heads the Department of Energy's petaflops effort. Thousands of those clusters would be loosely linked to produce a petaflops computer. The design will require extraordinarily fast communications between the clusters, however. "We're going to have to work through and overcome all these interconnection problems," says Smith. "That's not trivial."

Petaperipherals, petaprogramming

All three of these concepts, adds Smith, are only "candidate or example architectures, and we don't expect to be locked in by those three." Among the other NSF-funded projects, for example, is a University of Illinois design, in which the software can actually manipulate the hardware during operation, in effect continuously rewiring the computer to match its hardware structure to the algorithm in use at the time, thereby reducing latency. Another candidate comes from a Purdue University group, which is working on an architecture called variable granularity processor hierarchy, which relies on multiple small processors for most processing tasks but calls on a few very powerful but power-hungry processors when needed.

The NSF is also studying exotic technologies to store and retrieve the oceans of data a petaflops machine would generate. Caltech's Sterling, Likharev, and collaborators, for example, are exploring a hybrid machine that would link superconducting logic to a primary memory (the equivalent of RAM in a desktop computer) based on optical holographic storage. The memory would capture data in a medium that records interference patterns generated by intersecting laser beams. Ten to 100 terabytes of data could be encoded in those patterns and packed into a cubic centimeter of the storage medium. And when a "read" beam shines on the storage material, adds Sterling, "you can get all that information in big shots, megabytes per access."

The downside is that lab prototypes of holographic memory suggest it will probably take microseconds to deliver the data—100 times longer than the best silicon memories of the next decade. So the computer would need a hierarchy of different kinds of memories, from fast silicon to slower holographic ones. "The trick is to have the memory system itself bundle up large bunches of data and send it all at once to the processors," says Sterling.

Then there is the problem of displaying the output of a petaflops computer to the user, which could mean somehow communicating 100 terabytes of data. Ian Foster of Argonne describes this as a "problem of trying to provide a thick pipe from the supercomputer to the outside world," with the most likely solution being some kind of three-dimensional display, either virtual reality or what is known as a cave, in which the operator is actually inside a room-sized display device.

If all that isn't trouble enough, says Stevens, "At least half of the work, and probably 70%, is going to be software." A petaflops computer will need what by today's standards is an enormous degree of control over the internal data movement. "This is the hardest possible problem in parallel computing," says Stevens, "an efficient way of managing data movement throughout the machine." Part of the challenge is orchestrating the million-way parallelism of such a machine, but the other part is moving data between what is likely to be six to eight levels of memory, each one getting progressively faster as it gets closer and closer to the processors. "We want to start software development in the next couple of years," says Foster. "There's no point having great hardware unless you have pretty great software."

All of this will take another step toward reality next month, when the NSF-funded architecture projects will present their findings at the Frontiers of Computing 1996 meeting in Annapolis, Maryland. At that point, says John Toole, director of the National Coordination Office for High-Performance Computing, "they'll bang each of these ideas up against computer architects, software people, applications people, to really get some synergy about what the real meat of problems will be for the long term." Then it will be up to the funding agencies to decide whether to support a full-scale petaflops project—an effort that is likely to cost $400 million each year for a decade, twice the spending on supercomputers of the current high-performance computing initiative.

If the government decides against it, industry isn't likely to step in on its own, say the petaflops researchers. Many supercomputer-makers have fared poorly in recent years. And the kind of focused effort necessary to achieve a petaflops within 10 years, says computer scientist Geoffrey Fox of Syracuse University in New York, is not the kind of effort that is likely to be driven by market forces. "Rather it's like the atom bomb or the space program," he says. "Something outside of an industrial endeavor, but of importance to the nation."

–**Gary Taubes**

NEUROBIOLOGY

Glimpsing Myelin's Protein Glue

A nerve signal can zip from your spinal cord to the tip of your toe in less than 25 milliseconds. But such rapid nerve transmission is only possible because the axons, the long neuronal projections that carry the signals, have very good electrical insulation. Layers of tightly packed cell membranes wrap the axons much like gauze wrapped around an injured finger, forming an insulating sheath known as myelin. Two papers appearing in this month's issue of *Neuron* now help explain how these tightly wrapped layers are glued together in peripheral nerves and how some mutations weaken the glue, leading to neurological disease.

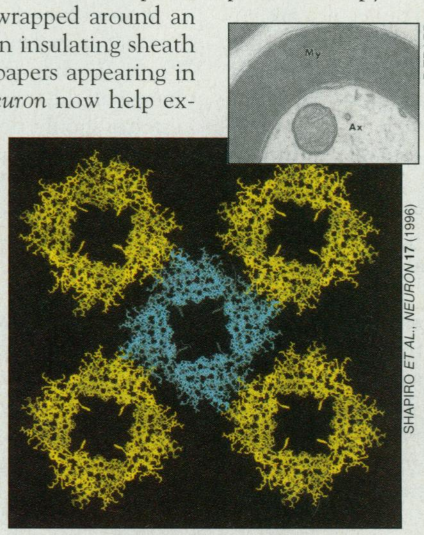

B.TRAPP

SHAPIRO *ET AL., NEURON* **17** (1996)

Sticky structure. The myelin surrounding nerve fibers *(inset)* is held together by interlocking P_0 tetramers.

By determining the crystal structure of part of the glue, a myelin protein known as P_0, a team including Lawrence Shapiro and Wayne Hendrickson of Columbia University in New York City and David Colman of Mount Sinai Medical School, also in New York, has found that the protein molecules apparently interlock to form a sort of molecular Velcro between the myelin membranes. And in work described in an accompanying paper, Laura Warner and James Lupski at Baylor College of Medicine in Houston and their colleagues use the new structure to analyze the effects of some of the 29 known disease-causing mutations in P_0, including five new ones they discovered. While the work doesn't open immediate avenues to improved therapy for the diseases, it does provide insights into how the P_0 mutations can cause symptoms ranging from poor coordination to paralysis.

"Here is an example in which nature has already done the function part of a structure-function analysis for you," says neuroscientist Greg Lemke of the Salk Institute. "There are very few examples like that where you have a large panel of mutations that have differing effects, and you have a very nice, high-resolution structure you can map them onto."

Healthy myelin appears under the electron microscope as "beautiful perfect spirals with dozens of wraps of membrane, each with the exact same spacing," Shapiro says. Researchers have known for several years that P_0 seals the membranes together with that perfect spacing. Without it, as Melitta Schachner of the Swiss Federal Institute of Technology in Zurich and her co-workers found in 1992 when they knocked out the P_0 gene in mice, peripheral nerves have virtually no myelin, only remnants of loose, floppy membranes poorly wrapped around some axons.

Exactly how P_0 bonds the myelin membranes together was unclear, however, and in 1992 the lab teams of Colman and x-ray crystallographer Hendrickson joined forces to try to find out by solving the protein's structure. Before the structure could be determined, however, the protein had to be crystallized, and that promised to be a stumbling block. P_0 spans the membrane of Schwann cells, the specialized cells that form myelin by wrapping themselves around the axons of peripheral neurons, and membrane-spanning proteins are virtually impossible to crystallize. So Shapiro, a graduate student with Hendrickson, crystallized just the part of P_0 that protrudes outside the cell.

Subsequent x-ray analysis of those crystals showed, Shapiro says, that they were formed of "layers of molecules on the order of 45 angstroms [Å] thick." Because that is the spacing seen between the myelin membranes in electron micrographs, he adds, "this led us to suspect that we might in fact be looking at a reproduction of the layer that forms the glue between the membranes in the extracellular space of myelin."

The analysis also revealed the structure of the individual molecules making up those 45 Å layers—and suggested how they might act as the myelin glue. The structure suggests, says Shapiro, that the P_0 molecules are anchored to the Schwann-cell membrane like balloons on strings, bunching together in groups of four called tetramers. Each tetramer interleaves with four tetramers protruding from an adjacent membrane, linking the membranes together.

But because the strings are flexible, that interaction alone would not hold the membranes at a fixed distance. That is where an additional feature of the protein apparently comes in: Protruding from the top of each protein balloon, opposite the string end, is the amino acid tryptophan, whose hydrophobic nature makes it want to bury itself in fatty cell membranes. That tryptophan may firmly anchor each tetramer into the opposite membrane. And together with the tetramer-tetramer interaction, says Colman, that would "rigidly specify the distance" between the membranes at roughly 45 Å. Structural biologist Pamela Bjorkman of the California Institute of Technology concurs: "All the distances … agree with the distances in myelin. [The structure] is quite convincing."

Knowing the structure gives researchers a window onto a group of human diseases in which myelin forms improperly in the peripheral nervous system. Among these are Charcot-Marie-Tooth disease (CMT), Dejerine-Sottas syndrome (DSS), and congenital hypomyelination (CH). These conditions produce different symptoms—ranging from paralysis at birth in the most severe cases of CH to the weakness and lack of coordination that develops in young adults with CMT—and they were originally seen as separate diseases. But in the past few years researchers have found that some cases of CMT and DSS are caused by mutations in the P_0 gene, and in the current work, Lupski's team has added one case of CH to the list. "There is a spectrum of severity," says Lupski, but the conditions in his view are essentially the same disease. And that raises the question of how mutations in the same gene can cause such disparate symptoms.

The new structure should help answer that. For example, among the mutations are several in which the amino acid cysteine replaces another amino acid in the protein. Like most of the disease-causing mutations in P_0, these mutations cause symptoms when present in only one gene copy. And in these cases the symptoms are severe—the patients become ill as children with a condition classified as DSS. Baylor's Warner says that the site of the amino acid replacement in the P_0 structure may explain why. Based on their location in the structure, the cysteines should protrude from the protein's surface, enabling them to form disulfide bonds with cysteines on other P_0 molecules, perhaps creating inappropriate aggregates of P_0 and disrupting the normal P_0 interaction.

Warner found another cysteine substitution that caused milder CMT-like symptoms. When Warner and Shapiro located this change on the structure, they found that the cysteine should point into the protein's interior, where it can't make mischief by reacting with other P_0 molecules. It might affect the protein's function in more subtle ways, or even prevent it from reaching the cell surface. And a mutation that just reduces the amount of P_0 or the effectiveness of the mutant P_0 molecules should not be as bad, says Warner, as one that actively disrupts P_0 interactions.

Such reasoning is not proof of how the cysteine mutations cause disease, but it is a significant advance, says Bruce Trapp, chair of neuroscience at the Cleveland Clinic, because it produces "testable hypotheses" about how structural changes make myelin come unglued. Researchers can determine the crystal structures of specific mutant proteins, for example, or they can express the mutant proteins in cultured cells to see how they behave. And as they learn how myelin can come apart, they may find new hints about how to repair it.

–Marcia Barinaga

WEATHER FORECASTING

Budgets Stall But Forecasts Jump Forward

Anxious farmers and sodden picnickers may not believe it, but weather forecasts are actually getting better. By giving computer models such tweaks as a sharper picture of atmospheric properties or more realistic processes for turning water vapor into rain, meteorologists have made their forecasts much more accurate over the last 10 or 20 years; their computer models can now make useful predictions of the atmosphere's chaotic behavior up to a week in advance. Most Americans aren't reaping the full rewards of this improved forecasting skill, because official daily forecasts stop at 5 days. But in a year, the official forecasts may reflect the computational successes.

The operative word is may. The U.S. National Weather Service (NWS) plans to extend its daily forecasts to a full 7 days ahead by October 1997 and its less precise 6-to-10-day outlook out to 2 weeks. Yet this forecasting milestone, which should benefit everyone from natural-gas suppliers estimating heating needs to forest firefighters looking for relief, could be diminished by federal cost-cutting that would take human forecasters out of the picture and delay the arrival of the next powerful forecasting computer. "We have increasing pressure from the public to improve our forecasts," says James Hoke, director of the NWS's Hydrometeorological Prediction Center in Camp Springs, Maryland. "At the same time, we're being cut back significantly budgetwise. It's a difficult time."

Hard times have hit the weather service just when decades of incremental increases in forecasting skill are finally adding up to major advances. "A day-3 forecast [i.e., for the third day ahead] now is as accurate as maybe a day-1 forecast was in 1960," notes Russell Martin of the NWS's Climate Prediction Center in Camp Springs, who keeps statistics on forecasting skill. And improvements can be found at longer ranges too. Day-5 forecasts of the high- and low-pressure centers that determine weather patterns are now nearly as accurate as day-3 forecasts of 20 years ago, says Hoke. Predicting the period 6 to 10 days ahead is harder, but since that type of forecasting was begun in 1979 "we've about doubled our skill with temperature," says Martin.

Forecasters themselves can't take all the credit for this—much of the improvement is due to increasing computer size and speed, says Martin. "Being able to run bigger, more detailed computer models has been the big