

PetaCrunchers

Setting a course toward ultrafast supercomputing

Computer simulation showing temperature fluctuations in layers of gas deep inside a sunlike star.

By IVARS PETERSON

Andrea Malagoli et al./Univ. of Chicago

Vast clouds of hydrogen gas. Swirling galaxies and brilliant quasars. Dust-enshrouded stars and superfast jets of ejected material. Bloated red giants and hot white dwarfs. Furiously exploding stars and wispy supernova remnants.

All these components contribute to the astronomer's view of the universe. Yet it's a limited, fragmented view. Telescopes and other instruments capture only the briefest of glimpses — mere snapshots confined to narrow slices of time.

Missing from this picture is the continuity of stellar evolution, from dusty cloud to nascent, fiery ball to mature star to death by collapse and perhaps explosion. Unable to observe the entire sequence on a human time scale, astronomers must find other means of assembling the puzzle.

Computer simulation offers a potential shortcut. Using high-speed, large-capacity computers, astronomers can test their speculations and theories: how gravity draws legions of stars into great spirals; why quasars burn so brightly; what drives a star to explode (SN: 2/18/95, p.106).

"Computer simulation is our only hope of turning astronomy into an experimental science," says Bruce A. Fryxell of NASA's Goddard Space Flight Center in Greenbelt, Md.

Unfortunately, today's most advanced number crunchers have neither the speed nor the storage capacity to handle more than a crude caricature of any given cosmic process. A millionfold improvement in computer speed and memory "would have an enormous impact on the field, permitting researchers to perform significantly more realistic numerical simulations," Fryxell contends.

Other fields have similar needs. Whether simulating turbulent air flow, modeling protein folding, mining huge stores of data for valuable nuggets of information, or visualizing interactions between brush fires, wind, and rain in a

Turbulence in explosively expanding material from a supernova.

Fryxell et al.

watershed, current technology lags far behind what researchers dream of doing.

In some ways, "we are still in the dark ages," says Thomas Sterling of the Center of Excellence in Space Data and Information Sciences at Goddard. "Our machines today actually do very little [compared to what we would like them to do]."

Last year, several federal agencies, including NASA, the Department of Energy, the National Science Foundation, the National Security Agency, the Ballistic Missile Defense Organization, and the Advanced Research Projects Agency, sponsored a meeting in Pasadena, Calif., devoted to exploring the feasibility of a giant leap forward in computer technology. Representatives of government, academia, and industry — self-characterized as a "constructive lunatic fringe group" — put forward their visions, quixotic and otherwise, of future computation.

This effort was followed by a second workshop, held in February in Fairfax, Va. "We're here to explore the far reaches of the computer frontier," Sterling declared at the meeting's start.

It was a chance for experts to peer at least 2 decades into the future, to do some creative, imaginative thinking while providing a realistic assessment of the opportunities, challenges, and critical elements of achieving truly high-performance computing.


Using pencil and paper, a person might take a minute or more to multiply 0.026 by 431.2 to get the answer 11.2112, manually keeping track of where to put the decimal point. Computers use so-called floating point arithmetic to race through the same calculation in just fractions of a second.

Floating point calculations have become so integral to computers that their performance is often measured in floating point operations per second (FLOPS). The new Cray T90 supercomputer can perform up to 60 billion calculations per second — operating, in other words, at 60 gigaFLOPS. Recently, a team from Sandia National Laboratories in Albuquerque and Intel Corp. in Santa Clara, Calif., set the world speed record for supercomputing, using two Intel Paragon computers to achieve a peak

performance of 281 gigaFLOPS.

At present, the federal high-performance computing and communications (HPCC) program has as a technical goal the achievement of computing at the teraFLOPS level by the end of the decade.

"But many of us think that's not the end," says Paul H. Smith, who heads NASA's HPCC effort. "There are significant applications that require more than



Modeling the impact of a comet with Jupiter's atmosphere.

just teraFLOPS computing."

The next logical step is petaFLOPS computing, a level of performance at least 10,000 times greater than the fastest of today's machines. It's an awesome leap into the unknown.


A report from the Pasadena workshop, which MIT Press will publish later this year, notes: "A petaFLOPS computer is so far beyond anything within contemporary existence that its architecture, technology, and programming methods may require entirely new paradigms in order to achieve effective use of computing systems on this scale."

Nonetheless, despite the daunting challenges ahead, "a petaFLOPS computing system will be feasible in 2 decades and will be important, perhaps even critical, to key applications at that time," the report predicts.

Reaching the goal of petaFLOPS computing in 20 years demands an immediate start in identifying and nurturing the incipient technologies that may ultimately determine the viability of such systems.

"Research agendas are becoming increasingly driven by immediate require-

Michael Warren et al./Los Alamos Natl. Lab. & Caltech



234

SCIENCE NEWS, VOL. 147

ments, with reduced attention to higher-risk, far-out ideas," Sterling maintains. "Without an over-the-horizon perspective, we may fail to embrace ideas of truly visionary merit because they do not address contemporary needs in the most cost-effective manner."

No computer can achieve petaFLOPS performance simply by having a single or even a handful of processors do lots of operations one step at a time, no matter how quickly. Instead, future computers will have thousands or millions of processors yoked together and working simultaneously.

Although such massively parallel computers exist today, no one is certain how many and what kinds of processors would have to be linked to reach the petaFLOPS performance level. Furthermore, at these speeds, even the tiny delays caused by the time it takes an electric or optical signal to travel from one place to another pose immense difficulties.

One approach to designing a petaFLOPS computer involves placing processors, which perform the arithmetic and logic operations, in the middle of memory chips, where the data are stored. In conventional computers, these two functions are normally found on separate sets of integrated-circuit chips, which must then be connected by wires.

"Processor-in-memory" (PIM) chips can serve as the building blocks of massively parallel computers, says Peter M. Kogge of the University of Notre Dame in South Bend, Ind. He estimates that a few thousand PIM chips, jammed closely together in a three-dimensional array, could attain petaFLOPS performance.

To demonstrate their approach's feasibility, Kogge and his coworkers at IBM created EXECUBE, a prototype computer made up of 64 relatively simple PIM chips. Each chip contains eight processors and 4.5 million bits of memory.

"A PIM-based architecture has the potential to achieve huge levels of performance with far fewer chips (and thus lower cost) than... other approaches," Kogge contends. However, the power demands of such arrays, as presently envisioned, could be so high that the heat generated by the device would readily melt it.

The PIM strategy represents a relatively modest step beyond conventional semiconductor-based electronics for computers. More radical, but potentially feasible approaches are also possible.

To sharply reduce power consumption and avoid overheating in electronic circuitry, some researchers are taking a fresh look at digital superconductor electronics. When chilled below a critical temperature, a superconductor loses its resistance to the flow of electric current.

Electrons can readily travel inside the material, allowing signal transfers at nearly the speed of light.

Because these circuits work with very small electric signals, superconducting devices can be packed together very tightly on a single chip, and the chips can be placed very near one another, says Konstantin K. Likharev of the State University of New York at Stony Brook. Such arrangements reduce delays when signals must travel within the circuitry or from one chip to another.

Furthermore, "superconductor fabrication technology is extremely simple," Likharev notes.

Likharev and his coworkers recently demonstrated that it's possible to use the tiny, isolated bundles of magnetic field — magnetic flux — that penetrate superconductors for rapidly storing and retrieving digital information. Earlier efforts to develop a computer using superconducting Josephson junctions failed, partly because the researchers — following standard practice in semiconductors — chose to encode binary data as high and low voltages.

"If you abandon information coding by voltage levels [and] use magnetic flux for this purpose, you can do everything very fast," Likharev argues. At the same time, power consumption goes down dramatically.

The real problem with a computer based on superconducting chips is refrigeration. Cooling with liquid helium is expensive. "This is why I don't believe that this technology will ever be in [personal computers] or even workstations," Likharev says. "It's something to be reserved for the high-performance end of computing."

Reaching petaFLOPS computing also means miniaturizing components beyond the fractions of a micrometer now readily achievable. "I don't think we can really build a machine that fills room after room after room and costs an equivalent number of dollars," says Seymour Cray of the Cray Computer Corp. in Colorado Springs, Colo. "We have to make something roughly the size of our present machines but with a thousand times the components."

Such an effort requires scaling things down from the micrometer to the nanometer range. Some computer designers are looking to molecular biology for working examples of what can be accomplished at this level.

Cray envisions two ways of riding the coattails of the recent revolution in molecular biology. Engineers might fabricate computing devices out of biological entities, he suggests. Alternatively, they could use biological processes to manufacture nonbiological devices — in effect, bioengineering bacteria to construct transistors.

Such schemes build on the notion of a

biological cell as an industrial park dotted with hundreds of protein factories, a smaller number of power plants, and an efficient railroad system for shuttling molecules from place to place.

It's a matter of learning how to harness these built-in capabilities.

Robert R. Birge, who directs the W.M. Keck Center for Molecular Electronics at Syracuse (N.Y.) University, has long studied the possibility of building biomolecular computers. He suggests that, in the near future, a hybrid technology involving both protein molecules and semiconductors could lead to computers just one-fiftieth the size of and up to 100 times faster than current ones.

Researchers are also exploring such options as using chemical reactions to process information. For example, John Ross and his coworkers at Stanford University have identified biochemical reactions that duplicate the basic logic functions from which practically any computer can be constructed. They can use various combinations of biochemical compounds and enzymes to get the results they want.

The potential for reaching petaFLOPS computing within 2 decades is there, Sterling insists. And good reasons exist for trying to achieve this goal.

This scale of computing power would permit tackling such problems as integrating data and simulation for designing new drugs or atomically precise nanostructures, then supervising the assembly of these products. Global climate simulations, models of large ecosystems, three-dimensional visualizations of complex physical systems, and tools for handling burgeoning databases — from satellite data to patient records — demand similar capabilities.

At present, the petaFLOPS computing effort is just a skeleton, Sterling notes. "We have to add flesh and muscle."

As one small step in this direction, he and his colleagues have just established an electronic database called PETA (PetaFLOPS Enabling Technologies and Applications) as an on-line reference index for the many topics, spanning a wide range of disciplines, that may have an impact on the realization or use of petaFLOPS systems.

"Researchers from around the country and the world working on concepts apparently quite unrelated can be joined in a single conceptual infrastructure binding them because of their potential contribution to realizing sustainable petaFLOPS performance," Sterling remarks.

The human brain itself serves, in some sense, as a proof of concept. Its dense network of neurons apparently operates at a petaFLOPS or higher level. Yet the whole device fits in a 1 liter box and uses only about 10 watts of power.

That's a hard act to follow. □