$200 million in mid-November to settle the suit could the Max Planck investigation finally conclude. While litigation was pending, UC did not answer Max Planck inquiries about whether Seeburg was allowed to take the samples, says Klaus Hahlbrock, Max Planck vice president and a member of the investigating committee. UC ultimately acknowledged that, at the time, there were no unequivocal regulations barring Seeburg from taking the samples, he says. Seeburg himself declared repeatedly that he felt he was entitled to do so according to "most scientists' ethical standards."

Although the committee didn't accept that argument entirely, the investigation focused on the alleged falsified information in the 1979 *Nature* paper. Seeburg's admission was hotly contested by co-author David Goeddel, then at Genentech and now chair of the South San Francisco biotech company Tularik, and several other former colleagues. The Max Planck committee took Seeburg's admission at face value, however. The committee concluded, says Hahlbrock, that "a falsified description in a publication cannot be tolerated, no matter if it dates back 20 years," and recommended that Seeburg be censured —a rare and rather exceptional measure. The censure does not directly affect Seeburg's position at Max Planck, says Markl, but it "will be put into his personnel record."

For his part, Seeburg told *Science* that he may donate part of the $17 million he and four other former UCSF researchers will each receive as part of the settlement to a charity or research foundation. But most of all, he hopes that his censure will be the final episode in this painful saga and that he will once again be able to "concentrate on doing science." **–MICHAEL HAGMANN**

With reporting by Robert Koenig in Berlin.

COMPUTER SCIENCE

Big Blue Aims to Crack Protein Riddle

IBM last week announced a $100 million research initiative to build a supercomputer 500 times more powerful than the current record holder. Dubbed "Blue Gene," the technology test-bed's initial goal will be to model how proteins fold into the three-dimensional shapes that allow them to orchestrate life within the cell. If successful, the 5-year effort could allow drug researchers to go right from the sequence of a disease-related gene to the predicted structure of its protein, in order to identify targets for therapeutic drugs. Down the road, Blue Gene and its kin could also revolutionize other computationally intensive disciplines, such as modeling climate change and the evolution of galaxies.

Researchers who model protein folding agree that Blue Gene's ability to run 1 quadrillion (10^{15}) mathematical operations per second (also known as 1 petaflops) will be a big step for the field. "Petaflops computers certainly make you salivate," says Stephen Mayo, a protein-folding expert at the California Institute of Technology in Pasadena. It won't be easy to serve up this feast, however. Supercomputers built by IBM and Intel for the national weapons laboratories currently reign as the world's fastest, at 2 trillion operations per second (2 teraflops). "But there's no way to get up to a petaflops using [the same] technology," says Monty Denneau, a mathematician at IBM's T. J. Watson Research Center in Yorktown Heights, New York, and Blue Gene's chief architect.

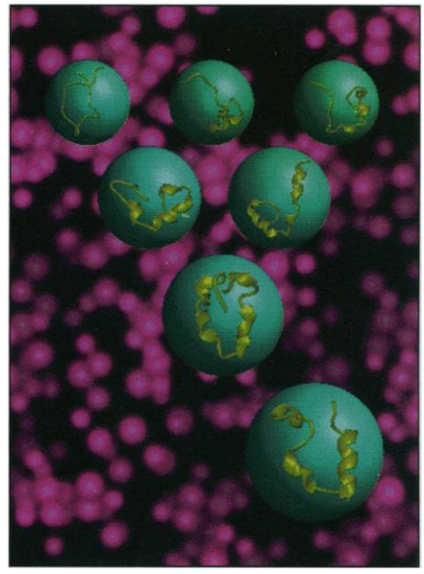

Nature makes it look easy. The bubbles illustrate stages in the folding of a protein—a process that researchers would like to predict.

The chief obstacle is power consumption and heat: Denneau says a petaflops machine that used the same amount of energy for each operation as current teraflops machines "would take a dedicated [power] reactor"—and would quickly immolate itself. To speed computation while cutting power consumption, IBM plans to come up with an "ultraminimalist approach" for both hardware and software, which will reduce the complexity of the processors but increase their ability to communicate and work in tandem. Both the new chips and the software are still on the drawing boards, but "I think the plan makes a lot of sense," says Arvind, a computer architecture expert at the Massachusetts Institute of Technology, who goes by a single name.

As a test for their new machine, Denneau and his colleagues have chosen one of the toughest challenges in biology. Inside cells, newly synthesized chains of amino acids take a second or less to fold into a functional protein. Every one of tens of thousands of atoms in the chain and surrounding water molecules pulls or pushes on its neighbors to determine the final shape. But even though researchers have measured the forces between atoms in great detail and can easily predict how a handful of amino acids will interact, precisely modeling the folding of proteins has been out of reach.

The most ambitious efforts, called all-atom simulations, calculate the interatomic forces and their effects for every possible pair of atoms in the protein chain. Even for a small protein, an all-atom simulation of just a fraction of the folding process takes months of supercomputing time.

IBM's answer, Blue Gene, will consist of more than 1 million processors, each capable of 1 billion operations per second (1 gigaflops), assembled on 64 racks. To reduce power consumption, IBM researchers are doing away with a type of fast but power-hungry on-chip memory called cache. In its place, they're moving onboard a slower type of memory called DRAM, which is traditionally located off the chip. This should bring big power savings by eliminating the need to send off the chip for additional data.

To compensate for the slower memory speeds, the IBM team is planning to turn to a technique known as multithreading, the computer equivalent of a multitasking commuter who eats breakfast, drives, and talks on the phone all at the same time. In this case, each processor will work simultaneously on eight separate problems. That way if a processor is waiting for a bit of data to come in from memory to complete one computation, it can still be working on others at the same time. Finally, because chip failures are inevitable in an array of a million processors, Blue Gene's software will be designed to reroute data to working devices if a processor or connection fails in midsimulation.

Even with a new petaflops machine, it will take about 1 year to simulate the complete folding of a typical protein. And even then the protein-folding problem may not be solved. Peter Wolynes, a protein-folding expert at the University of Illinois, Urbana-Champaign, explains that the all-atom approach of computing the interactions between pairs of atoms may not be enough. It may turn out that to get the right answer, researchers will have to compute the interactions among many atoms at once as they tug and push on each other, which would vastly increase the problem's complexity and require still more computing muscle. "My suspicion is that you won't need all that additional stuff," says Wolynes. But if it turns out you do, at the end of the day "you would have learned that you can't solve it."

–ROBERT F. SERVICE

CREDIT: KOLLMAN *ET AL.*, UCSF