High Performance Computing – Scientific Discovery and U.S. Competitiveness

Raymond L. Orbach
Under Secretary for Science
U.S. Department of Energy

Transforming Through Computational Simulation:
A Computation-Based Engineering Summit
National Academy of Engineering
September 17, 2008

High Performance Computing is the third pillar for scientific discovery – along with theory and experiment. And, experiment and computational science are indelibly linked to technology development and engineering. There is a clear feedback loop. Science is advanced by new technologies, and, concomitantly, drives technology development, and the Department of Energy is at the forefront of this nexus. Through the combination of high performance computing facilities, applications expertise, applied mathematics and computer science research, the Department is delivering computational science breakthroughs today, and leading the way for tomorrow's scientific discoveries in a wide range of areas. These include climate research, nanotechnology, energy and the environment.

These achievements have developed through President Bush's first and second terms, and are a direct consequence of his commitment to basic research. The first step was to improve performance of the U.S. computational and networking infrastructure, as well as existing codes. At the beginning of President Bush's Administration, the Department of Energy had largely switched from vector systems to massively parallel systems. And our application codes were getting, at most, 10-20 percent of theoretical peak – with many applications getting in the single digits. The largest computer available to the Office of Science was the NERSC facility, which at the time seemed quite capable at 3 Teraflops *peak* speed. Our highest *sustained* speed on, for example, fusion codes was 485 Gigaflops. Data was traveling into NERSC, and around the DOE system via ESnet at 622 megabits per second. Gains were being made with clock speed improvements in individual processors, but future gains would provide only marginal improvements. It was clear at that time that we needed a radical new approach to scientific code development to get more out of our investments in computing. What was not clear was the profound impact our new approach would have on science.

The Department of Energy undertook an innovative program for accelerating progress in computational science by breaking down the barriers between disciplines, and forming dynamic partnerships between applications – such as astrophysics or biology – and the computer scientists and mathematicians who deeply understand the hardware and software available. This program, called Scientific Discovery through Advanced Computing or SciDAC, has a remarkable history of success in advancing code development, improving the performance of scientific applications by up to two orders of magnitude. Today, NERSC is in the process of upgrading to a 355 Teraflop quad-core Cray XT4 system, and ESnet delivers 10 gigabit per second core service with Metropolitan Area Networks, such as the one in the bay area that moves data into and out of NERSC, at 20 gigabits per second. In addition, the Department now also offers two Leadership Computing Facilities – at Oak Ridge and Argonne National Laboratory. Together these facilities offer architectural diversity with a 263TF Cray XT4 at Oak Ridge and the world's fastest computer dedicated to open science, the 556 IBM Blue Gene P at Argonne. And SciDAC Fusion codes are getting up to 75% of theoretical peak on our 263 Teraflop Cray system at Oak Ridge.

The impact on science was even more dramatic. Our climate models can now spontaneously generate local phenomena such as hurricanes; we are learning how stars undergo supernovae collapse; improved modeling and simulation of sub surface flow allow us to better understand nuclear waste containment as well as facilitate carbon sequestration efforts; multifaceted collaborations between applied mathematics, computer science and fusion scientists are providing key elements in the research towards the success of ITER and the taming of nuclear fusion for our energy needs – and the list goes on…

A SciDAC application that is most relevant to our topic today is combustion. Transportation accounts for 60 percent of petroleum use by the United States— an amount equivalent to all of the oil imported into our country. One means to combat the cost of petroleum is to develop more efficient vehicles. If low-temperature compression ignition concepts employing lean, dilute fuel mixtures were widely adopted in next-generation autos, fuel efficiency could increase by as much as 25 to 50 percent.

Recently, a team led by mechanical engineer Jacqueline Chen at the Office of Science supported Combustion Research Center at Sandia National Laboratories used some of the Department's supercomputers to create the first three-dimensional simulation to fully resolve flame and ignition features including chemical composition, temperature profile, and flow characteristics. The simulations reveal details of these features on all size scales of a turbulent hydrogen fuel jet in a hot co-flowing air stream as it ignites. Simulating the complex hydrocarbon fuel, ethylene, required 4.5 million hours running on 30,000 processors on the Cray XT4 at Oak Ridge National Laboratory and generated more than 50 terabytes of data, which is more than five times as much data as contained in the printed contents of the U.S. Library of Congress. Indeed, the data itself has become a library that engineers are using to develop predictive models to optimize designs for diesel engines and industrial boilers with reduced emissions and increased efficiency.

Because of SciDAC's remarkable success, we have created the SciDAC Outreach Center to help us move the advances of this program out of the Department and broaden its impact. In addition to delivering SciDAC technologies to new communities, the Center provides workshops and training, important as we move more deeply into multicore computing. The Outreach center recently formed a partnership with the Council on Competitiveness to make U.S. industry more aware of this unique resource. It is our vision to share the advances that SciDAC has brought to the development of engineering and technology tools.

Recognizing the potential of SciDAC, in 2004, the Department also increased its investments in high performance computing systems by adding what we call Leadership Computing Facilities. These facilities, combined with our supercomputers at NERSC, give our researchers access to hundreds, soon thousands, of teraflops of computing power. As I speak, the Department's first Petaflop system dedicated to open science – a Cray XT5, is being installed at Oak Ridge and will be in full operation by the end of this calendar year.

These facilities should be thought of as similar to our light sources and large scale accelerators. They benefit from open peer-reviewed competition on a global scale for use. Supercomputers should be regarded similarly. We opened access to these systems though a competitive peer reviewed process called the Innovative and Novel Computational Impact on Theory and Experiment, or INCITE, program, opening access to the Office of Science's largest computing resources to all qualified researchers regardless of institutional affiliation or sponsor. This year, we made a quarter of a billion hours of computing time available through INCITE that will expand to over a half billion hours in 2009.

The response has been wonderful. Ever since we created INCITE, in 2006, we have been able to offer an order of magnitude more hours each year. However, the requests have always exceed what we have available by about a factor of three! I am particularly proud that this demand has been universal, with applications from, and awards to, industry and other agencies continuously growing.

The industry projects tell a very interesting story. Some industries turn to INCITE to inform their in-house purchase decisions, some use INCITE to prove their design concepts, some run very specific tests or validations, while others have multi-year proposals for very basic research focused on a real application that could deliver product results. All of their applications are peer reviewed, and treated identically to university and laboratory applications as long as they intend to publish their results in the open literature. Industry can also have the option of using our leadership computing resources for proprietary research, but then they must reimburse the government for the costs of using the machine in the same manner as our light sources. Our limited INCITE experience has already demonstrated that high performance computing can contribute to U.S. competitiveness.

Boeing and Pratt & Whitney, who are participating on a panel later this morning, are both INCITE awardees. Boeing's INCITE allocation is devoted to development, correlations, and validations of large-scale computational tools for flight vehicles, demonstrating the applicability and predictive accuracy of computational fluid dynamics tools in the real-life production environment. One experiment within this project was to investigate computationally what happens to a wing when a flap is suddenly deployed. Such "flutter analysis" has historically not included simulating both the structural wing response (the wing flaps) and the airflow around the wing. This requires coupling two complex codes together in a nonlinear fashion: one for the aerodynamics and one for the structural response of the wing. Such work helps to guide control surface studies; learn aerodynamic time lags and how to better design control laws, gain better confidence in control-surface free-play modeling, and help prevent unnecessary maintenance through more aggressive designs.

One result of Boeing's INCITE award was a rare validation of unsteady Navier-Stokes tools for "simple" unsteady pitch oscillation motion, and a demonstration of the applicability and predictive accuracy of a computational fluid dynamics-based aeroelastic analysis. These tools will complement wind tunnel testing for upfront cycle time reduction during airplane design. In addition to improving airplane performance and efficiency, it also contributes to safety.

With their INCITE award, Pratt and Whitney is exploring leading-edge 3D simulations of the complex physics of a combustor to reduce solution times and improve the design process. This process will create a "virtual test" that will reduce the need for physical testing. A gas turbine engine combustor is an annulus with evenly spaced fuel injectors at the upstream end. High fidelity combustor simulations normally save time by running only a single sector (a pie slice containing one injector and surrounding components). A symmetric boundary condition is imposed to make the sector act as if it has duplicate sectors on both sides. In real life though, we know that not all sectors are geometrically alike, and the differences could result in variation. Jet engines are already highly efficient and clean, so the need for a "greener" product means that these variations be understood in the design process. Further, thrust in military programs may be enhanced by effective manipulation of the combustor exit profile. Others such as Stanford have run multi-sector combustor CFD. Because of past INCITE successes, the P&W INCITE team received additional internal funding to evaluate whether multiple combustor sectors interact to create asymmetric flows.

Pratt and Whitney have applied the technologies developed, in part, under their INCITE award to their next generation, low-emission PurePower™ engine. The P&W INCITE award also provides an example of the power of SciDAC-like partnerships. The unique IBM Blue Gene architecture at the Argonne Leadership Computing Facility identified a bottleneck in an input/output algorithm used in Pratt & Whitney's ALLSTAR computational fluid dynamics solver. Analysis by P&W and Argonne computer scientists led to an alternate approach that not only alleviated the issue, but also reduced the run time of P&W's combustor design simulations by a full day. In addition within P&W, the scaling and algorithm improvements from INCITE were combined with a new compiler and hardware resulting in a unprecedented 3X speedup of the code.

Procter and Gamble, of Tide and Pampers fame, has long been a supporter of computational-based engineering. They are using their INCITE award to study gain insight into aqueous foams. Large scale atomistic molecular dynamics simulations of cavitational and plateau regions of foams, and resultant coarse grained simulations of multiple dynamic interacting bubbles, will lead to new approaches for materials evaluations. Fully matured, these methods will be used to evaluate bio-based replacements for petroleum-based ingredients. Using the IBM Blue Gene at Argonne, Proctor and Gamble has been able to run multi-million atom models of a pure, generic surfactant, a prelude to running a billion atom model of a real, commercial surfactant-soap bubbles and foam. This will be the first time the performance properties of this class of materials will be derived directly from the molecular structures of their ingredients. The billion atom model will be incorporated into Proctor and Gamble's future product development.

INCITE is making its impact in bio-engineering as well. Igor Tsigelny of the San Diego Supercomputer Center (SDSC) at the University of California-San Diego (UCSD) and Eliezer Masliah of UCSD used their INCITE allocation on the Blue Gene supercomputer at DOE's Argonne National Laboratory to determine how alpha-synuclein, an unstructured protein, can cause Parkinson's disease. Scientists call it an "unstructured" protein because it is not stable enough to remain in one form, or conformation; and, until this INCITE project, nobody knew which conformations actually triggers Parkinson's Disease. INCITE simulations showed that alpha-synuclein can aggregate into a five membered ring-like structure that penetrates the cell membrane to form a pore. Because the concentration of calcium outside cells is much greater than inside, calcium flows through the pore and triggers a cascade of events leading to cell death. When the cells that make the neurotransmitter dopamine die, Parkinson's disease is born.

Preventing alpha-synuclein from assembling into rings sets a goal for rational design of drugs that treat the disease in an entirely different way than do current treatments that merely compensate for lost dopamine. Drugs can be designed that cause the alpha-synuclein to pair, inhibiting ring-like structures, halting pore formation, and stoping the disease from progressing once it has begun or even preventing the disease from developing in the first place.

As so here we are today. SciDAC, the SciDAC Outreach Center, and INCITE are already transforming science, with the potential to transform U.S. industry.

And this is just the beginning. Exaflop systems loom on the horizon. Based on the three workshops we conducted last year, there are huge opportunities at this speed that will transform the landscape of our world, from integration of human behavior into climate codes, to bringing convergence to the partial differential equations that govern our physical universe. It will be critical to U.S. scientific and industrial competitiveness in the years ahead that we maintain our edge and leadership in high performance computing and computational-based engineering.