

David W. Turek
Vice President, Technical Computing OpenPOWER
IBM Systems Group
House Committee on Science, Space and Technology
Subcommittee on Energy
Supercomputing and American Technology Leadership
January 28, 2015

Good morning, Chairman Weber, Ranking Member, and members of the subcommittee. Thank you for the opportunity to speak with you about The Office of Science ASCR program, supercomputing and American Technology Leadership.

Supercomputing and National Competitiveness

Supercomputers have become critically important tools for driving innovation and discovery across an array of industries and scientific domains. By virtue of digital simulation and modeling, supercomputers have helped displace expensive, time consuming and sometimes dangerous physical models and experiments. This, in turn, has led to the development of dramatic insights into a host of very complex domains which have yielded profound commercial and scientific results. For example, supercomputers have been used to accelerate the pace of oil discovery, design new materials, predict the evolution of disease, discover fraud in consumer transactions, design fuel efficient cars and planes, provide insight into the operation of the brain, model the national nuclear stockpile, and contribute to a host of basic science discoveries. In fact, in a report from the US Council for Competitiveness, it has been noted that “modeling, simulation and massive data analysis are the ... game changing drivers for innovation.” The Council put an exclamation point on this by succinctly declaring that “to out compete, you must out compute”!

Of course, the problems of tomorrow are not the same as the problems of today, and there is a growing chorus of concern that without adequate investment in supercomputing, the ability of the US to continue to out-compete its sovereign competitors will diminish. It is our contention that investment in the Advanced Computing Science Research Programs in the DOE's Office of Science are instrumental to ameliorating this concern: maintaining innovative leadership in supercomputers goes hand in hand with maintaining scientific and economic leadership.

International Supercomputing Activities

The relationship between advanced computing and economic competitiveness is well understood internationally. In Europe, under the banner of Horizon 2020 and PRACE (Partnership for Advanced Computing Europe), there are a myriad of advanced computing efforts covering basic research (e.g. The Human Brain Project), support for university access, and the development of new computing approaches for complex scientific and engineering problems. In China, there has been a well-funded set of government programs focused on things like microprocessor design and software which are expected to contribute to the performance of domestic industries as a prelude to the establishment of a supercomputing export industry. Japan, this past fall, launched a government funded program for advanced supercomputing to achieve worldwide leadership (in a rudimentary way) within the next six years. There are even nascent plans for sovereign funded efforts to drive advanced supercomputing being contemplated in Russia, India, and Korea. Nothing will stop or otherwise impede these international efforts so competition with the US is here to stay. However, to maintain economic advantage it is very important for the US to be the leader in advanced supercomputing: there is a big difference between being present and being the best. It has long been the case that software engineers gravitate to do their work on the best and fastest supercomputers. In the absence of true supercomputing leadership the software engineers we depend upon to build the applications solving

some of most vexing business and scientific problems the US faces will migrate to wherever leadership exists. What we lose in terms of early access to these skills, someone else in the world will recognize as a gain at our expense.

State of Supercomputing in the United States

The problems for which supercomputers are designed are of such extreme complexity that conventional computer technologies are often stretched to the breaking point. However, it is when computer designers are working at the extremes of technology that innovation and invention is oftentimes spurred. For example, in the early nineteen nineties when commercial computing was struggling with data files measured in billions of bytes, supercomputers were being designed to work on files 10,000 times larger. Many of the technologies pioneered in supercomputers for big data have, over the course of time, been adopted in conventional computers helping usher in a new era of advanced analytics widely used in companies of all sizes to gain strategic business insight

Another critical issue facing supercomputer users is the amount of energy required for operation. In economic terms, the rule of thumb is that the fastest supercomputer in the world at any point in the last 50 years has cost between \$100 million and \$200 million dollars. We forecast that absent serious and sustained innovation in energy efficiency the cost for electricity for the fastest computer at the end of this decade or early the next could approach \$100 million per year! Clearly, the growing operating costs for supercomputers, left unchecked, could put a material damper on adoption and propagation throughout the economy and have a serious downside effect on the nature and rate of economic and scientific innovation. Conversely, if major advances in energy efficiency were realized one could envision a future where a single researcher could expect to have a desk-side computer equivalent in capability to the fastest computers in the world today. The ubiquity of this caliber tool in the hands of

large numbers of people would be a tremendous catalyst to accelerating the pace of innovation in the economy.

There is also the need for dramatic innovation in network technology, computer memory designs and storage technology. Today, one can move a file containing a billion bytes of data across the country in a few minutes under ideal conditions. Supercomputer applications often deal with a million billion bytes of data! This requires transcendent technologies and architectural designs, especially as supercomputer problems demand ever greater amounts of data. We need innovations to make memory dramatically more affordable to contain these volumes of data; it is commonplace today that much more of the acquisition cost of a supercomputer is tied to the cost and speed of memory than it is to the microprocessor. And we need innovations in storage technologies that affordably house these volumes of data but also are capable of making the data readily available to the supercomputer. Rotational speeds of spinning disks are relatively stalled with respect to the growth rate of data and promising new technologies are marginally affordable at best given the scale of data being entertained. Solving these types of problems holds huge promise even for a single computer user: with supercomputer inspired innovations one could imagine a handheld device holding the contents of the Library of Congress , continuously updated, with near instantaneous intelligent cognitive capability in any language, allowing the user to get answers to almost any question imaginable in real time.

Software is perhaps one of the greatest challenges facing supercomputer use today. The fastest supercomputers currently employ millions of computer cores orchestrated and managed by software to work in concert on solving very complex problems. However, much of the software in use is based on mathematical and algorithmic approaches that are more than 40 years old, and created at a time when there was no concept of current supercomputer design. It is imperative that software innovation be driven in lockstep with the innovations coming in computer hardware.

Addressing the challenges I've outlined will involve a host of players including supercomputer companies, technologies companies, software companies, universities, and the national laboratories. Much of the work is necessarily collaborative: a hardware company must work closely with software companies to produce the best solutions for example. All of the work is necessary to drive the pace of innovation to the requisite competitive level.

Supercomputing Collaborations with the Department of Energy

For more than twenty years, IBM has been engaged in a partnership with the Department of Energy to produce advanced supercomputing technologies in support of critical problems facing the National Nuclear Security Agency and the Office of Science. In the 1990s our collaboration with the NNSA resulted in advanced networking technologies enabling successful deployment of systems in support of the Accelerated Strategic Computing Initiative; later in the decade we extended the partnership through co-design of the Blue Gene system utilizing innovations enabling massive scale and dramatic improvements in system energy efficiency; and in the last decade we collaborated on the development of the Roadrunner system demonstrating the effectiveness of accelerators in the solution of highly complex science problems. In each case the innovations and ideas derived from our collaborations were adopted, copied, or extended by other computer companies helping to usher in an era of dramatic growth in the use of supercomputing. In a recent National Academy study, it has been noted that "virtually every sector of society-manufacturing, financial services, education, government, the military, entertainment, and so on-has become dependent on continued growth in computing performance to drive industrial productivity, increase efficiency, and enable innovation". In essence, an advanced economy needs advances in supercomputing to remain competitive.

Pathway to Leadership for the US

The longstanding collaboration with the Department of Energy is entering a new phase with the launch of the CORAL program targeted to produce systems for Lawrence Livermore and Oakridge National Labs beginning in 2017. These systems will be more than 100 times more powerful than the Roadrunner system and many times faster than the current fastest system in the world in China. These systems are also a prelude to the design and development of exascale systems in the early part of the next decade. (An exascale system will perform about one quintillion (1 followed by 18 zeros) operations per second.) The path we are on with CORAL will stimulate a new wave of innovation in supercomputing giving rise to new technologies and designs that we would expect to be broadly beneficial across the entire computing industry. We have to invent new designs for managing the massive amounts of data that supercomputers are asked to analyze, amounts of data so large that simply ingesting them into the supercomputer can take weeks. We have to invent new software techniques to program systems with millions of microprocessors working in concert. We have to find ways to make these systems reliable when the aggregate mean time to failure of all the parts implies continuous failure of the system. And we have to find a way to make these systems energy efficient so that they can become affordable.

We endorse funding for the Advanced Computing Science Research Program of the DOE because our historical experience working collaboratively with the DOE has proven to be of tremendous value in terms of stimulating the kinds of innovation we expect the CORAL and follow-on programs to require. Through the kind of public-private partnership that has characterized much of what we have done with the DOE over the last few decades we expect to see the US continue to lead the international supercomputing community for years to come.

Thank you for this opportunity to speak with you today. I would be happy to answer any questions.

