Lab Notes

Drosophila melanogaster, the geneticist's friend.

Channeling

You have to go through channels to get things done in any well-organized bureaucracy, and the nervous system is no exception. Nerve cells "fire"—conduct an electrical impulse through themselves—by creating a traveling ripple in their internal ion balance. To do so, each cell moves ions in and out of itself through channels—protein molecules spanning the cell membrane. Each molecule is designed to admit a particular ion: sodium, say, or potassium. The channels are such an infinitesimal proportion of the cell's protein—perhaps one molecule in a million—that the first one, a relatively abundant sodium channel, wasn't isolated and purified until the late 1970s. It took almost another decade to discover the gene responsible for producing it.

Last year a group led by Assistant Professor of Biology Mark Tanouye located the gene responsible for potassium-ion channels in *Drosophila melanogaster*, the fruit fly. Cellular conductivity studies had indicated that there were many different types of potassium channel per cell, implying that each individual protein would be correspondingly rarer. So instead of taking the conventional approach—isolating the protein, determining its amino acid sequence, and using this sequence to find the corresponding DNA sequence in the chromosome—the group took a novel tack. They zapped fruit flies with enough x-rays to jumble their genes just a bit. Some mutant offspring had aberrant potassium conductivity, and these flies were examined for visible chromosome damage to find the gene's general neighborhood. Then the researchers "walked" an overlapping series of DNA-binding probes along the chromosome to reach the gene's exact address. The gene resides within the "Shaker locus," a region named by earlier gene mappers because mutations therein produce twitching flies.

The group has since used the *Drosophila* gene to find corresponding genes

Right: A set of Drosophila chromosomes. An x-ray dose has broken the x-chomosome at the "Shaker locus," and reattached part of the x-chromosome to chromosome 3. Far right: In this close-up, the two sets of black blotches mark where a radioactively labeled probe has bound to both fragments of the Shaker locus. The lower blotches show where the x-chromosome (extending off to the left) has fused with chromosome 3, which curls away to the right and down.

in rats, and recently in humans, by assuming that vital parts of each gene would be sufficiently similar that a probe able to recognize and bind to one would also recognize and bind to the other. So far they've found six different channels in human brain tissue. The search also turned up a channel peculiar to rat heart tissue, and is closing in on the human equivalent. Potassium-channel-blocking drugs are given to cardiac patients to control heart arrhythmias. Unfortunately, these drugs block potassium channels throughout the body, causing all sorts of unpleasant side effects. A drug that blocked only heart-tissue potassium channels would create far fewer problems.

Surprisingly, while the channels in flies, rats, and people are quite similar, and are produced by similar bits of genetic code, the code comes in different formats. *Drosophila* uses one long gene that resembles a Chinese menu. Column A contains six initial segments, any one of which may be chosen when creating a channel type. Column B contains a single midsection, and Column C has four terminal segments. The fly creates some 20-odd channel types by mixing and matching segments. Mammals, however, have a separate gene for each channel type. Tanouye estimates that there may be as many as 100 different types of human potassium channel.

One hundred flavors might seem to be too much of a good thing, but it really isn't. There are four basic types of nerve impulse, or "action potential." The nerves running to the central nervous system (the brain and spinal cord) typically fire isolated impulses lasting about one thousandth of a second each, while central nervous system cells fire complex bursts lasting for hundredths of a second. Heart tissue "plateaus," maintaining an elevated action potential for half a second at a time to drive the pump stroke. "Pacemaker" tissue provides regularly repeated pulses, over and over. "There are different waveforms within each impulse category," says Tanouye. "But the rising phase, which is generated almost totally by a fast influx of sodium ions into the cell, is always the same. So to make each category and the small variations within it, the cell sculpts the falling phase by modulating the outward current of potassium ions. And a slow inward current of calcium ions keeps the potential high to make plateaus. That's why there are so many kinds of potassium channels, but only a few types of sodium and calcium channels. Each different cell has its own distribution of potassium channels to get the right waveform, which can be fairly complex."

Channels go through a three-step cycle: activation, which enables them to

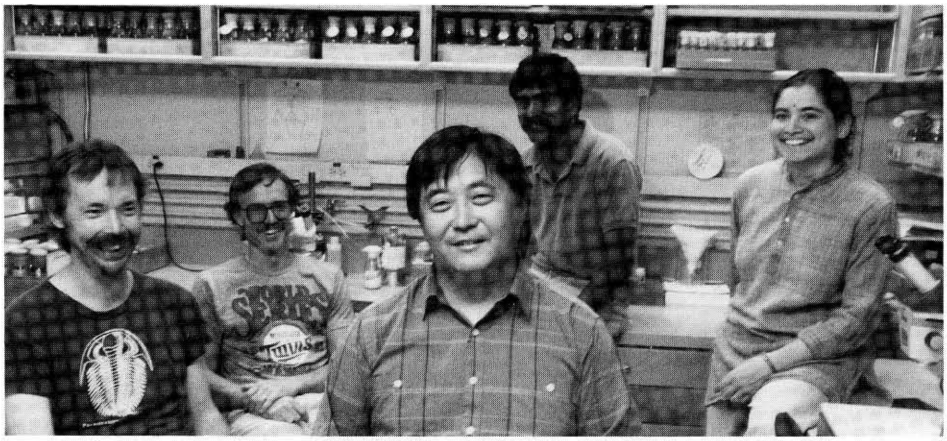

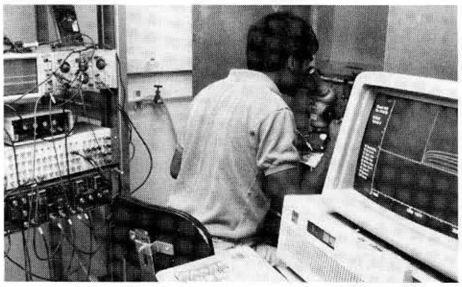

Above: Measuring an individual cell's conductivity is exacting work. A microelectrode is inserted into the cell, using the microscope to guide the hand. Conductivity profiles appear on the computer monitor at right.

Right: The Tanouye group. From left: William Trevarrow, Ross McMahon, Tanouye, Mani Ramaswami, and Mehda Gautam. (Missing: Mathew K. Mathew and Ken McCormack.) The flasks on the shelves in the background are home to various Drosophila strains.

pass ions; inactivation, which stops ion passage; and recovery, during which an inactivated channel resets itself to be activated again. Activation and inactivation are controlled, or "gated," by the voltage differential on either side of the cell membrane. Potassium channels vary in their gating voltage, and some channels need to have calcium or magnesium ions, or messenger molecules such as serotonin, present as well. Activation, inactivation, and recovery rates also vary. Tanouye's group has found that in *Drosophila*, choices from Column A (the so-called 5′ end of the gene) build in the channel's inactivation rate, while Column C (the 3′ end) sets the recovery rate.

The constant region (Column B) presumably encodes features that don't change much, such as ion selectivity. "We really didn't want to make a catalog of potassium channels per se, but we had to look at a collection of them to find the natural variations of structure and function in the constant region, as well as going after human channels of clinical significance."

Some things have already been learned. Other researchers have found a repeating amino acid sequence lying squarely in the middle of every ion channel found to date, in a region christened S4. S4 is believed to be the channel's voltage sensor.

Tanouye's group has found another region, overlapping S4 a little bit and continuing into the channel's interior, called a "leucine zipper." A leucine zipper contains the amino acid leucine followed by six others in a sequence repeated four to six times. When the protein coils into its natural shape, all the leucines line up along the coil like teeth in a zipper. Leucine zippers are believed to play a role in DNA-binding proteins, another hot area of molecular biology. What the zipper does in the ion channel remains a mystery, but Tanouye speculates that it may be part of the actual gateway. "We've found leucine zippers in every single potassium channel so far, and sodium and calcium channels also have zippers. So we thought, naively, if the S4 region moved a little bit in response to a voltage change, the channel might unzip so ions could go through. It's probably more complex than that."

To find out, Tanouye's group is now making channels with the zipper leucines replaced by the closely related amino acids valine and alanine. "These are really very subtle changes, to another hydrophobic amino acid that's somewhat smaller. But we've found that gating is strongly affected, in voltage sensitivity and other things. Now we're looking for the logical framework, the story of what this all means."□—*DS*

Bid Me Up, Scotty

"NASA's problem is to get the information needed to make the best use of scarce resources, and normal bureaucratic processes simply can't do it."

When the space station opens for business, it will have some room for commercial payloads. And if it's treated like the shuttle, that room will be allocated haphazardly to all comers. NASA tries to ensure that the best payloads fly, but their selection system can have unintended consequences. Relatively worthless payloads may go up while better cargoes languish in warehouses. And there's no built-in incentive to conserve the spacecraft's resources; thus the payloads that fly may squander what could be better used by others.

NASA doesn't always know which payloads are the best, nor if their design could be improved. It's not that easy to find out, says John O. Ledyard, professor of economics and social sciences. "A commercial payload's value includes both its immediate benefits—projected cash returns—and its long-term benefits—perhaps research results leading to marketable products in 20 years. Some long-term benefits are unforeseeable, but most firms have a pretty good idea of their payload's worth. They don't want to share this information because it's proprietary. And everyone wants a bargain, so if you just ask, 'How much are you willing to pay to fly this?' they'll say, 'Well, I can't afford much, but my payload is really important.' NASA's problem is to get the information needed to make the best use of scarce resources, and normal bureaucratic processes simply can't do it. A properly designed pricing strategy will."

Instead, NASA's payload-selection procedure has been divorced from its pricing policy. Engineers allocated shuttle space as best they could, evaluating payloads based on their own experience. Then the accountants sent a bill to cover launch costs. This cost-based pricing has its roots in "marginal-cost" pricing, developed in the 19th century to help set bridge tolls. The marginal cost for a bridge built to carry 100 cars a day is the extra cost of carrying the 101st car. The marginal cost of a shuttle payload is the cost NASA incurs beyond the cost of launching the shuttle anyway, sans payload. Bridges have been around for centuries and the rules for finding their marginal costs are well known, but the shuttle is so new that its marginal cost is still being debated. So NASA guessed at a price, and, to ensure a clientele, probably guessed too low. Low prices may be fine for abundant resources, but not for an infrequent-flier shuttle, or for a space station, where you can't just build another room over the garage. Resource allocation becomes first-come, first-served. Nothing prevents the first arrival from claiming *all* the resources, preempting the competition.

Ledyard, who in 1983 joined a group studying pricing policies at the

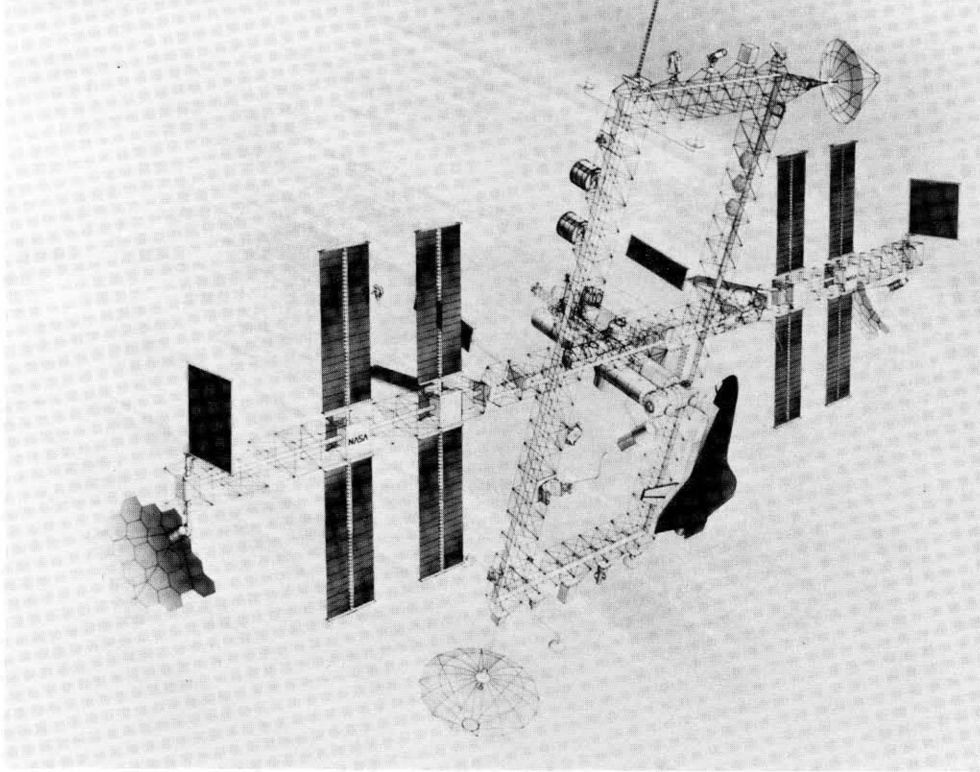

Low prices may be fine for an abundant resource, but not for an infrequent-flier shuttle, or for a space station, where you can't just build another room over the garage.

Jet Propulsion Laboratory (JPL), thinks there's a better way. "Any economist knows that pricing policy and resource allocation are intimately linked." Those gaining the most from a scarce resource will pay the most to secure its use, so auction it off. Assuming the bidders have *some* idea of their potential benefits, the bids become proxies for the payloads' real worth. The winning bids reflect the "opportunity cost" of the payloads that don't fly—the benefits lost to the unsuccessful bidders. Such a system is called "demand-based" pricing.

"The fact that the winners paid that price isn't as important, from the public-policy point of view, as the fact that they got on," says Ledyard. "The bidding has indirectly sorted out the good proposals from the bad ones."

In its simplest form, this isn't a particularly new idea. Cattle are sold at auction, as are soybean futures and van Goghs. When a single commodity is being sold, it's fairly easy to figure out how to bid. But a shuttle berth involves several "resources": weight, volume, electrical power, manpower, and other factors come into play. Each payload has specific requirements—a communications satellite might be large and heavy, but need no electricity and take only one man-hour to launch, while a compact crystal-growing project might draw lots of power and require constant human attention—and it's pointless to fly a payload if all its needs aren't met. It's impractical to auction each resource, as bids for any one item depend on the prices of the others. Even with a computer tracking all the various auctions, most people would suffer brain failure trying to plan their next bid.

Ledyard proposes an adaptive user selection mechanism, or AUSM (pronounced "awesome"). Each bidder submits a package, containing one bid for a list of resources, to a computer. Like a camper with more gear than will fit into a knapsack deciding what to pack, AUSM sorts through the bids to find the highest bid (or bids) whose combined resource demands can be accommodated. The highest bid always wins in a simple auction, but with the knapsack problem this isn't necessarily true; if 1,000 cubic feet of space are available, say, 10 bidders offering $100 each for 100 cubic feet will beat one bidder offering $700 for all the space. Thus many small bidders flying modest projects can collectively outbid a mammoth communications satellite. In practice, AUSM accepts every bid until all available resources are committed. Then prospective users must displace one or more payloads already on board by outbidding them.

The system could run for months, allowing users who've been bumped

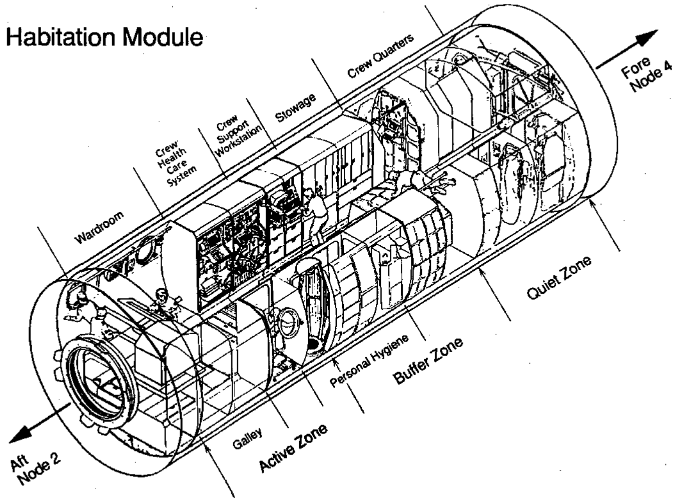

"You can get a huge bang for your buck when people start redesigning their payloads to fit better."

to refine and resubmit their bids based on the current roster of successful bids. Says Ledyard, "You can get a huge bang for your buck when people start redesigning their payloads to fit better. [A fixed bid and a scaled-down resource demand is tantamount to a higher bid, encouraging efficient resource use.] We can measure that bang experimentally."

Ledyard and Charles R. Plott, Harkness Professor of Economics and Political Science, use Caltech's Laboratory for Experiments in Economics and Political Science to test AUSM against other pricing systems, including cost-based ones like NASA's. The lab allows researchers to study economic and political behavior under rigorously controlled conditions. An experiment can include up to 20 people linked by a network of PCs.

In the first experiments, seven "payload managers" could choose to sponsor one of several possible payloads. Each payload needed a different mix of resources and promised various rates of short- and long-term return with an associated probability of failure. Managers could bid, rebid, alter their payloads, or even choose new ones as the computer noted their every move. After a set interval, the computer closed the auction and "launched" the shuttle with the winners' payloads aboard. The computer calculated how well these payloads performed in orbit, paid their managers accordingly, and began the cycle again. The managers were paid real cash, giving an incentive to succeed.

Unlike NASA, the experimenters knew every payload's true value (rate of return times probability of success). They measured a pricing system's ability to find the best payloads by the ratio of the value of the payloads that flew to the highest possible value attainable from any flyable combination of payloads. The cost-based mechanism a la NASA was about 65 percent efficient. AUSM was about 90 percent efficient.

Ledyard had spent two and a half years trying to sell the AUSM theory to NASA brass, engineers with a healthy skepticism of economics in general. It was an uphill struggle—a complicated issue challenging many vested interests. He'd penetrated several layers of bureaucracy with no end in sight when he, Plott, and the JPLers made one more trip to Washington. "Plott put up a viewgraph with the two data points on it and said, 'See, this is how it works.' And all the NASA people said, 'Wow! That's great!'" Ledyard recalls. "Suddenly they were willing to listen. The power of experimental analysis to convince people who otherwise don't understand economics is just amazing."

More proof came in a few months, when Ledyard and Plott ran a pricing experiment on NASA managers. The

"We simulated NASA's cost-based policy, and we warned them that the highest-priority bidder would try to grab everything. They said, 'Scientists don't act that way. That's crazy.'"

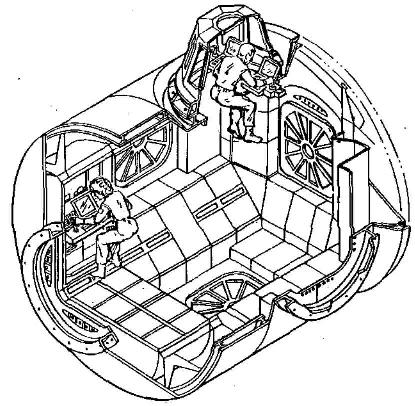

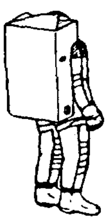

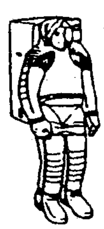

relative efficiencies held true, and NASA folks acted just like everyone else. "We simulated NASA's cost-based policy, drawing numbers from a hat for the first-come, first-served aspect, and we warned them that the highest-priority bidder would try to grab everything. They said, 'Scientists don't act that way. That's crazy.' And 15 minutes into the run, one guy was doing it. We asked him afterward, 'Didn't you know what you were doing?' and he said, 'I knew from the space station's perspective I shouldn't do it, but from my point of view, dammit, I *had* to!' Later, at a high-level NASA briefing, we were arguing that AUSM prevented this excessive demand of resources by guys who don't really need them. The person we were briefing said, 'We don't do that at NASA.' And this other guy stood up and said, '*I* did it.' There was no other way we could have proven it."

The next step will be to try AUSM on a real shuttle flight. There are still a couple of political hurdles to clear, but Ledyard is optimistic that it will fly one day. Meanwhile, AUSM's back in the lab for stress testing—seeing how well it holds up under various conditions.

The space station's clientele will probably be 90 percent scientific and technical, but AUSM would still be a boon to mission planners. Competition for resources favors payloads that use them most efficiently. And improved payload design could dramatically boost the space station's overall efficiency. AUSM can't evaluate purely scientific payloads like the Hubble Space Telescope now, but Ledyard has some ideas on how it could be done. As for the broader issue of the ratio of military to scientific to commercial use, he says that is a public-policy question. The allocation mechanism shouldn't decide policy or interfere with it, but should instead reflect Congress's, and ultimately the public's, will.

"AUSM would require a change in organizational culture," says Ledyard. "NASA sees allocation as its job, and pricing as a necessary nuisance imposed by Congress. NASA feels it would be nice if they somehow collected money, maybe, but it really has nothing to do with them. We feel we can significantly improve the allocation process while still raising some money for the government. What we're really looking at here is how you run government, good and bad ways to manage programs. Using economic data, generated under controlled conditions, in a policy debate is new for economists, but the opportunities are unlimited. And Caltech is remarkably well-equipped to worry about this kind of issue, because of our strength in integrating political science, economics, and experimental work."□—*DS*