The Always-Connected World and Its Impacts on Ocean Research
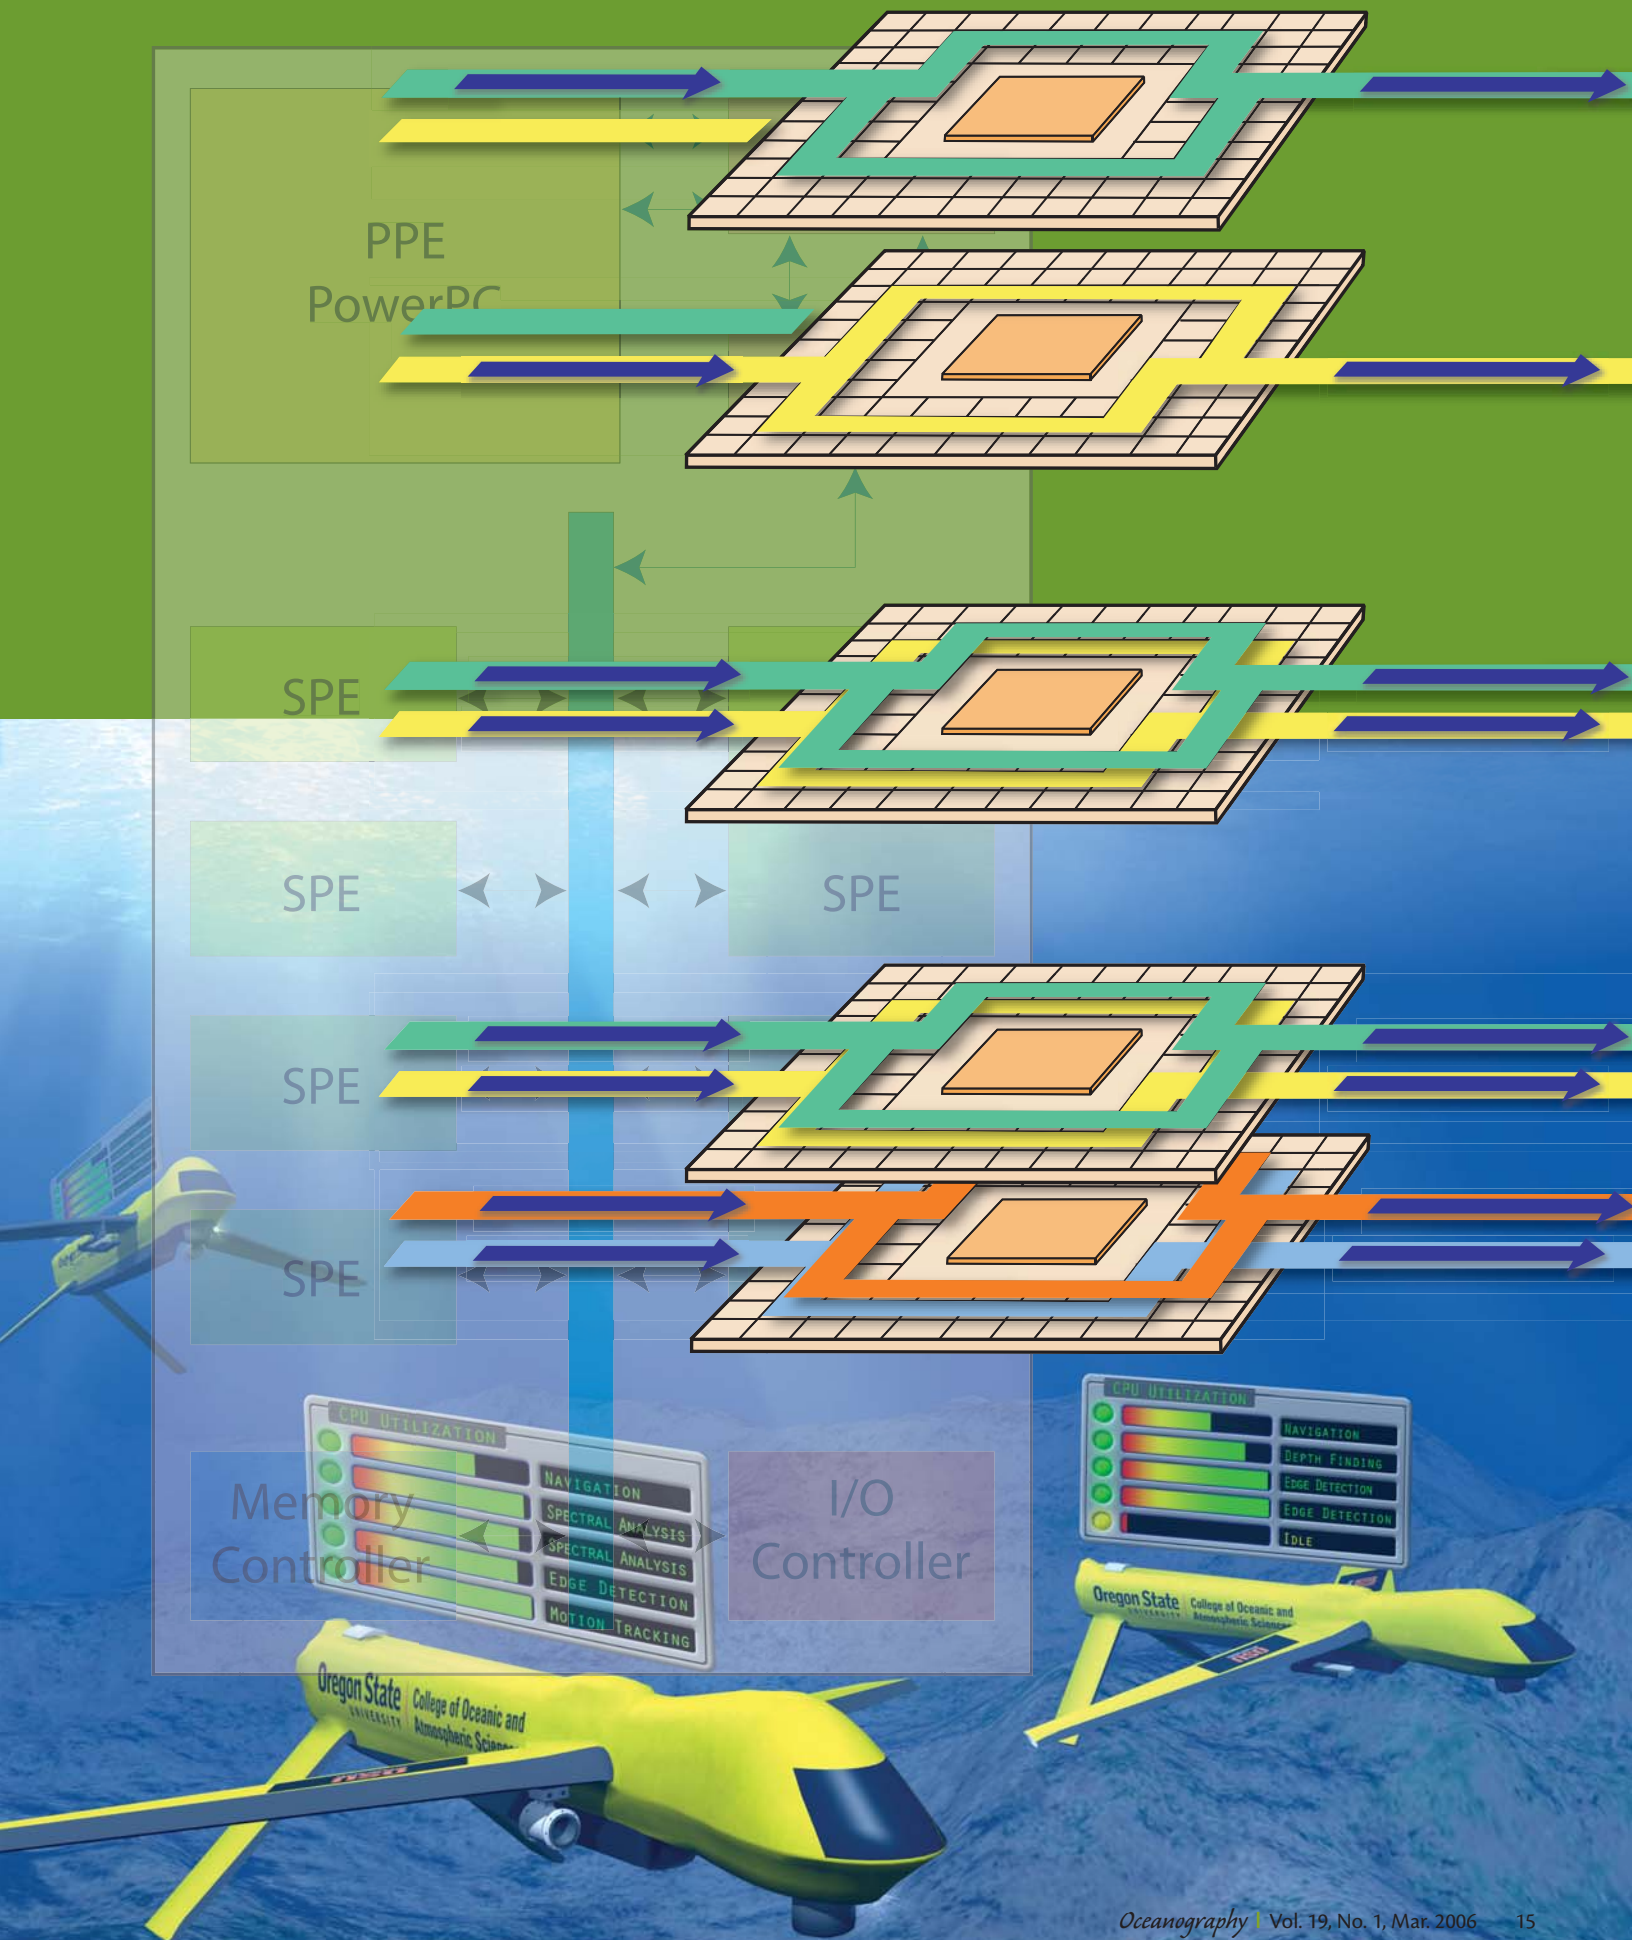
BY MARK R. ABBOTT AND CHARLES E. SEARS

The appearance of personal computers in the mid-1980s transformed the conduct of oceanographic research in ways that are now familiar as well as in ways that are now only beginning to be perceived and understood. The first Macintosh computers and IBM personal computers (PCs) were marketed primarily to the technical and educational communities, which eagerly embraced them. More powerful, UNIX-based scientific workstations also began to appear at about the same time, supporting more complex data analysis and even numerical modeling tasks. Within a few years, the reliance on academic computing centers to support modeling and data analysis had largely disappeared. The emphasis was on local, desktop computing supported by networked access to centrally managed services such as high-volume data storage, printing, email, and high-end computation. The inexorable march of Moore's Law (the doubling of the number of transistors on integrated circuits—a rough measure of computer processing power— every 18 months) brought more computational power to the desktop, blurring the distinction between the PC and the UNIX workstation in terms of performance. Network bandwidth in local area networks also increased rapidly, going from 10 megabits/s to 10 gigabits/s in 15 years (roughly a doubling every 18 months, equivalent to Moore's Law for the transistor

count). Disk capacity initially increased at a slower rate, though for the last several years it has increased faster than Moore's Law, doubling nearly every year. Along with dramatic drops in price, these forces have led to the now-familiar proliferation of computational, storage, and display capacity that was once the sole province of sophisticated computer centers at government laboratories and universities.

Software followed a similar path of growth in capability and decrease in cost. Of particular note to the research community is the expansion of interpreted languages, such as Matlab® and Mathematica®, which allow the researcher to construct sophisticated processing streams using simplified language constructs, standard building blocks, and functions. For example, multiplying two matrices is one line of code rather than a complex set of nested loops.

For oceanographers, as for scientists in general, these capabilities have greatly simplified the tasks of data collection, analysis, presentation, and publication. Moreover, PCs and workstations could be taken onboard ship to control data collection and data analysis in near real time. Oceanographers could conduct sophisticated analyses at sea and use the results to refine or revise sampling strategies to study specific processes and features of interest.

PPE
PowerPC

SPE

SPE   ◄   ►   ◄   ►   SPE

SPE

SPE

Memory
Controller

I/O
Controller

CPU UTILIZATION

NAVIGATION
SPECTRAL ANALYSIS
SPECTRAL ANALYSIS
EDGE DETECTION
MOTION TRACKING

CPU UTILIZATION

NAVIGATION
DEPTH FINDING
EDGE DETECTION
EDGE DETECTION
IDLE

Oregon State UNIVERSITY | College of Oceanic and Atmospheric Sciences

The development of the World Wide Web in the early 1990s, followed by the development of search engines such as Google™, transformed how we publish and search for information. Online access to journals has made it much easier to find articles of interest with simple keyword searches. Results from research cruises or real-time ocean observatories can be placed on the web soon after the data are collected and processed. The web and our expanded communications infrastructure have greatly enhanced our abilities for collaboration as well as outreach to the broader community.

Some aspects of oceanography have only been touched lightly by the microprocessor revolution. Ocean instrumentation needs a level of ruggedness and power consumption that are far more stringent than the typical microprocessor. Thus, ocean instrumentation has not benefited as greatly from the improvements in the price:performance ratio. Embedded processors are being developed for a wider range of consumer technologies such as cell phones, portable video and music players, and other personal devices, although adapting these processors to the specialized needs of oceanographic sensing systems can be difficult and expensive. Communication through the nearly electromagnetically opaque medium of the ocean makes networking extremely difficult for ocean-sensing systems. Thus, the impact of TCP/IP and nearly ubiquitous networking has been limited on the sensing side of oceanography.

High-performance computing (HPC) in oceanography has received considerable attention in the past five years (OITI Steering Committee, 2002). There has been increasing emphasis on new programming models (e.g., message passing, parallelization) to take advantage of rapidly evolving hardware architectures rather than just searching for the fastest piece of available hardware. The role of supercomputer centers is shifting as hardware resources that were once only available at the centers become more affordable to individual users or workgroups. With the appearance of Microsoft in the HPC marketplace, it is clear that we have entered the era of HPC as a commodity. The acquisition of a machine with a few teraflops of performance, a few terabytes of memory, and several terabytes of disk is not an insurmountable obstacle. Tasks that used to run only at supercomputer centers can now be accommodated on such machines. Moreover, the notion of computation as a utility is leading many vendors to develop large computational systems that are rented out to users on a time-share basis (which was in vogue forty years ago). The keys to success in this utility approach are the services that are delivered, not just compute cycles.

These changes in technology appear to have changed the way we conduct ocean research. But, their primary impact has been to speed up the individual components of the information collection, analysis, modeling, and distribution process. The fundamental workflow of oceanography looks much the same now as it did thirty years ago when computer centers ruled the academic computational landscape. Researchers collect data into files, return them to shore for analysis, share the data with colleagues, use the data to drive models (numerical, statistical, or descriptive), create graphs or images to visualize the results, and publish papers in journals. Whether the data were collected on strip charts for later digitization or collected on flash memory cards for analysis on a shipboard PC, the basic data and information flows are the same. The changes have been in the volume of data that can be collected, the speed with which data and information can be distributed, the complexity and time/space resolution of numerical models, and the sophistication of graphics and animations. Although these technology trends have extended our abilities to do ocean research, new challenges and

Mark R. Abbott (mark@coas.oregonstate.edu) is Dean, College of Oceanic and Atmospheric Sciences, Oregon State University, Corvallis, OR, USA. **Charles E. Sears** is Research Computing Manager, College of Oceanic and Atmospheric Sciences, Oregon State University, Corvallis, OR, USA.

> As oceanographers, we need to reinvest in training and testing so that we are not simply passive consumers of whatever IT tools appear in the market.

opportunities have emerged.

Our workflow is only loosely linked, with observations generally not integrated with modeling systems. Our information systems are generally put together in an ad hoc manner to meet the needs of a particular project and are thus often not enduring. Discerning and understanding the relationships among multiple data sets is challenging, especially in a quantitative manner. Interdisciplinary collaboration is increasingly important, yet it is difficult to describe data sets in ways that are understandable to a broad range of scientists, let alone between physical and social scientists. Visualization and analysis tools are not linked with the underlying observations, so tracking the various transformations and analyses is difficult. Our processes served us well when observations were made within the context of a cruise or experiment, when the primary focus was on a single discipline question, and when the results of the study were often not needed rapidly. With increasing awareness of the importance of the ocean to society (e.g., US Commission on Ocean Policy, 2004), oceanographers are being asked to answer questions, such as "how will climate change affect the patterns of upwelling winds in eastern boundary currents and how will this impact coastal ecosystems?" Such complex issues will require a high level of data and scientific integration. Our information systems must be designed to support these functions. Past designs of information systems have often focused on the implementation requirements (How much disk space do I need? How fast does my model need to run?), and although such questions will continue to require analysis and plan-

ning, more focus will be on the functional requirements (What types of services does the system need to deliver?)

A report on emerging trends in information technology (Ocean ITI Working Group, 2004) documented these issues and noted that the oceanography community needed to take a new look at workflows in relation to new technologies that are now network-centric. Ubiquitous networks now allow components to be connected that formerly were isolated, such as sensors. This connectivity does not simply mean that we now have digital communications where once we relied on physical transport of data (such as disks or strip charts), but rather we now can conceive of cooperating systems of components that can do far more than just move data from one component to the next. With increasing "intelligence" now available at the fringes of the network (e.g., in sensors, data storage devices, computational elements), the vision of an adaptable, dynamic network that enables a broad range of new services is becoming a reality. (See Box 1 for a description of the impacts of new CPU designs on sensors and decision support services.) Thus, our conventional approach of designing a point solution to an information technology problem is increasingly irrelevant, and oceanographers must develop robust information models that capture the functional requirements of the services they wish to deliver. The hardware, while important, is less critical than knowing the functions that need to be performed and the relationships among system components. Such an approach will be better able to take advantage of new capabilities that emerge on the network. Unlike the

approach to IT (Information Technology) that was in vogue twenty years ago, oceanographers can no longer rely on a computer center or a single vendor to deliver a solution. Researchers must keep abreast of modern tools and processes, which is especially challenging in an era of rapid change. Wilson (2006) notes that this has not been the case, especially in software development. There is a price to pay for having capable and adaptive IT infrastructure.

For an example, consider computer security. Ten years ago, the accepted approach was to establish a firewall around the critical pieces of the network, but allow free movement among components inside the firewall. It is now clear that such systems do not fail gracefully; if someone gains access from the outside (or a disgruntled person on the inside), then it is a simple matter to wreak havoc on the network. Modern approaches rely on a layered approach of authentication and access policies to reduce the chance of catastrophic failure from a single security breach. This approach requires a clear understanding of network resources and functions so that appropriate access policies can be developed. It recognizes that network-centric systems need more dynamic security systems than a single point solution.

As the oceans community begins to develop and deploy real-time observing systems, whether of moorings, gliders, or coastal radars, it will need to develop new workflow models to use the real-time, distributed nature of these systems. These observatories will evolve and the types of services they provide will evolve as well, extending their reach to non-research uses. The development of use-case

New processor technologies will serve as a specific example of the changes confronting the ocean research community. Since the early 1970s, computer-processing rates have increased substantially, supporting the development of new ideas and spurring innovation that incorporates these new capabilities on an annual technology lifecycle. Key to this increase in compute capacity has been the continued success of Moore's Law, which states that the transistors used in microprocessors and memory chips continue to become faster over time. In addition to Moore's Law, more transistors can be incorporated into an individual CPU through advances in today's chip fabrication processes.

Software engineers have viewed a computer as a single processor executing a stream of sequential instructions that fetch data stored in memory. Designers have optimized chip designs by implementing innovative ways to increase the number of instructions issued on every clock cycle even as the processor clock frequency increases. Early designers realized that to do so would require increased complexity in their respective architectures. Thus, the trend was to improve execution in existing architectures by increasing on-chip memory caches and improving instruction pipelines. In doing so, today's superscalar processors were born. They are capable of executing more than one instruction in each cycle by fetching more than one instruction from cache, then determining if more than one instruction can be executed at a time. An emergent property of these superscalar designs is the use of instruction-level parallelism to improve the performance of these architectures while maintaining an in-order, sequentially executed instruction path. This approach has allowed programmers to follow essentially the same programming model (see figure to left).

Today, we are seeing new microprocessor trends that will shift the way we design, deploy, and support next-generation information technology solutions. Continued advancement around improvements in instruction-level parallelism, pipelines, and superscalar designs requires additional transistor counts, which require more power and the associated advanced technologies to dissipate heat. The predictable speed increases that we have lived with for the past thirty years are slowing; microprocessor designers must find new ways of utilizing transistor counts more efficiently.

Some chip designs are moving to multiple core architectures within in a single chip, with each core running multiple processing threads. As with all multi-processor architectures, communications and latency are key determinants of overall processing throughput. These cores behave as programmable I/O (input/output) devices or attached co-processors. Each of these co-processors can be optimized for complex functions such as signal processing, video, and audio.

Some chips, such as the IBM/Toshiba/Sony "cell" chip, rely on task-specific processing elements to achieve the necessary

A. Single thread processor

Threads

or

B. Multi-thread processor

C. Dual-core, multi-thread processor
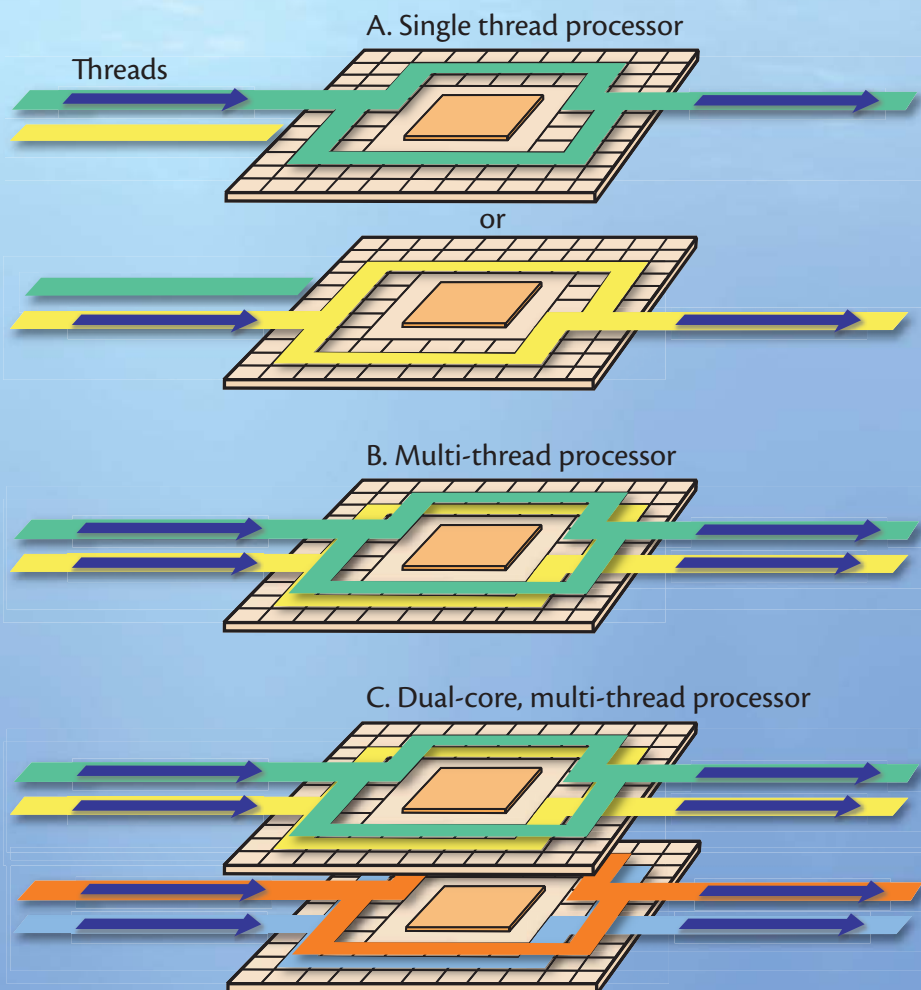
(A) *Two separate processing threads are decomposed into sets of sequential instructions. A single-thread processor must complete an entire thread before starting the second thread. (B) A multi-thread processor can accommodate more than one thread at the same time. (C) Adding multiple, multi-thread cores leads to more capable systems, although at the cost of more complex programming models.*

A cooperating network of sensors and computation could work with a more capable "control" node (Node D) to provide adaptive sampling and computation. The boxes above the autonomous underwater vehicles (AUV) represent processing units with "cell" capabilities. The units are network-aware as well as environment-aware, and they can adjust their processing systems to reflect the surrounding conditions or their scientific mission as well as draw upon the processing capabilities of neighboring nodes. The entire network would be used as part of a decision support system.

Cooperative Sensor Web

node A

node B

node C

node D

INPUT

OUTPUT

Sensed
Observations

Decision
Support System

throughput. Future versions will incorporate an increasing number of task-specific cores capable of running tens of processing threads concurrently. In a sense, this architecture does not look much different than architectures from twenty years ago where a single computer had a separate, dedicated array-processing board. In fact, many of the issues are the same, including understanding I/O from the specialized hardware and integrating optimized software libraries for the special purpose hardware within the processing flow. The challenge now is that there is greater complexity and flexibility available for user applications.

The gaming industry has taken advantage of these new architectures to significantly enhance performance. For example, the Sony Playstation 2 relied heavily on the graphics processor to run radiative transfer simulations to generate realistic imagery in a variety of atmospheric and underwater situations. The cell chip (which will be part of the Playstation 3) will similarly take advantage of task-specific hardware. However, these algorithms and tasks are stable; for example, providing realistic underwater light fields is based on well-known radiative transfer equations. In contrast, algorithms used by the scientific community are much more diverse and evolve more rapidly. Thus, the typical scientific application will find these new multi-core architectures much more challenging to use effectively.

Along with increased performance and a greater number of processing elements, next-generation chips will be inherently network-aware. Individual components will both advertise and seek out available services so that, for example, a particular process that requires specialized hardware can locate this service on the network. This capability will lead to a higher level of adaptability in the hardware/software system. Moreover, the deployment of cooperating and adaptive networks of systems will lead to new capabilities that will have a wide range of applications in the ocean sciences.

Next-generation chip architectures will be a perfect match for the Single Instruction, Multiple Data (SIMD)-like processing that is common in gaming applications, such as im-

age rendering. In this case, a simple processing kernel operates on a single data element, allowing the entire process to be parallelized. These kernels can then be linked within the chip to build up a continuously flowing data structure. These structures can be mapped onto various processing elements within the chip or operated in a data-parallel manner (figure above). However, one cannot simply port existing software applications to these new architectures and realize their potential. Effective applications of next-generation chip architectures will require a thorough analysis of the problem, subdivision of tasks, and allocation of these tasks to the programming elements. Managing the communications and scheduling of these tasks will be a challenge for the research community.
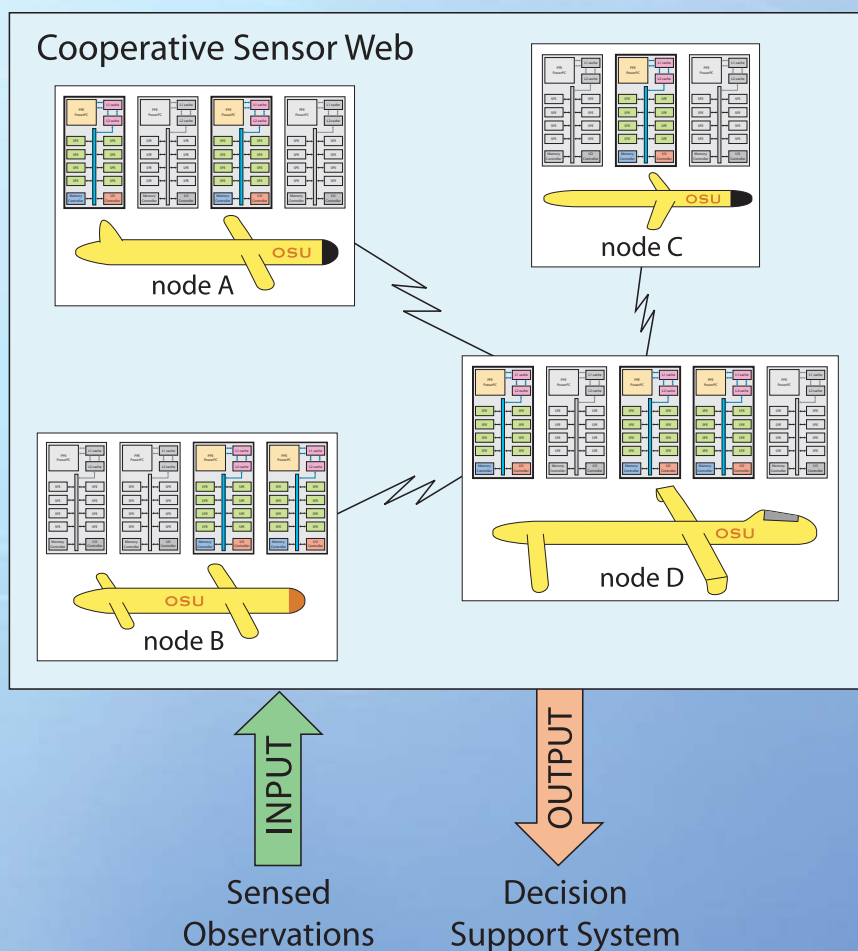
scenarios will be crucial in this context. However, simply linking together existing tools that were developed for a different set of workflows that were based on a pipeline processing model will likely be unsuccessful. The research community must begin to understand the IT capabilities that are being developed outside the science world for larger market segments. Unlike the IT world of twenty years ago, the research community is not the leading edge in regard to applications. Thus, one of the unexpected effects of the IT revolution has been the spread of IT capabilities outside the research realm, thus diminishing our impact. The second challenge is the need to recognize that interconnected networks of people, sensors, and systems require different approaches to system design.

The concept of networked collaboration to explore complex issues begins to resemble a gaming metaphor where the players share insights by examining outputs and making decisions. Although one could focus simply on the hardware and the fact that many of the next-generation chipsets are designed to handle more realistic visualization and complex games, the more interesting concepts are the uses of gaming for decision support and for real-time ocean observatories and prediction systems. For scenario development and decision support, a gaming approach would engage a broader range of researchers as well as managers and policy-makers to investigate the impacts of both natural and human-induced changes on the outcomes of various ocean processes. For example, we can imagine a coupled coastal ocean ecosystem model being forced by different patterns of upwelling winds, with

managers implementing various fisheries management regulations in the face of an uncertain environment. This is a broad definition of decision support to address issues such as, "How vulnerable are coastal ecosystems to shifts in temporal patterns of wind forcing?" The approach is based on risk assessment, not just answering a science question.

Gaming could also be applied to real-time ocean observatories and sensor webs. With sufficient local processing in each sensor or platform, we can imagine self-forming and cooperating networks designed to observe specific processes. If such a network were coupled with data-assimilation modeling, then the possibilities become much more interesting. Suppose an upper-ocean model projected the formation of a strong convergent front. This event might create a cooperative network of gliders to study the biological response, although we will also have to incorporate probabilistic behaviors in both the sensors and the models in our networks (Figure 1). As our observing systems move towards continuous observations but with discontinuous communications to the shore, such a level of autonomy and adaptability will need to be developed. Previously, we would consider networks to be composed of known resources with known services, but in the future we might have a system composed of potential resources. The sensor/model network would behave more as a community of producers and consumers where producers communicate their services, consumers communicate their needs, and the two can reach an agreement on a transaction.

There are many technical challenges to developing such communities or net-

works of potential resources. Standard ways of describing and identifying services and needs, establishing a level of trust between producer and consumer, and developing a framework to conduct the negotiations between the two parties are just some of these challenges. If our systems were unchanging over time, then this would be straightforward, but we expect these systems to evolve continuously, especially with next-generation processors that can be adaptive. Fundamentally, the focus will be on delivering services, not on data sets.

The ubiquitous-network, ubiquitous-data environment will continue to present challenges to the research community. Many of our tools and workflows are inherently designed around an environment of sparse data; wading through a sea of complex, variable data streams (as opposed to finite, discrete data files) will require new skills and new approaches. Moreover, the next generation of IT will be driven by these needs, which are not necessarily driven by science. Understanding how to leverage these capabilities will be difficult. Simply describing an IT solution in terms of flops and gigabytes will not be sufficient.

The oceans community is moving into an era where our traditional modes of data collection, analysis, and publication will be challenged. These modes have served us well and they will persist. However, we must add new modes that support a broader range of interdisciplinary interactions, including interactions with non-oceanographers and non-scientists. Our observing systems will be much more dynamic and adaptable, and the systems will be integrated into our overall analysis/publication sys-
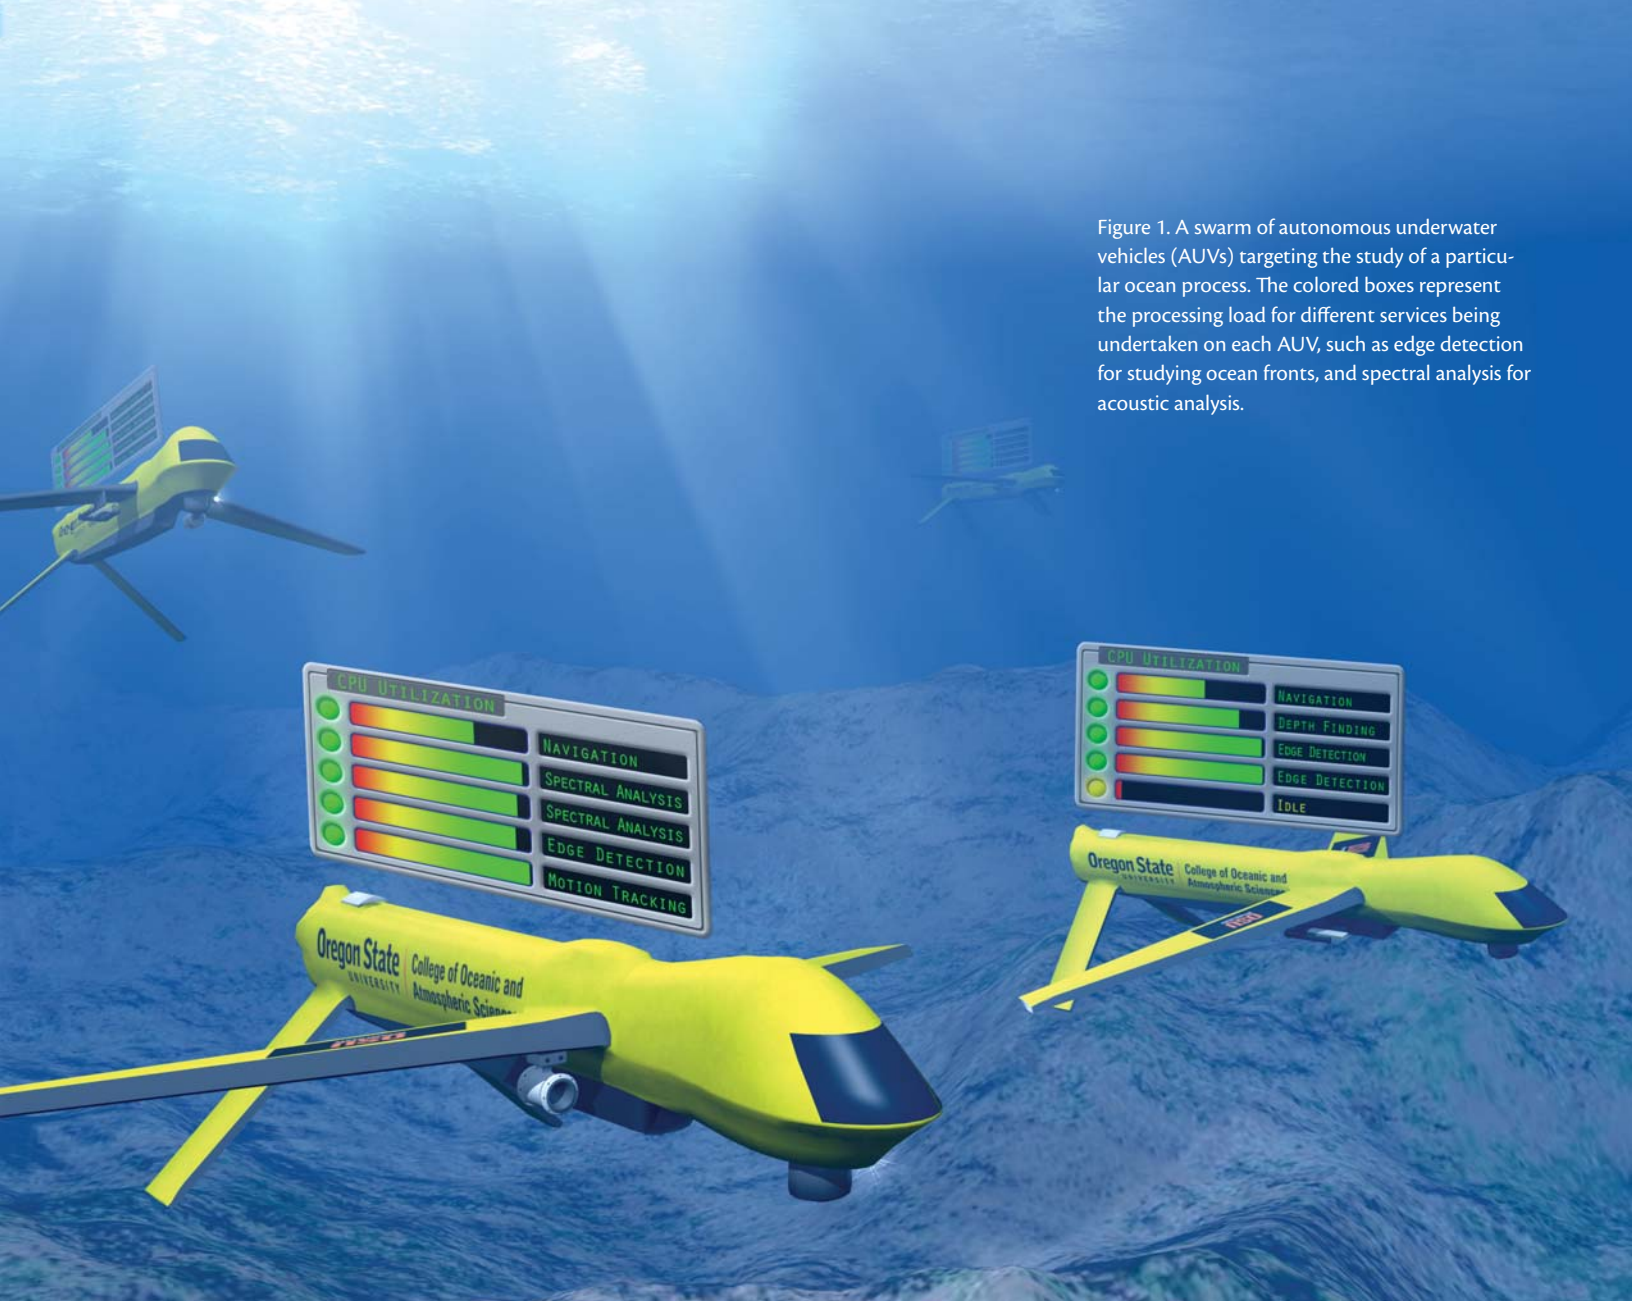
Figure 1. A swarm of autonomous underwater vehicles (AUVs) targeting the study of a particular ocean process. The colored boxes represent the processing load for different services being undertaken on each AUV, such as edge detection for studying ocean fronts, and spectral analysis for acoustic analysis.

tem. Understanding our workflows and the associated functional requirements is essential if we are to take advantage of these advances in IT. Moreover, most of these advances will not be driven by the needs of science but by other, larger segments of the market. As oceanographers, we need to reinvest in training and testing so that we are not simply passive consumers of whatever IT tools appear in the market. The next twenty years will see as many changes as we have witnessed in the past twenty.

REFERENCES

Ocean ITI Working Group. 2004. Trends in Information Technology Infrastructure in the Ocean Sciences, 24 pp. [Online] Available at: www.geo-prose.com/oceans_iti_trends [last accessed February 13, 2006].

OITI Steering Committee. 2002. An Information Technology Infrastructure Plan to Advance Ocean Sciences, 80 pp. [Online] Available at: www.geo-prose.com/oiti/report.html [last accessed February 13, 2006].

U.S. Commission on Ocean Policy. 2004. *An Ocean Blueprint for the 21st Century*. Final Report. U.S. Commission on Ocean Policy, Washington, DC., 522 pp. [Online] Available at: www.oceancommission.gov/documents/full_color_rpt/welcome.html [last accessed February 13, 2006].

Wilson, G.V. 2006. Where's the real bottleneck in scientific computing? *American Scientist* 94:5–6.