Computing superpower

—by Ilse Genovese

The bloggosphere is abuzz with exciting news—IBM is building a monstrous $200 million, next-generation supercomputer. Commissioned by the National Science Foundation, the computing giant will be the fastest computer in the world, the first to break the petaflop barrier.[1]

The first petaflop will be about twice as powerful as today's dominant model, a basketball-court-size beast known as BlueGene/L at the Energy Department's Lawrence Livermore National Laboratory in California which can perform 596 trillion calculations per second.

To be located in Chicago, the petascale supercomputer would reportedly cost $200 million to build and require over $400 million to maintain it over its five-year lifetime. It'll be used for only a small number of "Grand Challenge" science projects, such as simulating global warming and playing "Crisis" at 60 fps.[2]

The computing muscle of the new petaflop machine will be akin to that of more than 100,000 desktop computers combined, experts say. A computation that would take a lifetime for a home PC and that can be completed in about five hours on today's supercomputers will be doable in as little as two hours.

The difficulty in building the machines is tremendous, experts admit, and the amount of power these machines require is pretty mind-boggling. But!! the scientific results that we can get out of them are also mind-boggling and worth every penny and every megawatt it takes to build and operate them.

Progress in building chips doubles the power of microprocessors about every 18 months. But this alone does not explain the leaps in supercomputing, scientists say. Around 1995, supercomputer design changed dramatically by exploiting the notion of "grid computing." In essence, grid computing is an ability to "gang" hundreds of thousands of computing nodes together, either in a single machine or over the Internet so as to find solutions to complex problems fast. This, of course, means devising better ways of having the networked computers communicate with one another and dividing up the work.

IBM's BlueGene/L built in 2004 for the Department of Energy. [Source: www.lbl.gov]

Present-day supercomputers churn out work at pretty fast speeds, in some applications almost real time. For example, a detailed weather map, which would take a personal computer several days to compile can take a supercomputer just a few minutes to create, with the user monitoring the transformation as it occurs.

Three other technologies are particularly promising in achieving phenomenally fast computing in peta- and exascale supercomputers. The first of these is optical computing, where light are used to carry information, instead of slower electrons.

The second technology is known as DNA computing." Here, calculations are done by recombining DNA in different sequences, with those sequences that are favored and persist representing the optimal solution. The neat thing about DNA computing is that solutions can be deduced even before the problems have actually appeared.

The third technology is called quantum computing which exploits the properties of atoms (or

[1] For you folks keeping track of these things, a petaflop is a thousand trillion mathematical operations per second. Yeh, that's fast!

[2] Your quad-core XPS tower just got a lot less impressive!!

The Earth Simulator at the Japan Marine Science Technology Center. This most powerful of "vector" computers shocked some into believing that Japan could overtake the U.S. in many areas, just as it did in climate science. [Source: www.riken.go.jp.]

nuclei) designated as quantum bits, or quibits. A quantum computer will be capable of performing a computation by working on many aspects of a problem at the same time, on many different numbers at once, and then using these partial results to arrive at a single answer.

Deciphering the correct code from a 400-digit number would take a supercomputer millions of years. However, a quantum computer that is about the size of a teacup could do the job in about a year.

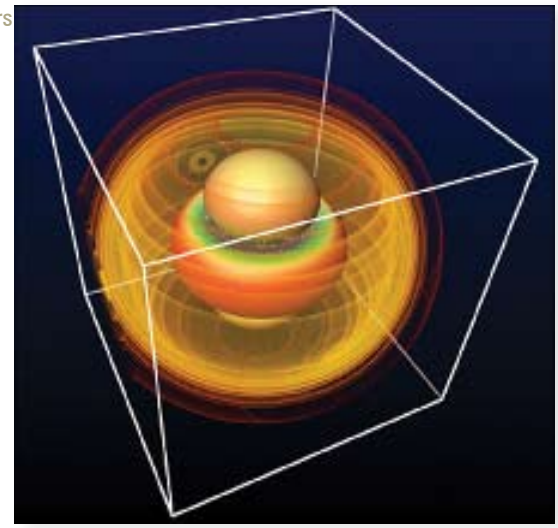

A visualization of gravitational waves generated with the Cactus Computational Toolkit, an astrophysical application requiring vast computing power. [Source: www.riken.go.jp.]

The petaflop

A leading candidate to become the first petascale machine is IBM's Roadrunner being developed at Los Alamos National Laboratory in New Mexico. Scientists say that Roadrunner and its cousins will make possible dramatically improved computer simulations. With Roadrunner's computing superpower they hope to shed new light on such subjects as climate change, geology, new drug development, dark matter and other secrets of the universe, as well as other fields in which direct experimental observation is time consuming, costly, dangerous, or simply impossible.

In fact, supercomputers and their simulations are becoming so powerful that scientific discovery might be seriously impeded or even compromised without them.

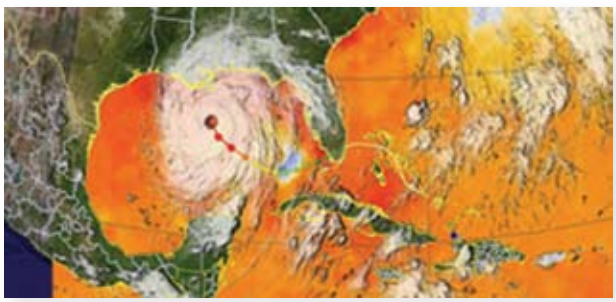

Simulating Katrina path. [Source: www.noaanews.noaa.gov]

A particularly fruitful area of computer modeling has been the study of global climate change. [See also p. 21.] In this field of study, computational science enables scientists to model atmospheric and oceanic conditions, and, importantly, how changes in each affect the other.

Petascale computers also will make it possible to predict the effect of an earthquake on every building in downtown Los Angeles. The increased detail could help shape building codes and be a valuable tool in evacuation planning and disaster preparedness.

Increasingly, supercomputers are being used for security purposes, such as the analysis of electronic transmissions for codes by the Echelon network of supercomputers and satellites. The computing power of the petascale machines will helps assess the performance and security of nuclear weapons—without conducting a real-life test.[3]

Outer space offers another infinitely mind-boggling opportunity for petascale computers—to help mankind understand some of the most baffling "inner workings" of the universe. ■

[3] Nuclear weapons are the quintessential example of an area where experimentation is not possible any more.