

SUPERCOMPUTING and INDUSTRY



Rod Adkins, Vice President of IBM Systems and Technology



Peter Ungaro, President and CEO of Cray Inc.

Many SciDAC applications will need access to petascale computing resources as they move forward, and it is the industry leaders, such as IBM and Cray, who are ultimately responsible for meeting that need. Peter Ungaro, President and CEO of Cray Inc., and Rod Adkins, Vice President of IBM Systems and Technology, discuss with *SciDAC Review* the role industry will play in the future of supercomputing.

SciDAC Review: *What are the characteristics of your company's planned petascale architectures that would make them attractive to the SciDAC applications?*

Cray: Cray's goal for petascale computing is consistent with the goals of the SciDAC program—to arm scientists and engineers with the tools they need to reach the next stage of scientific discovery. In order to do this, Cray is focused on developing petascale systems that provide the highest sustained performance on real applications, allowing researchers to do real work and helping them become more productive through features that will make the system easier to use.

At Cray we design all of our systems for one thing—supercomputing. We purpose-build our supercomputers to deal with the challenges of scaling applications to the petascale level. In particular, synchronization and communications among processors become more important as you scale up. The larger the system, the more impact the network and software have on the overall sustained performance. The Cray XT4 system, introduced in November of 2006, was designed from the ground-up to

scale to terascale levels of sustained performance. Both the hardware and software architectures are optimized to provide an application environment that can take advantage of tens of thousands of processors—providing more bandwidth per computation than competing architectures—and this trend will continue with Cray's future systems.

We have also taken steps on the software side to further enhance our scalable architecture. We have incorporated a Light Weight Kernel (LWK) on the compute blades of our current systems, which keeps the operating system (OS) out of the way of the target application, dedicating all of the processor cycles to the application. This creates a “jitter-free” OS that enables many thousands of processors to synchronize quickly, resulting in superior application performance.

Our next-generation system, codenamed Baker, will leverage all of these technologies and take them one step further with a new system interconnect which will support globally addressable memory and a wider variety of programming languages. This will extend the scalability and usability of our systems to support SciDAC's applications at the petascale level.

IBM: IBM has extensive experience with the wide variety of SciDAC applications, as many of the DOE science labs have been deploying our large systems in production for the last decade. Insights and experience gained from the SciDAC set of users have helped significantly influence the architecture and design of IBM's PERCS (Productive, Easy-to-use, Reliable Computer System), which is based on the Defense Advanced Research Projects Agency (DARPA) High Productivity Computer Systems (HPCS) program. The PERCS system will be the basis for the general purpose petascale systems from IBM.

IBM's solution provides a holistic approach for petascale design which includes focus on significant innovations in the core computer architecture, system design, system interconnect, and all aspects of software that are critical for petascale enablement. A number of the salient attributes of our systems will directly benefit SciDAC applications at petascale.

Our balanced system architecture and design features—well-defined proportions of compute capacity, caching hierarchy, memory capacity, memory bandwidth, interconnect latency and bandwidth, file system capacity and bandwidth—help address a broad set of supercomputing applications. The basic building block for the PERCS system will be the Power 7 processor core with significant integration that helps simplify the complexity of the petascale system environments and significantly improve performance and efficiency of these large systems.

Another important attribute is a robust, scalable, and high-performance software architecture. IBM has invested significantly over the last decade in building a robust software architecture that includes its leadership General Parallel File System (GPFS), state-of-the-art transport protocol stacks, programming models, libraries, schedulers and resource managers, and systems management infrastructure. This software architecture has proven scalability, reliability, and performance. IBM's investment in PERCS will help enhance the scaling, performance, and reliability necessary for petascale systems. The software stack is being enhanced to ensure the system is resilient and can continue to serve applications even in the event of multiple failures.

We also place a significant focus on the productivity of petascale systems. IBM is investing significantly in enhancing the programming models, newer languages, application development infrastructure, debug tools, performance tuning tools, and significant enhancement to compilers. Much of the compiler enhancements are focused on hiding the architectural complexities from the end user while enabling applications to fully realize the potential offered by enhancements in the computer architecture.

“IBM has extensive experience with the wide variety of SciDAC applications...”

Rod Adkins, IBM

For a General Purpose solution, IBM's approach is to allow for a wide variety of applications to be enabled by seamlessly leveraging the large set of existing Independent Software Vendors (ISV) applications on the UNIX platforms. This implies support of a standard UNIX operating system on which these applications run. In order to control the jitter experienced by parallel applications from a commercial UNIX OS and its associated daemons, IBM has patented technology for synchronizing daemons. In addition, IBM has lighter-weight versions of the UNIX operating system that can be run on the compute nodes to reduce unnecessary interference. In addition, there is work under way to efficiently exploit the background simultaneous multithreading (SMT) capability to hide the daemon activity at lower priority on these background threads so that they can run without interrupting the running application. This approach allows for the petascale systems to benefit from both a large breadth of ISV applications and libraries being available while effectively controlling the impact of OS jitter.

IBM is also working on two specialized petascale systems based on Blue Gene and Cell processor-based accelerators. This provides for an efficient hybrid model of supercomputers. Concepts and ideas from Blue Gene- and the Cell-based hybrids that provide significant enhancements to applications performance are absorbed into the mainstream HPCS-PERCS-based systems.

These fundamental salient features of IBM's petascale solution through HPCS-PERCS will directly impact the

SciDAC applications by significantly enhancing their capabilities to address the grand challenge problems confronting the SciDAC community.

What do you think will be the biggest challenge that users will face in using architectures with over 100,000 processors?

Cray: We expect that the 100,000 processors will be comprised of 12,000 to 25,000 quad- or octa-core sockets. While this architecture will support the current generation of Message Passing Interface (MPI)-based applications, achieving the best performance possible will generally require the user to redesign their application to use some combination of shared memory parallelism with MPI or partially (or fully) utilize partitioned global address space (PGAS) languages, such as Unified Parallel C (UPC) or Co-Array Fortran (CAF). As a vendor we must supply users with the software to allow them to effectively perform message passing within a shared memory construct. We must also continue to advance the state-of-the-art in performance tools so that users can more effectively utilize the multilevel cache on these new systems. Our focus is on productivity, which includes both portability and performance.

Users will also likely be faced with reliability issues with new petascale architectures, unless those systems are specifically designed to scale. Cray supercomputers have a dedicated hardware reliability subsystem that monitors all of the environmental variables, checks the nodes for errors, and communicates over a separate network. The system reconfigures itself if an error is uncovered. At Cray we are working on how to make systems more resilient to failures in any one component.

IBM: Scaling to over 100,000 processors will represent several fundamental challenges for the user community.

Applications will need to be designed to scale to over 100,000 processors. This will imply end users having to significantly redesign their core data structures and the algorithms used to parallelize the applications in a load balanced fashion. In order to reduce the burden on the programmer, IBM supports multiple programming models such as OpenMP and MPI to be enabled in a single program. In addition, IBM is also enabling PGAS models such as UPC. The compilers are also becoming smarter with auto-parallelization enhancements, automatic exploitation of SMT, and multicore architectures. A key observation is that the cost of memory is becoming an increasingly dominant part of the overall system. This will force the end users and system software developers to be increasingly efficient in their usage of memory in their applications and system design.

As the scale of the applications increases, there is increased complexity in ensuring that the load balancing is efficient among the tasks of a parallel application. Race conditions become more difficult to debug and avoid. Efficient synchronization becomes an increasingly dominant problem in the design of the application. IBM is working on a robust set of debugging and performance

tuning tools that will be critical in ensuring that the end users have the necessary tools for efficiently harnessing the power of these petascale systems.

Communication and OS limits must also be addressed. Many of the applications and the software subsystems depend on the Transmission Control Protocol/Internet Protocol (TCP/IP) stack which has architectural limits, such as TCP port number, that will need to be redesigned. Similarly, for file systems to scale to over 2 billion files will require fundamental architectural limits to be redesigned.

Ensuring that applications run uninterrupted on these petascale systems will be critical especially for long-running applications. Given the large number of components and parts, Mean Time Between Failures (MTBF) will be smaller than for current systems. System administrators will need the necessary tools and infrastructure to ensure that they are able to detect problems quickly and take appropriate action before they manifest themselves into bigger problems that may cause either the running applications to be terminated or for the system to be brought down. IBM is building resiliency in its hardware and software to ensure that the system has the ability to recover from transient and some permanent errors transparently. In addition, tools for rapid fault-detection, isolation, and repair will also be provided.

What specialized software will be needed to run SciDAC codes on your machines?

IBM: IBM has significant specialized software that it enables with its supercomputer solutions. Much of this specialized software is enabled transparently to the end user. The programming models utilize advanced transport protocols underneath that effectively exploit the emerging high speed networks. The GPFS file system enables standard POSIX (Portable Operating System Interface for uniX) semantics while ensuring the best scaling and performance for a cluster/parallel file system. The scientific libraries—Engineering Scientific Subroutine Library (ESSL) and Parallel Engineering Scientific Subroutine Library (PESSL)—are designed to be easily usable by end users while the underlying routines are carefully tuned for our platforms. The suite of compilers (xlc, xlf) effectively exploit new architectural constructs such as SMT multicores, among others.

Cray: Cray already provides with its systems a mature software stack that has been designed for High-Performance Computing (HPC) codes, and it supports a wide variety of programming models. This means that you will be able to take an application that today runs on a Cray XT3 or XT4 and transfer it seamlessly onto our petascale system, giving application developers the confidence that they won't need to do major rewrites with every new generation. As with our hardware architecture, these software tools have been designed for scalability. This includes fully scalable and integrated performance analysis and debugging tools that enable programmers to rapidly test and fine-tune their applications on extremely large processor counts.

Libraries are also an important component. Our current and future systems will contain high-performance versions of the most important libraries used by the SciDAC applications. Given the need to have a combination of message passing and shared memory parallelization, we will supply versions of all libraries that also use MPI and shared memory parallelization. We see this feature as an important differentiator in the HPC market.

How will compilers be used to take advantage of multicore architectures?

Cray: Cray's compilers will support both a pure MPI model across individual cores and a hierarchical model with shared memory parallelism across cores and MPI across sockets. In the hierarchical case, the compilers will be able to automatically extract parallelism within an MPI process, or take advantage of OpenMP directives. While Cray's

“Cray's goal for petascale computing is consistent with the goals of SciDAC—to provide the tools needed for scientific discovery.”

Peter Ungaro, Cray Inc.

parallelizing compiler technology is quite sophisticated, user directives and hand tuning can provide performance advantages for many codes. Cray has established several Centers of Excellence around the world to support users of our largest system, and to help achieve the shared memory parallelism needed to optimize performance at scale. One such center supports users of the Cray Leadership supercomputers at ORNL.

IBM: IBM is investing heavily in compilers to adapt to the increasing complexity of computer architectures, and to hide these complexities from the end user while exploiting newer architectural capabilities to the extent possible. Examples of this include SMT exploitation through assist threads and in intelligent pre-fetching to ensure a cache miss-free execution, development of multithreading capabilities for automatic parallelization (through OpenMP and otherwise) such that multiple cores and SMT threads can be exploited through fast synchronization constructs among the multiple threads of execution, development of advanced power optimization techniques, incorporation of barrier synchronization constructs in hardware, and the intelligent use of memory move engines.

How will application codes that use over 10,000 processors be debugged?

IBM: IBM envisions that the Parallel Tools Platform (PTP), based on Eclipse, will be the first major step in significantly



Figure 1. Mare Nostrum, the Spanish Ministry of Education & Science's powerful IBM supercomputer, is used to further research in protein folding, drug development, and climate change.

enhancing the debugging capability at scale. In addition, IBM will be supporting the parallel debugger. IBM is also investing in the development of static analysis tools that are very important for detecting potential race conditions—deadlocks that can be quite difficult to reproduce and debug at run time.

In terms of detecting data corruption, IBM is providing extensive trace utilities in all critical subsystems to help capture sufficient information to debug the root causes of difficult data corruption problems, especially with Remote Direct Management Access (RDMA). IBM also provides hooks for conditional watch-points that help monitor memory locations that are being over-written and corrupted. IBM is also providing a suite of tools that help distill and filter large amounts of trace data to useful information necessary for efficient debugging at scale. In addition, IBM has a large portfolio of Rational-based tools for detecting memory leaks and other utilities for debugging.

Cray: Petascale systems will present significant challenges for debugging. The traditional way of debugging applications will simply collapse under the weight of tens or hundreds of thousands of threads of control. With thousands of such threads, direct examination is untenable.

Under our Cascade development program, part of the DARPA HPCS program, we are developing new techniques for debugging at massive scale. Our debugger strategy has a “data centric” approach, which will allow the user to focus on the data of the program, rather than the control. This approach will include a scalable debugger manager, an introspection mechanism, and comparative debugging techniques, which will provide the functionality of comparing a properly working version of an application against a version that is not working. Data centric debugging techniques will allow users to narrow down the problem location without studying the details of the individual control threads.

What role will Open Software play in your supported software stack?

Cray: We see the Open Software community playing an increasingly important role in both Cray's and other supported software stacks, by supplying efficient high-level libraries, where the low-level kernels can be optimized by the respective vendors for their selected instruction sets. Partnerships will be needed to make this work correctly. For example, few libraries currently have OpenMP-aware versions, and this presents an opportunity for the Open Software efforts to supply the infrastructure for efficiently executing on the new multicore systems.

The Cray XT4 system uses a derivative of SUSE Linux which gives users access to all of the tools and applications available under the Open Source Linux/AMD64 environment. Meanwhile, a lightweight kernel runs on the computational nodes, allowing applications to run more efficiently. This will be the same environment as Cray's petascale system so users will be able to build their needed environment today and bring that over to the petascale systems when they are ready.

IBM: IBM supports the open source software model in areas in which it makes sense to do so, and where there is an active community of developers intent on enhancing the software for the good of the entire community. We are a very active supporter of Linux and invest heavily in enhancing Linux in ways that are beneficial to IBM customers. In addition, much of the tools infrastructure for application development is based on the open source platform Eclipse, which is itself open source.

IBM is working in concert with LANL to enhance the HPC tools infrastructure based on Eclipse. In addition, we are building an ecosystem around this platform, and expect that as the community rallies around Eclipse, the entire society will benefit.

We contribute to an open source storage management project called Aperi. IBM is also a significant contributor to the Operations Research Library (ORL) routines through the Computational Infrastructure for Operations Research (COIN-OR) project. We participate in a number of open source projects, and evaluate various options that exist in the open source model. Depending on the project and the community, we will continue to make decisions on how to enhance our participation in such open source projects. IBM believes that open source is an important tool for advancing collaboration and innovation. However, we do believe that there is still a critical role for more traditional value add software licensing, which helps overall value differentiation of our systems.

What is your approach to delivering higher performance within a reasonable power profile?

IBM: More and more of the power consumed by the compute servers is funneled into “wasted” leakage current. Lowering the junction temperature reduces this leakage current, so we are exploring direct water cooling on high-



Figure 2. BlueGene/L, among the most powerful supercomputing systems on Earth, was developed by IBM for LLNL.

power chips to reduce non-productive power. We are also exploring the appropriate balance between the highest single thread performance, which implies high-frequency and power, and a lower frequency design point, such as Blue Gene, that optimizes the Compute Power/Watt. In addition, IBM is investing in significant dynamic power management technologies in both hardware and software.

Cray: Cray has a long history as a leader in cooling technologies for its system cabinets. In the petascale system timeframe we will have integrated cooling in the cabinets where the air that exits the cabinet is virtually the same temperature as it is when it enters—this is much more efficient than emitting the heat into the machine and data center for air handlers to extract. We are also working with AMD to leverage their power-saving features and continue to look for ways to reduce the power of these ultrascale systems.

What opportunities exist for new approaches to minimize power requirements?

Cray: There will always be a tradeoff between power and performance. In the extreme, you can design a single lightweight integrated chip with low-performance processors, network, and system functionality all on the same chip. But in our opinion, this is not likely to provide the best overall performance for the wide variety of applications required by SciDAC. For broader applicability, separating these functions but providing a matched, high-bandwidth interconnect provides a higher performance system with controllable cost-benefit tradeoffs between power and performance. We are also working with different application accelerator technologies to look at the power efficiency

advantages of more specialized systems. Finally, deploying Linux lightweight kernels on the compute nodes of our current and future systems results in higher efficiency so that fewer compute nodes are needed to achieve the same application performance as on competing products. Fewer compute nodes translates into lower overall power consumption.

IBM: We believe that Active Power Management (APM) is another approach to minimizing power. With this approach, the software provides hints to the firmware/system management to turn off the power to units not in use. Since the power needs to be turned off completely and then turned on again, the granularity of the usage model needs to be fairly coarse to be effective.

Regarding the architectures of the future...

What breakthroughs do you foresee in the amount of memory that can be added to machines?

IBM: We foresee that memory packaging density improvements, possibly chip stacking, will be required to improve the amount of memory that can be added to machines. It will also reduce the latency to access these larger arrays.

Cray: Although some architectures are highly dependent on memory per node, Cray systems are, again, designed from the start with scaling in mind. Our petascale system will have globally addressable memory so our machines can have up to petabytes of physical memory that would then be available to all the processors in a system.



Figure 3. The Cray XT3 at ORNL.

Do you foresee computing platforms composed of heterogeneous architectures?

Cray: Absolutely! That is the heart of our Adaptive Supercomputing vision, our long-term development effort in which hybrid supercomputers with a variety of processor types will adapt the processors to each application, or portion of an application, with minimal user intervention.

Different applications achieve optimal performance on different types of processors, but most high-performance computers offer only one type. In order to take full advantage of the system, applications today need to be tuned for the computing platform, often requiring significant effort on the part of the user. With Adaptive Supercomputing, the concept of heterogeneous computing—access to multiple process technologies—takes a great leap forward by providing an integrated view of hardware and software to support multiple processing technologies in a single system.

In accordance with this vision, all new generations of Cray supercomputers will combine standard microprocessors (scalar processing), vector processing, multithreading, and hardware accelerators in one high-performance computing platform using a Linux operating system, starting with the Cray XT4 and Cray XMT supercomputers we announced in November 2006. We are also working on powerful compilers and other software which will automatically match an application to the processor technology that is best suited for it—adapting the system to the application rather than requiring the user to adapt the application to the system. This complexity is hidden from the user, allowing scientific and engineering problems to be solved more quickly—and programmers and end users to be more productive.

IBM: Yes, we believe that the broad application scope for petascale systems, coupled with the environmental constraints in the areas of power and cooling require multiple system architectures in some cases. Examples of innovative architectures to optimize large-scale systems for important application classes include Blue Gene and Cell. Blue Gene has delivered break-away performance per watt for a number of scientific applications through a combination of power-efficient design and innovative interconnect technology. We believe that Cell holds the promise of delivering similar breakthroughs for large-scale problems that can exploit an accelerator-type model for application deployment. IBM is delivering advanced software capabilities to allow these focused architectures to work in concert with general-purpose clusters. This will provide petascale compute capabilities that can be used by a growing number of application developers.

How will programming paradigms need to shift?

IBM: This is a complex question. Parallel programming will need to adapt to advancing technologies, such as SMT multi-cores, RDMA transports, advanced CPU architecture (with asynchronous move engines, fast synchronization constructs, and remote atomics), hardware acceleration of collectives, maturation of Partitioned Global Address Space (PGAS) programming models, as well as integration of message passing, OpenMP, and PGAS models in a single application.

Cray: Our goal with Adaptive Supercomputing is to create a comprehensive hybrid-processing environment where the system can adapt to the application transparently. Users will not need to change their application to fit a particular

hardware environment. They can write their code in a natural and intuitive way. Notably, even the computational accelerators Cray is developing as part of our DARPA HPCS program do not require modification of the application.

When will we have a petascale computer on our desktop?

Cray: I'm not the best person to make a prediction like this one! If you look at it, the entire HPC industry has been increasing system-level performance faster than Moore's law—but it will have taken about a decade to go from a teraflop to a petaflop. Thus, I expect that sometime within the next few years we'll have a teraflop on a chip—so, maybe 15 or so years to see it on a desktop. But I don't recommend betting your lunch money on my prediction!

IBM: The more suitable question might be: How will the individual user gain easy and seamless access to petascale computing? The engineering problems associated with creating a petaflop on the desktop will be daunting in any reasonable time frame. Perhaps a better way to think of the future is in terms of a more virtualized environment in which the engineering and environmental issues of enabling supercomputing for the desktop are finessed within the Internet.

Petascale systems present many challenges for integration and testing. What type of relationship with government customers will be most useful to enable testing and evaluation of petascale systems?

IBM: We have always found that true partnership is essential for successful deployment of large, complex systems. As we cross the petascale mark, this will become even more evident. The most effective relationship between IBM and our government customers has been one of shared risk, open communication, and alignment with regard to success criteria. Of course, there is still much to learn as systems grow to multi-petaflops. Greater flexibility to react to problems, and leverage opportunities for improvement will be required throughout the procurement and deployment cycle. A thoughtful balance between competitive pressure and value recognition will be required to ensure a sustainable technology stream that makes sense from business and technical perspectives. These relationship characteristics will enable government customers to continue to push the envelope in applying massive computational capability to problems that are critical to our national interests, while at the same time encouraging vendors to maintain strong investments in future supercomputing innovations. We have many examples of such partnerships which have proven to be very productive for both IBM and our customers.

Cray: Systems are so large now that you can't really deploy them in-house for testing and evaluation—this makes close partnerships with customers critically important. Our ongoing partnership with Oak Ridge National Laboratory is an excellent example of this, and we fully anticipate that this



Figure 4. An artist's rendition of the Cray XT4 supercomputing system.

will be the scenario for petascale systems in general. Having such a close relationship with customers where we can gain access to their systems part of the time is beneficial to all of us.

Regarding High-Performance Computing in general...

Is the use of HPC becoming more widespread? That is, do you think that the community is large enough to sustain High-Performance Computing?

Cray: The overall high-performance computing market has grown over the past several years, largely due to the introduction of highly leveraged systems using commodity servers. While this move has expanded the overall market for HPC, it has stifled innovation at the high-end. In fact, programs like DARPA's HPCS, SNL's Red Storm—part of the National Nuclear Security Administration (NNSA) Advanced Simulation and Computing (ASC) program—as well as the DOE Office of Science's Leadership Computing Initiative are indicators that the HPC community needs supercomputing systems specifically designed for the demands of complex HPC applications. Advances at the high-end trickle down and adoption typically spreads into a number of industries, such as automotive and aerospace, energy, climate/weather, life sciences—and it is also spreading into commercial data centers. This often occurs through collaborative programs between government and industry, such as the DOE Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program, as well as the NNSA and National Science Foundation (NSF) industry partnership programs. Cray is the only company solely focused on supercomputing with purpose-built systems that will provide high-end users with the systems that support the most complex and demanding HPC applications.

IBM: We foresee the HPC community extending far beyond universities and government laboratories. For quite some time, the techniques and technologies of HPC have been

used commonly in areas outside the academic community and government, in areas such as automobile and aerospace design, financial portfolio modeling and risk assessment, new drug discovery and logistics, and supply chain modeling. In recent years, this community has expanded with the emergence of applications in areas including bioinformatics and digital media. Overall market growth has been in the double digits for the last several years and has been outgrowing conventional areas of information technology (IT) quite comfortably since 2000. We are convinced that there is sufficient business and opportunity for innovation and that is what motivates IBM's strong participation. Others might have a narrower view of the opportunity, but we believe the myriad ways HPC is being applied makes it an appealing market segment for us. However, extreme scale supercomputing continues to be a reasonable size niche market.

What needs to be done to increase the HPC user community?

IBM: The same things that are required in any IT segment. The technology must be engineered for easier usage and made more affordable progressively. Also, newly designed products and services must demonstrate a clear strategic advantage for users. IBM already does this by offering a diverse set of architectures designed to map the right tool to the problem at hand. The biggest problems that we have seen in the HPC market is when companies try to force-fit a solution or technology where it simply does not belong.

Another important factor is the development of the ecosystem around HPC. This includes education in universities, seeding universities with supercomputing capability, and designing academic coursework that encourages research and exploitation of supercomputers by students. In addition there needs to be a renewed focus in enabling new and growing sets of emerging applications. Unfortunately the HPC community is fragmented and obsession with top 500 lists have distracted the community from more strategic and important issues at hand. The DARPA HPCS program has helped refocus the vendors and the community around the theme of productivity and new application enablement.

Cray: As envisioned by DARPA's HPCS program, the biggest barrier to increased use is ease-of-use, or "user productivity." We must make high-performance computing systems easier for people to use. This is the essence of the DARPA HPCS program and Cray's developments in that program will be utilized on our petascale systems for the SciDAC community.

What countries are emerging as your competitors in the HPC arena?

Cray: Since we sell supercomputers around the world, I don't think of countries as our competitors, but companies. As far as international companies focused on supercomputing, most of our competition comes out of Japan, with companies such as Fujitsu, Hitachi, and NEC.

That said, there are several countries that clearly have demonstrated the technological wherewithal to compete in high-performance computing on a global scale outside of the United States. Of course, Japan remains the most formidable competitor, but there are major efforts afoot in countries such as China and India, and within the European Union.

IBM: We tend to not view countries as competitors, given the international nature of our business. Instead, we try to understand the needs of clients in each country and work to ensure that our products and services are aligned with those needs. The nature of HPC today is so centered on international collaboration that the imposition of national boundaries has very little meaning in terms of how we approach opportunities.

Has the National Nuclear Security Administration (NNSA) Advanced Simulation and Computing (ASC) program been a driver for HPC in the U.S.?

Cray: Most definitely. The ASC program was one of the biggest mission-oriented programs to invest in a major way in using HPC—not only in purchasing some of the world's largest supercomputers, but also in building scalable applications to take advantage of them. At Cray, our collaboration with Sandia National Laboratory to design

High-performance computing systems must be made easier for people to use.

and build "Red Storm" is a prime example of how the ASC program is a driver for the industry. From this collaboration came our commercial Cray XT3 and current XT4 system architectures. I believe that the entire HPC community has benefited from the ASC program, worldwide.

IBM: Our participation in the ASC program goes back to the very beginning of the program when it was defined in 1996. It clearly has been a catalyst for not only developing new technologies, but producing and introducing innovative technologies at a faster pace than the ordinary speed that commercially-driven innovation would imply. We have always seen the requirements of the ASC program as a leading indicator of more general market requirements, as long as three years in advance. This has increased our competitiveness significantly across a broad spectrum of market opportunities that mature and manifest themselves at a slower pace. In essence, we are ready to compete for the next significant commercial opportunity because we already have the technology in place when the opportunity appears, mainly because of the ASC program.

Thank you both for contributing to this issue of SciDAC Review. ●