



Photo: SGI

Scientific Computing

The future of science will be driven by the improving performance of accelerators, telescopes, microscopes, spectrometers, and computers. The progress of scientists will depend on

how well these instruments leverage each other and the investigation process. Huge data sets threaten to overwhelm researchers in physics, chemistry, biology, climatology, and cosmology; bigger data sets are on the way as spatial resolution and sampling frequencies improve. Extracting knowledge from this "data tsunami" is the first priority of scientific computing.

The interactive visualization of massive data sets is perhaps the most useful approach. Auditorium-level visualization allows for multidisciplinary viewing and interaction between investigators; remote viewing and control can be achieved without having to transmit any of the data set. Interactive immersion within the data itself encourages discovery and exploration of singularities, discontinuities, artifacts, and irregularities that may not be obvious otherwise. New and insightful questions easily arise in this viewing environment because it is very intuitive.

Very large, complex models are readily viewable in this manner: the complete structure of a Boeing 747 aircraft, the seismic data of an underground oil and gas reservoir, and the 60,000 CT scans of a 2000-year-old Egyptian mummy.

Multi-level, multi-physics models can be merged into a common operating picture to enhance decision making, analysis, and discovery. This has been of real value to defense and homeland security analysts, and will soon become a necessary tool for natural disasters' prediction and response planning.

Nature is complex, non-linear, and non-homogeneous. The modeling, simulation, and visualization of her secrets require immense amounts of compute power, software capability, and data capture. Fortunately, we have Moore's Law and parallelism on our side. The rate of advance in compute power exceeds anything else in human experience.

More recently, emerging satellite and fiber-optic networks can be used to build computerized scientific instruments that straddle the globe. As for the architecture that binds the myriad of compute elements together, there will always be multiple choices. However, given that scientific data is generally real-time streaming data, it is essential that all data paths have high bandwidth and low latency characteristics.

Having large amounts of memory will help sustain modeling and simulation activities that may be taking place contemporaneously with data collection. NASA's space shuttle turbine, for example, creates 50 terabytes of data per second during shuttle-launch simulation. And the recently installed Altix supercomputer at Japan's Atomic Energy Research Institute includes 13 terabytes of globally-shared memory to simulate nuclear fusion.

Modeling, simulation, and visualization together create a third branch of human knowledge, perhaps on equal footing with theory and experiment, and the predictive power of mathematical models is accelerating scientific discovery. Rapidly increasing compute capabilities allow for finer granularity and model refinements: climate models now take into account terrain, rivers, lakes, vegetation, dust, cloud cover, and ocean salinity. The computer industry is striving to meet the challenge of nature's additional complexities. Japan's Earth Simulator, built by NEC and first operated in 2002, was considered a "Sputnik" event at the time. It delivered 35 teraflops of sustained performance, but was later superseded twice by more powerful US-built machines (first by SGI at NASA Ames, then by IBM at Lawrence Livermore National Laboratory). More importantly, the Earth Simulator has triggered a race to petascale performance levels, most likely achievable by the end of this decade.

In the next generation of supercomputers, new processing elements will augment the traditional ones, all fully integrated onto large amounts of globally-addressable shared memory, and all under the supervision of a single operating system, most likely Linux. Such machines will "reconfigure" themselves on the fly, and optimize their compute elements for the code at hand. They will form the high-performance nodes of larger grids, both national and international, that make up the compute platform and cyberinfrastructure for science in the next decade.

Beyond that, Moore's Law will reach its limit as semiconductor line widths approach the dimensions of a single atom, in which case "quantum computing" may kick in to boost performance for many machine generations yet again. By the middle of the 21st century, we will be ready to model the human brain and mimic its extraordinary thinking abilities. Consistently, scientists and engineers will find a way to out-compute in order to out-compete.

Bob Bishop

Bob Bishop is Chairman and CEO of SGI. His 40-year globe-spanning career involves active partnering with scientists and engineers in their efforts to accelerate discovery and innovation.