and powerful new visualization tools akin to the generic mapping tool (GMT). The participants also envisioned dedicated hardware support in the form of frequent upgrades of powerful community modeling clusters, to keep pace with the rapid, 3-year development cycle of computing hardware.

There is great motivation for addressing the substantial difficulties involved in developing modern parallel modeling tools for geophysics. Off-the-shelf PC clusters now allow, in principle, sufficient spatial resolution for modeling the full plate tectonic/mantle convection system, and promise breakthroughs in understanding the dynamics of Earth's core. Further advances anticipated in the next few years should allow us to model seismic frequencies of 1 Hz or less globally in numerical wave propagation codes, thus holding the potential for rapidly unifying the study of a range of complicated solid Earth processes.

At the same time, ambitious new data acquisition initiatives, such as USArray, will gather information on Earth's interior in far greater detail than ever before. These new initiatives will necessitate modeling efforts, directed for example at efficient data assimilation methods, on a massive scale. Such efforts will be difficult to meet at reasonable cost within the traditional infrastructure of the NSF's Supercomputer Centers. Arguably, the most cost-effective approach is to support a number of comprehensive modeling frameworks aimed at efficiently utilizing low-cost topical Beowulf clusters.

Japan's Earth Simulator

One computer in the world is about 10 years ahead of Beowulf machines. This remarkable computer, the EarthSimulator (ES) in Japan (www.es.jamstec.go.jp), has 640 8-processor compute nodes, for a total of 5120 processors. Each node has 16 GB of shared memory, for a total of 10 TB of memory. The peak performance per node is 64 GFLOPS, and the total peak performance is 40 TFLOPS, about a factor of 100 larger than what can be accomplished today on a large Beowulf. Access to the ES is still highly restricted, and computer resources are allocated based on strict parallelization and vectorization benchmarks. Recently, we were able to test software on the ES.

Jeroen Tromp, in collaboration with Seiji Tsuboi and Dimitri Komatitsch, for example, has been running a global seismic wave propagation code on the ES. On a 2-year-old Beowulf with 75 dual processor nodes and 75 GB of memory, one can simulate global wave propagation at periods greater than 18 s, and a typical simulation lasts 48 hr. On 48 nodes of the ES, one can model periods of 8.75 s, and a simulation lasts 10 hr. On 432 ES nodes, one can perform simulations at periods of 3.75 s, and a simulation is expected to last about 24 hr.

To put these numbers in perspective, typical normal-mode summation codes that calculate semi-analytical synthetic seismograms for one-dimensional Earth models are accurate up to 6 s. In other words, the ES allows us at this point to simulate global seismic wave propagation in fully three-dimensional Earth models at periods shorter than current seismological practice for one-dimensional spherically symmetric models.

In our view, this experience suggests a dual strategy for NSF supercomputing over the next few years. On the one hand, modeling capacity is probably best served by powerful local supercomputers based on cost-efficient Beowulf clusters that can be upgraded frequently to keep pace with the rapid advance of modern microprocessors. On the other hand, NSF centers should focus on enabling modeling capabilities that can only be met on ES-class machines. The necessary advances in geophysical modeling software, including seismic wave propagation codes and geodynamic forward and inverse models, are already being developed to assist this effort.

—Peter Bunge, Princeton University, N.J. and Jeroen Tromp, California Institute of Technology, Pasadena