The *HIGH-STAKES RACE* in **HIGH-PERFORMANCE COMPUTING**

Bragging rights to the world's fastest computers have belonged to China for the past several years, with the U.S. seeming to slip behind. This matters, because so many areas of science and industry now depend on very fast computers. These include modeling combustion processes for improved engine design, predicting the properties of novel materials, and maintaining national security. Increasingly, academic and industrial scientists also depend on high-performance computing to help analyze masses of data from X-ray probes of complex molecules or other data-intensive areas of research including artificial intelligence.

The Department of Energy is now funding development of the next generation of U.S. supercomputers—so-called exascale computers because they will be able to perform a billion billion computations per second. These computers may enable the U.S. to regain the lead from China, but it turns out that raw computing speed is not the only thing that counts. Today's supercomputers attain their speeds by employing hundreds of thousands of separate processing units that work in parallel. And for the kinds of practical applications mentioned above, it is specialized software that enables effective use of these computers—software that must divide up a computational problem into many different parts, dole out those pieces to different processors, and reassemble the answers to give a useful result. And thanks in part to early support from the Basic Energy Sciences (BES) office of DOE, the U.S. leads in developing software for such massively parallel computers.

The Summit supercomputer nearing completion at Oak Ridge National Laboratory in Tennessee will have more than 36,000 separate processor units operating in parallel and is expected to process some 200 million billion operations per second. *(Oak Ridge National Laboratory)*

Back in 1997, the fastest U.S. computer had just over 500 separate processors, each with a single calculating unit or "core." Until about 2005, computer speeds grew both by adding cores and by increasing the processing speed of individual cores. But since that time, the growth in computing capacity is now almost entirely about ever more massive numbers of processors that operate in parallel. Today's best supercomputers have 200,000 processors, each typically with two or more cores. The planned exascale computers could have many millions, perhaps billions of cores. Managing an efficient computational process with that many separate units is complex, to say the least. Managing efficient data transfers between the processors is one of the most challenging tasks for developing parallel software that can make the best use of these powerful machines.

The U.S. research effort on massively parallel supercomputing software began in the 1980s, when it became clear that parallel computers would be the future. Computer scientists realized that software designed for single processors could not be simply transferred to a parallel machine. So BES funded a multi-decade effort to develop massively parallel software from the ground up—new computer architectures, new algorithms, new ways to manage the dataflow, new ways to organize data storage—designed to benefit several different areas of science.

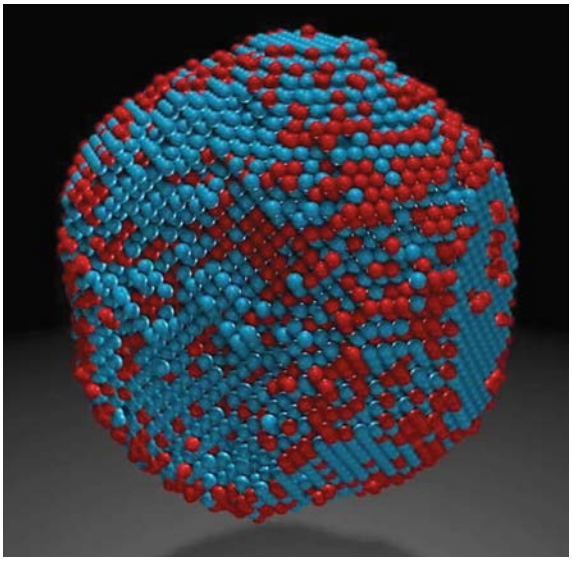

A magnetic nanoparticle composed of 23,000 iron (red) and platinum (blue) atoms is a promising material for next-generation data storage devices. A supercomputer analysis of its magnetic structure—calculated separately atom by atom—is helping to optimize its properties.
(Jianwei Miao / University of California at Los Angeles)

Theoretical Chemistry. One pioneering effort centered at the Pacific Northwest National Laboratory focused on computational chemistry and involved a collaboration of theoretical chemists, applied mathematicians, and computer scientists. It took five years, but the result was a host of innovative capabilities, such as powerful new ways to model and thus predict chemical processes. That required modeling individual molecules and how they combine chemically—in effect, how their electrons break and form bonds. The resulting computer software or code, called NWChem, has gone through successive upgrades, and has been widely used both in academia and in industry (more than 70,000 downloads of the code just since 2015). It has dramatically expanded the ability of chemists to predict the properties of a broad range of molecules and materials, which has accelerated research into new biofuels, solar cells, and advanced batteries.

Other software codes adapted from NWChem have enabled scientists and industrial chemists to model the intermediate, transitional stages of chemical reactions—which are hard to measure experimentally—and thus to model the specialized substances called catalysts that influence how those reactions proceed. The petrochemical industry depends heavily on catalysts in order to turn raw materials into complex plastics and other products. The new codes—and the growing power of parallel supercomputers—have enabled industry to develop advanced materials much more quickly and even, in at least one case, to design an entire chemical plant on a computer.

Properties of Complex Materials. A second major effort was centered at Oak Ridge National Laboratory, with a focus on predicting the properties of metal alloys and magnetic materials. Alloys often possess superior strength compared to pure metals, but their structure is complex because of the random locations of different types of atoms. Modeling alloys requires solving the fundamental equations of quantum mechanics not just for a few atoms, but for thousands and now millions of them. Magnetism is an inherently quantum phenomenon, so modeling magnetic materials also requires atom-by-atom calculations. BES supported a several-decades-long research effort to develop parallel software codes capable of such calculations. The result—the LSMS code—has transformed materials science and is proving especially useful in designing new nanoscale materials. One such effort involved a tiny

particle (23,000 atoms, about the width of a piece of human DNA) of an iron-platinum alloy for which the position of each atom was known. Researchers used the Oak Ridge code to precisely model its 3-dimensional magnetic properties. That capability enhances the potential of such materials as candidates for next-generation data recording and storage devices that could greatly increase storage capacity while lowering costs.

Combustion. A third computational challenge was to model what happens during combustion. The performance of internal combustion engines, for example, depends on how quickly and completely combustion occurs when fuel and air are injected into the cylinders—which in turn depends on both the fuel chemistry and how the air and fuel mix. A successful computer code needed to model both the turbulent motions of gases or liquids and, simultaneously, how they modulate the rate of chemical reactions. BES supported development at Sandia National Laboratories of a direct numerical simulation code of these combined processes, known as S3D. It has given researchers and industrial engine designers new insights on when to inject fuel to optimize the mixing and combustion process. Today, onboard computers in vehicles use that knowledge to adjust the injection process in real time as engines warm up or conditions change. The combination of better design and real-time adjustment has led to a remarkable decrease in pollutant emissions from diesel engines over the past several decades, while also increasing their efficiency.

The development of specialized software for parallel computers has now spread into many areas, supported by many federal agencies. Of course, raw speed remains important in high-performance computing and there have been significant, recent U.S. innovations in the types of processors used in computers—especially the graphics processors used in video games and in self-driving vehicles. The new Summit supercomputer nearing completion at DOE's Oak Ridge National Laboratory will combine both very high speed conventional processors developed by IBM and the newest generation of graphics processors developed by Nvidia.

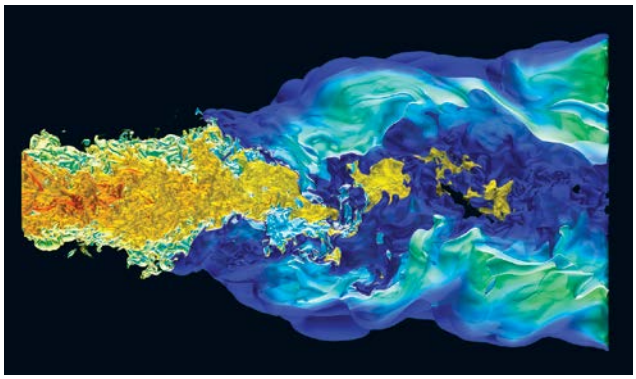

Snapshot of a supercomputer simulation to understand the processes that control how fuel burns when injected into an engine. Such simulations using specialized software developed by Sandia National Laboratories, combined with laboratory experiments, have helped engineers dramatically reduce pollutant emissions from diesel engines while also improving their energy efficiency.
(Hongfeng Yu / University of Nebraska at Lincoln)

The Summit computer also exemplifies the kind of collaborative, co-design process initiated by the BES software development projects. When the Oak Ridge Summit team decided to incorporate Nvidia's graphic processing chips in the design, they asked Nvidia to add some major new features to the chips to make them compatible with other aspects of Summit's design. Nvidia agreed to do so, and the resulting effort (including new software to run the chips) became a joint process. The result was not only powerful graphic accelerators in the Oak Ridge supercomputer, but new market opportunities for Nvidia.

The Summit supercomputer may well put the U.S. back in the lead in raw speed—an expected 200 million billion operations per second—as well as enabling still higher speeds for specialized artificial intelligence applications. Equally important, however, is the breadth of software for highly parallel computing—catalyzed initially by a sustained BES effort—that gives the U.S. a significant edge in high-performance computing.