

High-End Computing and NATIONAL SECURITY



George Cotter recently answered our questions about hardware, software, science, and national security.

SciDAC Review: *First of all, congratulations on your election to the National Academy for your leadership in the research and development of high-end computing (HEC) and communications for national security. What does this recognition mean to you?*

George Cotter: I am convinced it is a vote of “excellence” for the work we do. I’ve been fortunate to be in a leadership role here, but to me it means great recognition for the work that a relatively small agency has done, in a field of great importance to national security, industry, and science.

You have had a hand in shaping the evolution of supercomputing. What are the top one or two impacts in shaping HEC in which you feel you played a personal role over the years? Is there one that you wish you had to do over again?

I think to stay focused on the extreme high-end problem was key to shaping that evolution. Learning from Federal Coordinating Council for Science, Engineering, and Technology (FCCSET) studies and setting up our Supercomputer Research Center in the 1980s and 1990s, and shaping federal efforts with workshops and studies which tried to define the architectural and technology challenges that had to be addressed. Petaflops I and II showed clearly that simply scaling up commodity systems would not work. We put serious follow-on effort into the Hybrid Technology Multi-Threading architectural study. The President’s Information Technology Advisory Committee report (1998) made the case for a federal research and development (R&D) program that gave solid, consistent support. Unfortunately, few federal agencies understood the message.

With the help of Congress, we took up the challenge of keeping the Cray product evolution alive, aiming for balanced, heterogeneous architectures. We convinced Congress that

industry on its own would not pursue advanced technologies—our High-Performance Computing and Simulation (HPCS) partnership with Defense Advanced Research Projects Agency (DARPA) ultimately demonstrated this. Fortunately Congress listened, and chartered the Department of Defense to do the Intelligence High End Computing study, which was finally implemented in Fiscal Year 2008.

But I failed to see that the high-performance computing (HPC) market would broaden far more than it deepened. Perhaps the greatest disappointment was the failure of the Federal Networking and Information Technology Research and Development (NITRD) program to set the bar high enough to drive architectural and technology directions. I was a major mover to create the HPC Users Forum, in an attempt to collectivize the Feds. It failed to do so.

Over the years, you have worked with a great many pioneers in the field of supercomputing. Looking back, what are the most important partnerships which you have had?

We have enjoyed first-class relationships with all the great HEC architects—Seymour Cray, Greg Papadopolous, Steve Wallach, Burton Smith, Thomas Sterling, and now Steve Scott. I’m sorry I can’t include Tad Watanabe in this list. It has been far less a personal issue and much more their recognition that the NSA must remain a leader in the field and that our users and developers are worth listening to. I also believe that establishing direct and frequent interactions with Federal Centers of Excellence/National Labs has been very important, as NSA’s relative clout began to wane.

Recently, there is recognition of HPC as the third leg of scientific discovery, together with theory and experiment. What do you think about supercomputers as major scientific facilities?

I think it is a shame that this recognition has been so long in coming, and more importantly, is so slow in producing change. A critical number of dramatic breakthroughs in scientific discovery are probably needed. Science aside, what’s more depressing is that HEC modeling and simulation is not yet practiced in most major defense weapons acquisition programs. In our work, HEC modeling and simulation is absolutely key to success; unfortunately, it is impossible to show-and-tell. Here is a great opportunity for federal agencies; one would think that a number of these agencies could focus on demonstrating, indeed dramatizing, scientific discoveries and building much stronger national support for HEC. The

current NSF study of the dependency in four science fields on HEC took too long to organize and was too carefully insulated from the HEC practitioners. I suspect we'll see glowing words mostly about the needs of experimentalists. I hope I'm wrong.

The DOE SciDAC program has tried to facilitate strong coupling between hardware, software, algorithm, and science. What do you believe is the major accomplishment of the SciDAC program?

Aiming high, I believe, in the interest of great science, and having the courage to take the lead for government. I hope DOE Office of Science will continue to gain the confidence of the Administration and of the Congress. It benefits the National Security community in many ways, as well.

What are the most important advances that are needed in HPC to solve problems of national importance?

We must find innovative ways to design multi-core processors. Simply replicating identical cores on a device will be just another form of massively parallel processing; it transfers the task of increasing performance to the user. This is the reason server vendors have been slow to adapt large-scale multi-core processors. There have to be ways to take advantage of tightly integrated hardware parallelism, by incorporating multiple processor architectures on a single device. Of course this challenges compiler designers. Ultimately the degree of processor integration will come down to this.

Large-scale HEC systems still tend to be unbalanced. The major weakness is lack of all-optical interconnections from I/O to the device level. We once had a project called Lightning, in which we demonstrated that we could color switch 256 optical streams; it could have been a major advance over I/O systems still in use today. Something such as this will be needed as blue ray optical disk farms enter the HPC realm. Incidentally, I like the idea of inserting processors in disk arrays, to give us far more data handling agility than we enjoy today.

To focus on the phrase, "to solve problems of national importance," the HPC community needs much greater credibility than it has today. The HEC community knows there are physical objects that can be modeled, examined, and manipulated far more readily in HEC systems than in experimentalists' laboratories. Our top national science officials really should address this issue with a "top 10 list"—major challenges and awards to go to the team that solves one such challenge through computational science. Have the National Research Council generate the list.

In many ways, we are at the cusp of a transition from single-core processors to multi- and many-core processors, and perhaps heterogeneous processors with increasing use of accelerators. If this is the immediate future, what should we do to get ready for it?

I would argue the case for a deeper architectural perspective on the integration question. Recall how the "attack of the killer micros" led to simplistic scaling views—"MPI uber alles." Generally, extreme high-end users have not studied their range of applications, nor understood the relative contributions of processor architectures, the potential for accelerators, even special purpose processors and other components. These are easier issues for us, since we design and build one-offs.

For extreme high-end systems, the cusp is transition to heterogeneous but balanced architectures; and balance in this case will require significant change in operating systems, compilers, and perhaps languages. Scalar, vector, and multi-threaded processors in single integrated systems are here to stay, with pipelined superconducting processors not far behind. The open question is how tightly we can achieve architectural integration—and this will depend on truly innovative system software development. Japan seems clearly headed this way, in their one-off Kei Soku 10 petaflops program. Frankly, I foresee that the marketplace will support only limited numbers of extreme high-end systems, and suspect that customization and perhaps scaling-down will be the correct strategy for a couple of companies.

As computers have become more powerful, it has stretched the infrastructure requirements. Are you surprised that space, power, and cooling have become critical factors in deploying leading-edge systems?

Not really surprised, since we had to address this issue over 15 years ago with Supercomputer Systems Incorporated, Steve Chen's MP architecture—four megawatt machines. We built a special building for just this purpose, although we clearly did not go far enough, as Oak Ridge National Laboratory (ORNL) clearly understands. However, it goes beyond infrastructure requirements: the power, space, and cooling factors are becoming major life cycle support cost issues also. How far out-of-hand all of this gets and if it begins to affect our goals and HEC and special purpose device aspirations, that is worrisome. I'm pushing hard on technology breakthroughs—superconductivity for example—but here we are talking about fundamentals of industrial policy, and that's tough.

What are the major differences in the top issues with HEC, from your early days to today?

In the late 1980s, I firmly believed that high-end users would prevail, and as mentioned, failed to see that the broadening of the market base would, in fact, divert industry from meeting the needs of a relatively shrinking top tier. The Japanese Earth Sciences machine demonstrated that ultimately the market was bifurcating and that our high-end system and technology needs could only be sustained with government investment.

I was late in recognizing that the trickle-down technology theory had given way to a trickle-up mass market dynamic. So really good R&D just can't make it into our systems easily. Naively, early on, I expected software evolution would keep parity with hardware evolution, but industry has let us down in failing to bridge into academic research. Individual users can only observe; they cannot solve this. Witness our own torturous efforts on UPC.

Today, we seem to be most successful when we form small alliances and leverage through a combination of development projects, contracts, and partnerships.

What do you believe are the fundamental computer architectural impacts—positive and negative—that you have seen over your career?

Of course, Seymour Cray's introduction of vector processing architectures was wonderful for the time and suited the peculiar needs of a small but thirsty HPC community.

Seymour understood the issues of balance—processor, memory, communications. Bear in mind that users had both scalar and vector processors at hand, with global addressing and lots of memory, and that started a very important trend which continues today.

As we reached the limits of vector scalability—cost and functionality—the advent of the massively parallel processing (MPP) offered another path to greater performance. Fortunate indeed were the installations which had both.

Thomas “Beowulf” Sterling gave the masses something in entry architectures that spread HPC to a great number of academic institutions, spawning interest and growth. It was like a teenager’s first used car; you can jazz it up, but it is still a Beowulf.

And the Japanese Earth Sciences machine was a stunning accomplishment. It raised the bar for all truly high-end users and, although a one-off machine, it led vendors to quickly adopt many of the better features in commercially available machines. This has now been followed by the Riken Kei Soku 10 petaflops project, a massive heterogeneous architecture, with massive incorporation of Tokyo University’s Grape architecture adapted for specialized but important science applications.

On the negative side, massively parallel single instruction, multiple data architectures were the Edsel of the HPC world. There was amazing hype but cooler heads could not find the applications to fit the architecture.

I view the commodity-based cluster architecture trend with mixed emotions. It certainly allowed lots of institutions and sites to join the HPC world, but it also led many first class installations to keep from pushing limits of performance. A high position on the TOP500 list and a banner on a web page and it became very hard for most of these installations to admit to the very poor scalability of these architectures. The big question is whether this mid-range market will chase massive multi-core cluster architectures, as the next challenge to increasing performance.

We spend a lot of time discussing hardware and computer design these days. However, system software plays a very large role in the usability of supercomputers. What is your vision for the software stack in the future, specifically operating systems and compilers, given that the coming generations of computer users are going to be facing systems with millions of cores?

System software innovation must occur. The truly great computer architects worry as much about this as they do about hardware and computer design. The great parallel advances in systems software that occurred for vector systems, for well-balanced MPP architectures, and which more recently is occurring in multi-threaded systems, really makes this clear. One of the great lessons out of Petaflops I was that systems software issues have equal importance, leading to Petaflops II. Compiler and operating system design will certainly determine the degree of hardware integration we attain in heterogeneous architectures.

I suspect truly effective use of massive multi-core architectures will be hard to achieve, except through years of hard work to tune applications to fit the hardware suite, not unlike recent experience with commodity clusters. It’s possible

that intensive study of well-known algorithms can lead to standardized library suites, but frankly, vendors shun such labor and we continue to have trouble transitioning university innovation into the market place.

As supercomputers begin to be used more widely in addressing national and global issues such as climate change, medical R&D, energy security, and terrorism—to name just a few—their usability by greater segments of the scientific community becomes necessary. Given that you foresee a change to heterogeneous systems with accompanying change in operating systems, compilers, and languages, what can we do in the near term to increase the usability of these systems?

In the near term, most of us will not be permitted the luxury of a poorly used, say, petaflops-capable machine that we are not ready for. As systems grow in size, performance, and cost, we have to work much more closely with vendors and collaborators, in parallel with the development of the system, to make effective use of the machine. We should select applications now, taking on the tasks that historically vendors have not addressed until full-scale machines are available. Scalable simulators may be a sensible investment for communities of interest. For structured programs such as DARPA’s HPCS, this could be relatively straightforward; for scientific communities of interest—such as weather, climate change, and medical tomography—some coordination function by lead institutions would be necessary.

Do you have a longer-term strategy you would recommend?

For the longer term, I actually anticipate sets of applications to have considerable effect on system design. Certainly, heterogeneous architectures will be customized for application suites, which lessens the risk that users will take inordinate time to learn how to get performance out of the system. Given the scale of such machines, finding the sweet spot in applications performance cannot be a hit or miss proposition. It’s entirely likely that no two vendors’ very large scale machines will be alike.

Many times in the past, solution algorithms have been responsible for as much as 50% of the performance gain in large problems. Do you think that is true today? And what is your vision for computer speed advances due to algorithm improvements, given that the coming generations of computer architectures will have millions of cores?

What is remarkable is that this works in both directions. Algorithm improvements almost always lead to design innovation, even to special purpose devices, where needed. And this is precisely the direction that multi-core computer architectures should take. The real value of this architecture trend will be seen in the variety of innovative processing architectures and communications/streaming designs that can be implemented on a single device. This of course is not what the processor vendors envisage, but they will feel the competitive heat from the huge vendor server market, so we’ll see innovation.

Congratulations too, on your new position in the Office of the Director of National Intelligence (ODNI). Can you give us some idea of this new role? Is it a new position and a signaling that the Director

of National Intelligence (DNI) sees a greater role for supercomputing among all the intelligence agencies?

I have been influential over the past year in raising interest in HEC at the DNI level, and thought I could usefully spend a year or so in the ODNI as a Senior Advisor to the Deputy Director National Intelligence/Acquisition. This office has primary responsibility for overseeing all intelligence community acquisition activities, intelligence community collection architectures, and Science and Technology, including the newly formed Intelligence Advanced Research Projects Agency.

Yes, the DNI recognizes the criticality of supercomputing to my agency and several others and I'm expected to help find that greater role, and to foster greater collaboration across the intelligence community in HPC and other technologies.

What do you see as your greatest challenge in your new position, as it relates to supercomputing and computer technology?

The greatest challenge will be to bring experts and non-experts into collaboration and to get meaningful and lasting dialogue started.

Earlier you commented that it was important for the U.S. to maintain a strong HEC industry, but you also commented that there will likely be only a few extreme HEC systems in the U.S. at any given time. Won't industry likely want a bigger market, before they commit the type of top talent it takes to design and develop these extreme systems? What do you see as the strategic approach government should take going forward, to maintain a strong domestic HEC industry?

What industry will want is a fair return on investment. The only way that is possible in an American extreme computing market almost totally dominated by federal users is if the government underwrites a fair share of design and development cost. Note that this includes advanced research and development of technologies that otherwise would not be undertaken by industry, or vendors. Sharing and maintaining competition are not necessarily in harmony, so it will be a challenge if the government-market relationship moves in this direction.

How has the government's view of HEC changed over the years?

From FCCSET in the 1980s to NITRD in this millennium, it has morphed from a set of focused, somewhat altruistic federal goals to a budget-driven and somewhat politicized cross-agency bureaucracy. The good news is that partnerships form, relationships and associations form and are enduring, and some very good technical work gets done.

What should the U.S. do to maintain pre-eminence in HPC?

The U.S. will have to maintain domestic critical mass in the HPC industry if we are to remain pre-eminent. I do not believe that we can totally outsource HPC component manufacturing or full system integration, and somehow remain pre-eminent in HPC. There is entirely too much interplay between research, development, design, and manufacturing to expect we can hold onto the intellectual capital, strategic vision, and implementation factors. If we don't maintain critical mass domestically, we will certainly lose the industry to foreign competition.

What fundamental infrastructure changes are required to protect the nation's cyber infrastructure?

By fundamental infrastructure changes, I assume you mean federal-state-local relationships including critical private sector elements. Firstly, there must be broad national acceptance that there is substantial risk and that risk will only increase over time. Secondly, there must be recognition that protection can never be certain in the open society we know and love. Thirdly, sea changes in all relationships must occur; this will be evolutionary, but we can only be successful by building on mutual trust and cooperation. Finally, most of us are convinced that major technology efforts must be undertaken, to underpin important policy and cultural developments; that is, to make sure policy does not get too far in front of reality. Over the next few years, Congress and a new administration will be consumed with the problem. I credit the current DNI with having the foresight and courage to bring senior policymakers together on this major national challenge.

What are the key challenges for the next 20 years?

- Maintaining a strong domestic HEC industry through government funding, as necessary
- Establishing a viable national technology development strategy to support extreme HEC
- Clarifying critical dependencies on HEC in the science, technology, and national security communities
- Building far more effective partnerships between academia, government, industry, and the HEC vendor development and user communities
- Recruiting, training, and developing our successors
- Implementing some cohesive national programs that creatively cement the points above

What advice do you have for young scientists just entering the R&D field in HEC?

For many years I managed a senior technical development program and my advice to all who completed the program was to continue to invest 25% of their time in learning. For new entrants to the R&D field in HEC, that time investment should be 50%, even if you have to do it on your own. HEC R&D is as close to leading-edge as anyone will find, in any field. Materials science, advanced technology development, and computational science linked to HEC are wonderful opportunities. It is a small and selective field, so there are many ways to relate directly to advances in basic science; the excitement should never end.

With such a broad and important role in government, how do you find time to balance your personal life? I hear that you are a great cook, a marathon runner, a cyclist, an avid boater...have I left out any other major hobbies?

Yes, skiing. Just back from a great week at Steamboat Springs—56 inches of new snow over seven days. Exercise, stay in shape, stay focused, deeply engaged with friends and the younger generation—'tis easy.

Thank you for taking the time to answer our questions.