eton per hour. When he incubated the corals in the dark, the calcification rate of the algae-harboring *Galaxea* dropped dramatically. But in the *Tubastrea*, it remained high. Undercutting conventional wisdom even further, Marshall has found in more recent experiments that in the dark, calcium uptake in *Galaxea* was actually confined to areas that did not contain the algae, while in algae-loaded regions it ground to a halt. Rather than giving calcification a daylight boost, the zooxanthellae seem to be actively repressing it at night.

This repression might make sense if the algae—not the coral—were running the show. According to a model proposed by Ted McConnaughey of the U.S. Geological Survey in Denver, the process of calcification might be enhancing photosynthesis, rather than the reverse. Calcification, says McConnaughey, produces free protons that can be harnessed to convert bicarbonate in seawater into carbon dioxide to fuel the algae's photosynthesis. Marshall speculates that perhaps the algae manipulate the calcification process to increase the supply of these protons during the day when they need them for photosynthesis. At night, when photosynthesis can't occur anyway, the algae conserve the host's metabolic energy by shutting down the system.

Marshall concedes that no one knows how the algae might accomplish this. Moreover, the model doesn't address the more profound question raised by his results: If reef-building corals don't have higher calcification rates, why are they so patently successful—and their algaeless counterparts so unsuccessful—at building reefs? "According to this, *Tubastrea* should be all over the place. But it's not," says coral biologist Len Muscatine of the University of California, Los Angeles. He speculates that azooxanthellate corals like *Tubastrea* might be building more fragile skeletons which do not survive to become reefs, or perhaps the species is subject to more intense predation.

At the same time, he and others have some reservations about Marshall's conclusions. Muscatine points out that skeletal surface area, rather than weight, might have been a more accurate yardstick for comparison of calcification rates. And such data might bolster—or break—Marshall's hypothesis.

–James Shreeve

James Shreeve is a science writer in Takoma Park, Maryland.

COMPUTER SCIENCE

Microprocessors Deliver Teraflops

SAN DIEGO—In supercomputing's ancient past (about 10 years ago), the fastest machines were exotic beasts with custom-built processors, so high-strung that some of them could survive only when they were kept cool in a bath of freon. Now the world's fastest computer is being put together with truckloads of the same microprocessors that power high-end PCs.

At the Supercomputing '95 conference here in December, researchers from Intel Corp. in Beaverton, Oregon, showed off parts of several "nodes," or modules, from what is informally called TFLOPS: a massively parallel computer that will be able to perform mathematical operations at a peak rate of 1.8 trillion per second, or teraflops—about an order of magnitude faster than any existing computer. "All the silicon is in hand and working," says Intel Vice President Edward Masi: 9072 copies of Intel's next-generation PC microprocessor, called the P6 or the Pentium Pro.

The biggest remaining challenge in the $45.5 million project, funded by the Department of Energy's (DOE's) Accelerated Strategic Computing Initiative, is "the sheer logistics of building such an enormous system," says Masi. Although most of the machine will be pieced together from off-the-shelf components, it requires specialty chips to handle the stupendous volume of data flashing between the P6s. Meanwhile, Sandia National Laboratory in Albuquerque, New Mexico, where the TFLOPS computer is scheduled to be delivered in December, was obliged to create a custom operating system to control the behemoth without stealing too much memory from its main task—performing massive calculations.

All that computing power is meant to turn TFLOPS into DOE's latest nuclear testing ground. As a result of the Clinton Administration's decision to end nuclear testing in the United States, the weapons laboratories are relying on computer simulation to ensure the safety and reliability of existing weapons. The effort, part of DOE's so-called "stockpile stewardship" program (*Science*, 6 October 1995, p. 20), has a virtually bottomless appetite for computer speed and memory.

Researchers must not only do "virtual," or computational, testing of the explosions of existing warheads and new designs; they also face the greater challenge of modeling the astronomical number of different ways in which aging, stockpiled nuclear devices decay, how they can malfunction, and how they might fare in accidents. "Say the airplane [carrying nuclear weaponry] crashes and the fuel tank catches on fire. We have to model that," says Edwin Barsis, the recently retired director of computational sciences and mathematics at Sandia.

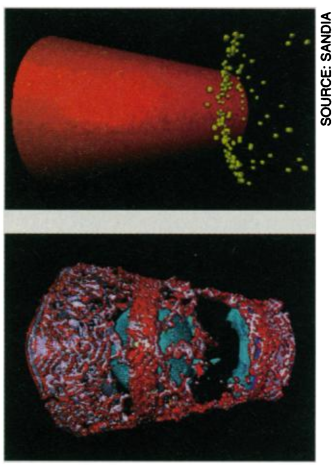

Cyberwar. Warhead-bearing missile payload collides with impactors in a parallel simulation.

TFLOPS will tackle these problems with a kind of "Russian doll" strategy, centered on the P6 processors. The P6s are grouped in pairs that share a common memory, forming nodes, or computation units; pairs of nodes share communications pathways on a single CPU board. Arrays of boards fit into about 80 cabinets, each about the size of a refrigerator.

Coordinating the 50 billion bytes of data that will flash through the computer each second are special-purpose network chips, one for each board. "There aren't any [existing] high-volume networks for machines like this," says James Tomkins, a computer scientist at Sandia. The computer's performance will also benefit from a software structure known as the lightweight kernel. Storing a complete copy of the processors' UNIX operating system on every node, says Tomkins, could soak up as much as 20% of the machine's total memory. The lightweight kernel reduces the burden by limiting the full operating system to only a handful of nodes, leaving the rest to subsist on a minimum set of instructions. Other areas of parallel computing are likely to benefit from these innovations, which "push that high end of computing," says Tomkins.

Users too will have to make some adjustments to wring the best performance from TFLOPS's parallel design. In some cases, says Charles McMillan, a computational physicist at Lawrence Livermore National Laboratory, this means rewriting codes completely, and in other cases, building in a kind of congruency. During the simulation of a nuclear explosion, for example, fluid-dynamics and chemistry calculations have to ask for data on changing gas temperatures at similar times from similar storage locations around the far-flung machine.

Researchers will also need to devise ways to spot-check TFLOPS's predictions against laboratory data from, say, laser-fusion experiments. There's good reason for such quality controls, says Sandia's Grant Heffelfinger. "It used to be that [supercomputer calculations] provided only physical insight" into how a weapon would work. The goal now, he says, is "a detailed mirror of reality."

–James Glanz