Research in Progress

Cruisin' for Credit

Most drivers hope to avoid traffic jams. Not juniors Peter Hughes and Brad Solberg. They cruised the freeways this summer looking for trouble. Their goal: to develop a better computer program to predict traffic jams.

Hughes and Solberg's SURF (Summer Undergraduate Research Fellowship) project compared freeway observations with computer simulations. They based their simulations on existing computer models, in which traffic tends to move in waves. These waves are normally invisible to the average driver, but can be seen as the sequential flashing on of brake lights as traffic approaches a slow driver ahead, for example. The students' observations helped document the wave model, and enabled them to contribute some details of their own. For example, drivers have differing reaction times and preferred driving speeds, just as cars have different maximum accelerations and wind resistances.

Hughes and Solberg's model recreates these variations in detail. The model encompasses eight vehicle types, including four kinds of passenger cars, assorted trucks, buses, and 18-wheelers. The model uses 18 variables to describe a car-driver combination. This gives sufficient detail to populate a model freeway with the hotshot lane-changer in the brand-new, cherry-red Porsche Targa, the elderly driver in the '67 Rambler who drives at 40 mph in the center lane, and everyone in between. Additional variables describe the number of lanes, visibility, and related freeway characteristics. Some of the variables, such as AMAX (maximum acceleration) and BRTIME (braking reaction time) were borrowed from previously published models. Others, such as SPAZ (acceleration uncertainty) and BOC (braking overcorrection factor) came from their own observations of freeway flow.

Peter Hughes (left) and Brad Solberg videotaped freeway traffic from strategic hilltops.

Before sending any electronic cars onto their silicon freeway, the two spent many hours on the real thing, watching the behavior of real-life motorists. These observations initially helped determine what variables governed a driver's actions. Later, when the model was being tested, the antics of a simulated driver were compared to these same observations to determine the model's accuracy.

The pair set up several observation posts overlooking the freeways. The posts overlooked points where traffic jams could be expected to occur, such as freeway interchanges and long uphill stretches, as well as areas where traffic flows freely. They videotaped traffic for hours on end. "We learned a lot just by watching the traffic," Hughes said. "We could see waves [of higher density traffic] combining to form traffic jams. The San Diego Freeway through the Sepulveda Pass was especially good to watch, because it's a long hill with few exits."

After several days videotaping traffic, the two students were almost ready to take to the roads themselves. But before any meaningful speed data could be collected, they had to be sure their speedometers were accurate.

They calibrated their speedometers by driving a measured stretch of road near campus at various speeds and comparing their times to calculated values. "I have a Datsun B-210, and Brad has a '72 AMC Matador. It turned out that my speedometer was off by a constant 10 mph, while Brad's was off by a linear factor," Hughes said. "We wound up using Brad's car mostly. At least it settled the argument about who has the best car."

Then, armed with a digital stopwatch and a carefully calibrated speedometer, they were ready to roll. They patrolled the Ventura, San Diego, and San Gabriel Freeways. They recorded their speed and mileage at ten-second intervals (five seconds during traffic jams). The speed/time data would be converted to speed/distance, acceleration/time, and acceleration/distance plots on the computer.

They drove the same stretches of freeway over and over again at different times of day. The accumulated plots from off-peak hours enabled them to determine the effects of hills, curves, interchanges, and other fixed features on traffic flow. Traffic jams were not hard to find, however, as these same stretches of freeway clog regularly every rush hour. A portion of the Ventura Freeway, in fact, has the dubious distinction of being the busiest freeway in the nation, according to the United States Department of Transportation.

Between time and speed recordings, they observed their fellow motorists' driving habits. Hughes summarized their observations: "There are a couple basic types [of drivers]. One is the guy who picks a lane and just sits in it and will drive at whatever speed he feels like. It's pretty variable how well people control their speed. Sometimes people will be driving along and will go down a hill and won't watch their speed. Other people are very careful. Then there are the types who drive just as fast as they can, and weave from lane to lane when other people get in their way. Then there is the driver who'll drive along, usually in the fast lane, just as fast as possible, and when he comes up behind someone he'll tailgate and honk his horn and flash his lights till the guy moves out of his way. We started seeing a lot less of that sort of behavior in the second half of the summer after the freeway shootings started. . . . That changed everybody's driving strategy. Fewer people were driving in the left lane, people were letting people in a little more often, and we were seeing people looking around a lot more."

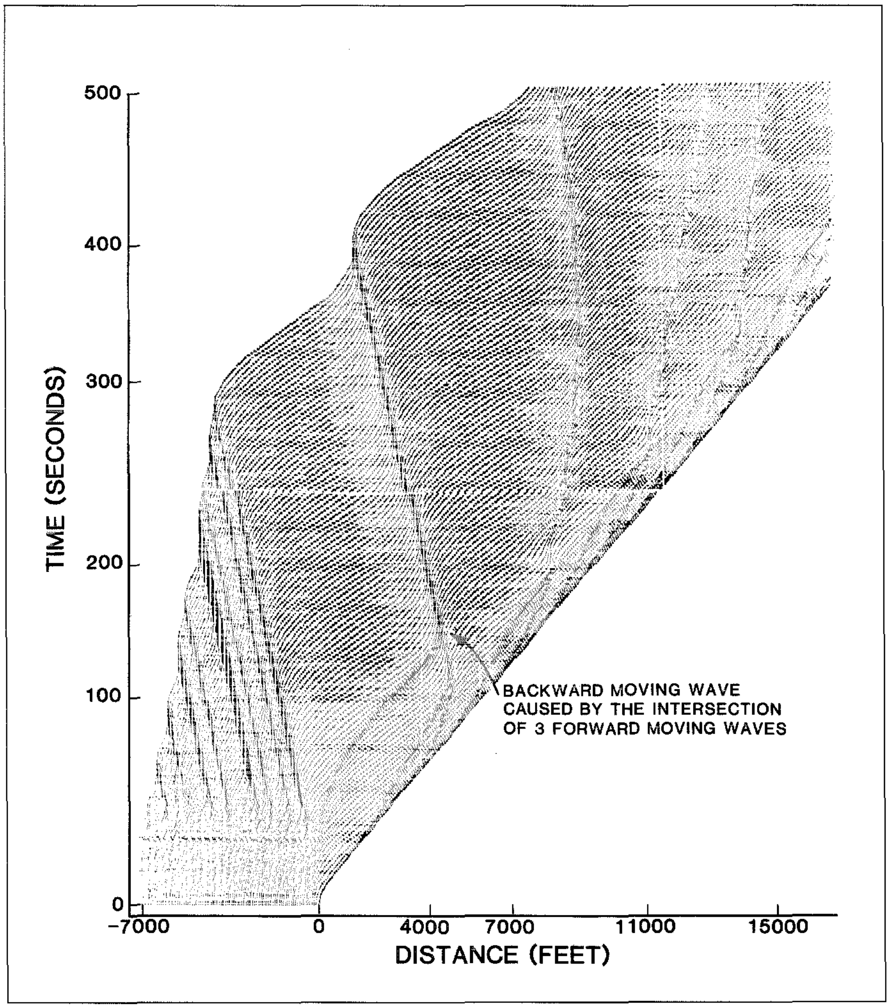

Time-distance plot for all the cars in a 100-car simulation. The simulation begins with all the cars stopped bumper to bumper. Each car is represented by a line. The plot shows how small waves of greater-than-average density combine to form very-high-density waves — traffic jams. The plot also shows the low-density areas following the jams.

The pair developed a set of variables to describe the driving strategies they observed. The variables quantify subjective responses such as driver frustration. Now the real work could begin: integrating the new variables into existing programs to produce realistic behavior.

They divided the model into two sections, a "car-following" routine and a "lane-changing" routine. "Previous models had also combined these problems," Solberg noted. "But ours added more individuality to both maneuvers."

The car-following routine simulates a driver's efforts to maintain a preferred speed. Hughes explains, "People will brake or accelerate to where they will avoid hitting the car in front of them but yet will approach their preferred speed. Most other car-following models we saw had a breakdown into two separate situations: when you're not constrained by another car [in front of you] and when you are constrained. We fit them together, saying that a person will see more and more effect as they approach the car ahead of them. People react a lot more to large speed differences [relative to the car ahead] than they do to small ones. It's a more-than-linear relationship, so we used a squared relationship." They finally derived an equation which included the square of the velocity difference between the two cars, the distance between the cars, the preferred following distance, and the "braking overcorrection factor."

"The other important change we made from what researchers had done in the past was to put in this braking overcorrection factor. It makes the equations asymmetric. When people are braking, they'll brake harder than they need to, whereas when they're accelerating they will accelerate just as much as they need to," Hughes said.

The lane-changing routine presented a special problem. "Since lane changing is an all-or-nothing response," Hughes said, "we had to find a way to make the probability of this response proportional to the stimulus." They created a series of four steps to determine when a car changes lanes. In the first step, the computer determines if traffic is flowing freely. If it is, there is no need to change lanes. If traffic flow is constrained, the computer calculates the desire to

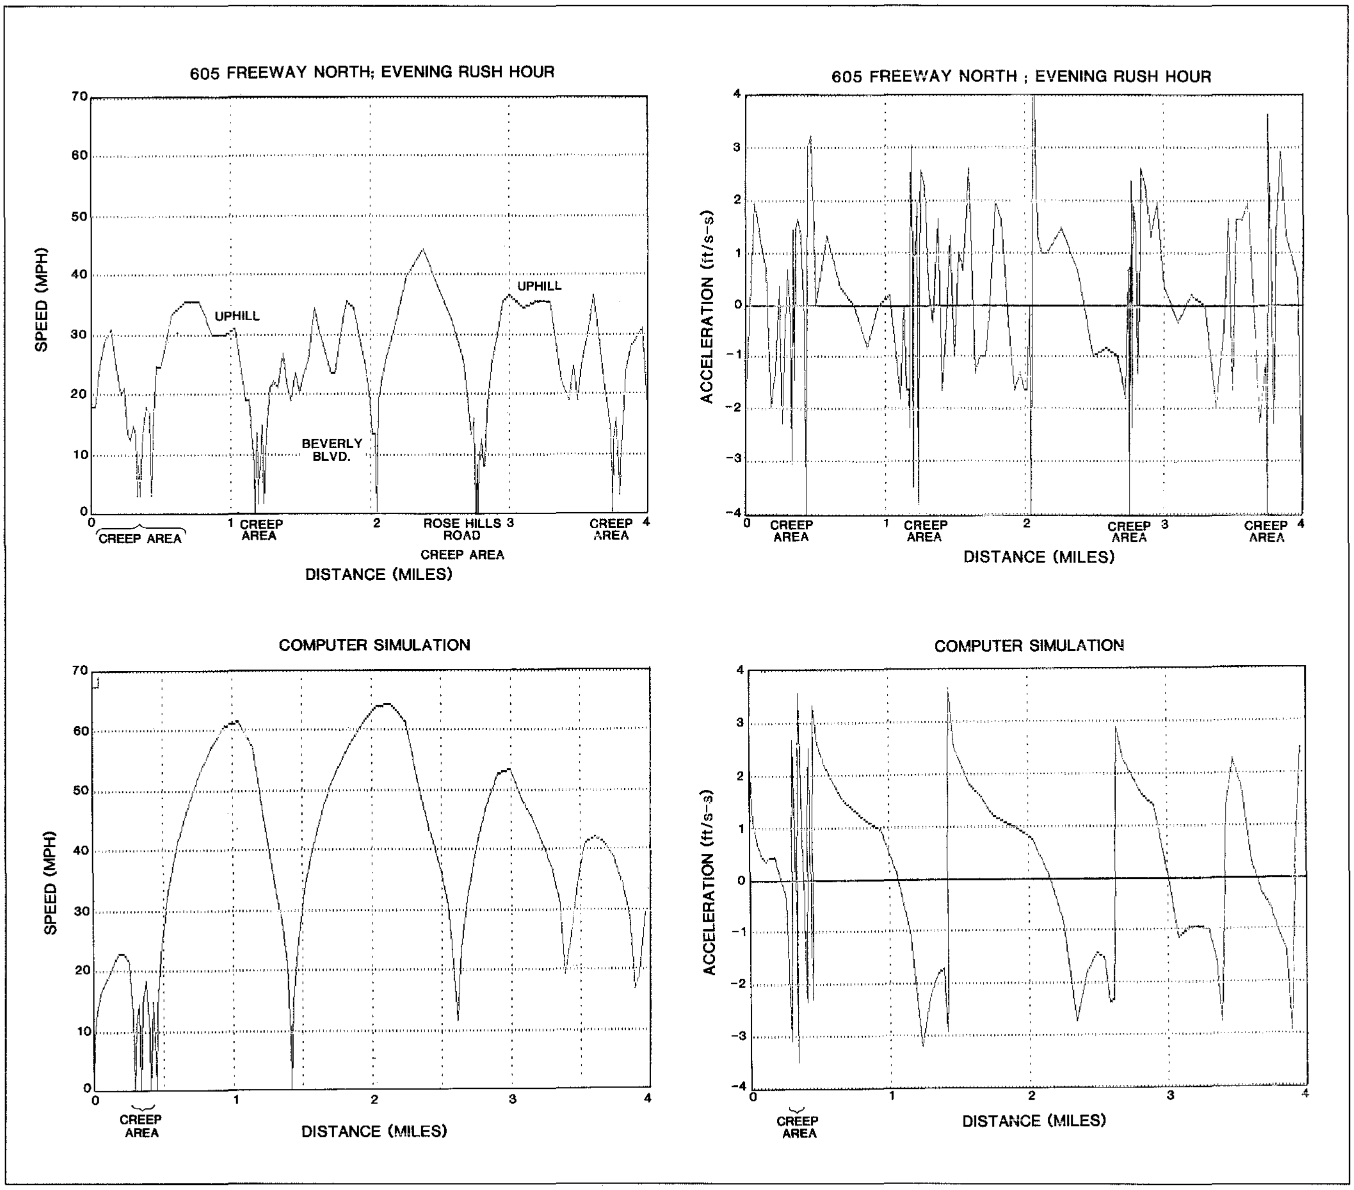

Speed and acceleration plots for a real car compared to the simulation. Both sets of plots show traffic slowing to a creep at intervals of approximately 0.7 miles. The acceleration plots show the braking overcorrection factor as a small hump following each large deceleration. The real life plot shows the effects of uphills and merging traffic.

change lanes as a function of the driver's "frustration factor" and elapsed time. After about 20 seconds, the driver is fed up and is ready to change lanes. Now the computer looks at the adjoining lanes, and if there is already a car there, or if changing lanes would cause a car in the new lane to brake unsafely, no lane change occurs. If the way is clear, the computer compares accelerations in the new lane and the current lane. If the current lane has the higher acceleration, no lane change occurs. If the new lane has a sufficiently higher acceleration, then the car changes lanes.

The model includes random variations, even for similar drivers in similar cars. The computer chooses the exact value of any variable based on a bell-shaped distribution around the mean value assigned to that class of car.

As the pair worked to refine their model, they compared computer generated plots to their field plots. "We set up a computer model and determined whether the waves formed were actually similar in shape to the waves we saw on the real freeway. Then we'd tweak the coefficients . . . We had a lot of people plowing into each other in our computer models when we were debugging them. In fact, that was our major way of figuring out that something had gone wrong. We had a collision flag in our program that went, 'SMASH CRUMPLE SCREAM MOAN' on the screen and the program would terminate. We figured that in normal freeway operation, even in a traffic jam, accidents probably won't happen if people are attentive," Hughes said. He added that they did not get involved in any accidents during their study, but that a few weeks later he was rear-ended by an inattentive driver on a surface street near campus. No injuries and little damage resulted.

The fruit of their labor is a computer simulation in which cars nearing slow traffic brake and change lanes realistically. The computer tracks all the cars on the freeway section, updating each car's position and velocity every 0.2 seconds. The data can be printed out as a trip chart for any given car; or as a snapshot of the entire freeway section at a given time; or at any point on the freeway for the duration of the test.

The simulation produces individual vehicle plots resembling field plots. A real car's velocity or acceleration, plotted against time or distance, gives a series of jagged, but reasonably regular, undulations. The peaks represent free-flowing motion between jams, while the valleys represent the jams. The braking overcorrection factor gives rise to the little hump seen at the end of each large peak on the acceleration plot. The peaks have a wavelength of approximately 0.7 miles.

When all the cars on the freeway are plotted on the same graph, a traveling wave pattern emerges. The plot shows a series of small waves of greater-than-average density. These ripples can travel forward, that is, along the direction of traffic flow, or backward. When two or more ripples meet, they merge to create a large, backward-moving wave of much higher density. A low-density zone follows, duplicating the area beyond the slowdown where cars burst free into empty lanes. This wave replicates those traffic jams, which, once a motorist has fought through them, have no visible cause. The model created these jams every 0.7 mile or so, duplicating Hughes and Solberg's freeway observations. The ripples also pile up at natural choke points, such as uphill grades, merging traffic areas, and stalled cars.

The computer plots tend to be simpler and more regular than field plots. Hughes hopes to make the model more realistic this year by expanding a driver's lane-changing options to include forcing his way into the adjoining lane; moving to cut off a car entering his lane; and losing his nerve partway into a lane change, aborting it. When these modifications are added to an expanded system, the simulation will enable transportation planners to simulate traffic conditions on their own freeways.□ — *DS*

The PASADENA Effect

Michael Pravica, Daniel Weitekamp, and Russell Bowers (l to r) prepare parahydrogen.

THE PASADENA EFFECT (Parahydrogen and Synthesis Allow Dramatically Enhanced Nuclear Alignment) produces greatly enhanced NMR signals that will allow scientists to determine the structure of reaction intermediates previously invisible to them. Predicted by Assistant Professor of Chemistry Daniel P. Weitekamp in June 1986, it has since been verified experimentally by C. Russell Bowers, a graduate student of his.

The discovery has many uses, as Weitekamp explains. "Looking at the structure and kinetics of catalytic sites, both in solution and at solid surfaces, will be the hot things. Many catalytic reactions occur on solid surfaces. When you want to design a catalyst, you have to figure out how it works first. Then you know how to design it better. Also, we can transfer the effect to nuclei that normally have very weak signals, such as ^{13}C and ^{103}Rh."

An NMR (Nuclear Magnetic Resonance) signal is generated when a nucleus changes its spin state. All other things being equal, the probability of making a transition from one allowed spin state to another is directly related to the relative populations of the two states. Nuclei, like Daniel Boone, want elbow room. The less crowded the destination, the more nuclei will make the transition and the greater the resulting signal. In the PASADENA effect, a sample of hydrogen gas is forced to choose a particular nuclear spin state. When that sample reacts with an asymmetric carbon-carbon double bond, the product molecule preserves that spin state long enough to be visible to NMR. The nuclei's headlong flight from the overpopulated spin state generates an NMR signal far out of proportion to the molecule's concentration.

According to Weitekamp, "The PASADENA effect signals are up to several thousand times bigger than the signals you could have gotten just by taking the product molecule and putting it in a magnetic field. If you had this molecule in a magnetic field, you'd have some population in each of four energy states, and the fractional differences between the amount in each of these four states would be something on the order of 10^{-5}. If you had 100,000 molecules in the highest state you'd have 100,002 molecules in the lowest state. With the PASA-

DENA effect we have no molecules in the highest and lowest states, and all the population just in the two middle energy levels until we turn on the radio waves. That causes a very intense transition. That makes for very large signal enhancements."

The four states Weitekamp describes are shown on the energy level diagram below. The hydrogen's energy (nuclear spin) states before the reaction determine the product's energy states. Quantum mechanical considerations divide the hydrogen molecule population into four energy states. Here's why.

Hydrogen, like all nuclei with an odd mass number or an odd number of protons, has a nuclear spin. By the laws of quantum mechanics, hydrogen's spin is limited to one of two possibilities: +½ (also known as the "up" or α state), and −½ (the "down" or β state). Since hydrogen gas (H_2) is a diatomic molecule, the spin states of both protons must be considered. The lowest energy state, where the rotational quantum number equals zero, is called parahydrogen; it is described in wavefunction terms as αβ − βα. (Remember, α is up and β is down. Which atom is which doesn't really matter, since the molecule is symmetric and floating freely in space.) There are three other states accessible to room-temperature hydrogen gas. All have one quantum of rotational energy and are, in order of increasing energy, αα, αβ + βα, and ββ. These three states together are called orthohydrogen. Room-temperature hydrogen interconverts among all four states, and the concentration of each is approximately equal.

If the gas is cooled, the proportion of parahydrogen increases. At absolute zero, the hydrogen will be 100 percent para. "It's been known for 50 years that you could isolate this pure nuclear spin state," Weitekamp observed. "But nobody's ever done anything with it. This nuclear spin state is seemingly very dull from the viewpoint of magnetic resonance because it has no magnetic moment. One spin points up; the other points down; so it has no observable magnetization."

It's the magnetization that makes hydrogen nuclei (protons) visible by NMR. The spinning, positively charged protons act like tiny magnets. When they find themselves in a strong, externally-induced magnetic field, they tend to align themselves with the field. Up spins are parallel to the field, down spins antiparallel. The parallel alignment is the lower, and therefore slightly more populated, energy state.

Protons can switch alignments (resonate) to a higher energy state by absorbing a quantum of the correct energy. A proton in the higher state could emit the same quantum and flip back to the low energy state. The net effect, absorption or emission, depends on which state is less densely populated. The required energy depends on the strength of the magnetic field. For the strongest fields practical in the laboratory, these quanta fall in the radio frequency (rf) range. The NMR technique places a sample in a strong magnetic field and bombards it with rf radiation. When a proton changes alignment, it generates a signal in a detector.

The NMR spectrum is a plot of absorption peaks versus field strength. Protons in different chemical environments resonate at slightly different field strengths, because other atoms in the vicinity can "shield" the proton from the magnetic field's full effect. The field interacts with nearby electrons, creating small secondary magnetic fields oriented opposite to the primary field. The shielded proton feels a lesser field and so resonates at a higher applied field. An interaction between nuclei, the J-coupling, can split a peak into a multiplet, a group of several closely spaced peaks.

But if parahydrogen has no net spin, it won't be affected by the magnetic field, so what's all the fuss about? "We use it as a chemical reagent and break the symmetry that prevented it from having any interesting spectroscopy," Weitekamp said. Molecular hydrogen reacts with carbon-carbon double bonds, forming a product with a carbon-carbon single bond and a hydrogen atom on each carbon. If the carbons were bound to different substituents to start with, the two hydrogen atoms would wind up in different magnetic environments. Suddenly, the hydrogens would not be equivalent any more. Now it matters which atom was up and which was down, because the two up-down combinations will be at slightly different energy states. The spins no longer cancel, and the protons suddenly become susceptible to the magnetic field. Half of the population winds up in each of the two available orientations, as shown on the following page, ready to make those intense transitions.

Once the theory had been worked out, the effect was simple to demonstrate.

Bowers prepared para- enriched hydrogen gas by cooling it. (At room temperature, all four spin states have almost equivalent populations.) As the temperature drops, the proportion of the lower-energy para- state increases. At liquid nitrogen temperatures (around 70K), half the molecules are in this one spin state. Unfortunately, the interconversion process is very slow. It takes several days for the ortho- and para- states to equilibrate

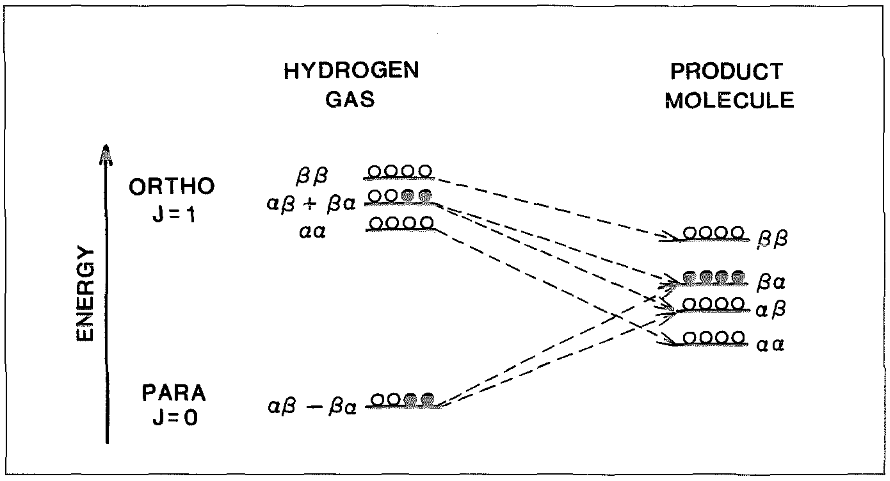

Normal hydrogen is equally distributed among four energy states (left). Hydrogen adding across an asymmetric double bond transfers its energy states to the product as shown by the dashed arrows. The black circles show how half the population in each of the αβ − βα *and* αβ + βα *states combine to form two new states, depending on which hydrogen is* α.

without catalytic help.

Bowers passed ordinary hydrogen through a coiled tube immersed in a liquid nitrogen bath. The tube contained a nickel catalyst to accelerate the interconversion process. In the meantime, a probe containing a premixed solution of substrate and catalyst was inserted into the room-temperature heart of the NMR's liquid helium-cooled superconducting magnet. When everything was ready, Bowers pushed a button to bubble a single pulse of parahydrogen-enriched gas through the probe.

The experiment proceeded in a carefully timed sequence. The hydrogen pulse lasted for 1.1 seconds. A delay of 0.8 seconds followed, allowing bubbles, which would distort the NMR signal, time to escape. Next, a 6.0 μsec pulse from the rf transmitter stimulated spin-state transitions, which were recorded for the next 1.25 seconds. The instrument collected 8,000 digital data sets for Fourier frequency analysis in that small interval. This tiny length of time was sufficient to generate a product molecule peak as big as one that would take a day to acquire under ordinary conditions.

Weitekamp's group has observed the effect in several substrates having carbon-carbon double bonds. In the spectrum shown here, acrylonitrile (CH_2=CHCN) has been converted to propionitrile (CH_3CH_2CN). The hydrogens must be in magnetically inequivalent positions in the product molecule for the effect to occur. The cyano group provides the shielding difference needed to break the magnetic symmetry. Wilkinson's catalyst — tris(triphenylphosphine)rhodium(I) chloride ($Rh(PPh_3)_3Cl$) — was used to bind dissolved hydrogen molecules and add them across the double bond.

The experiments confirmed Weitekamp's predictions. The tiny quantity of reaction product, too small to be seen by ordinary NMR, identified itself with a readily detectable signal. The signal disappeared after a few minutes, as the effect wore off and the molecule slipped back into a normal population distribution.

Another set of experiments used the catalyst but no substrate. The two hydrogen atoms occupy asymmetric positions on the catalyst, so Weitekamp hoped to obtain an NMR signal from the reaction intermediate. As predicted, the reaction intermediate gave a strong signal.

The NMR signals did not merely appear; they gave new information about the molecules' structure. Unlike conventional NMR, the PASADENA effect allows transitions to lower energy states (emissions) as well as absorptions. Emission peaks appear inverted, below the baseline. Note that only half of the peaks in each multiplet are inverted. Which half are inverted depends upon the sign of the "coupling constant," J, which governs peak separation within the multiplet. The coupling constant measures how strongly nearby nuclei interact with each other. In the product spectrum, the upfield half of the CH_3 and CH_2 multiplets were inverted, showing J had a positive sign; this agreed with theory and with less direct experimental evidence. In the spectrum of the catalytic intermediate, however, the downfield half of the multiplet was upside down, indicating J had a negative sign. Quantum theory had not predicted J's sign in this case, so the discovery came as quite a surprise.

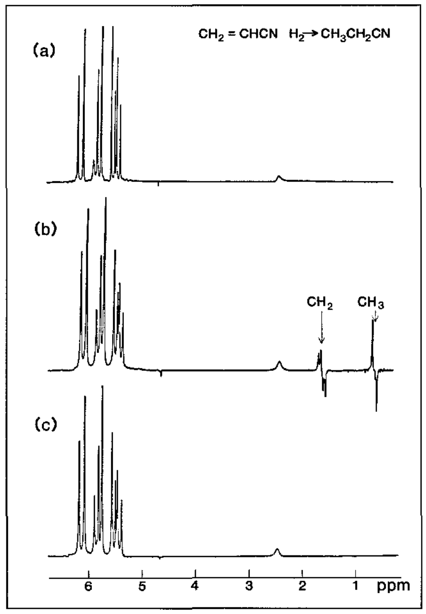

NMR spectra of catalytic conversion of acrylonitrile to propionitrile. Peaks clustered around 6 ppm are due to acrylonitrile. (a) Before adding parahydrogen. (b) The PASADENA effect makes trace quantities of propionitrile visible. The middle peak of the CH_3 triplet has been spilt into two equal peaks of opposite sign that cancel each other. (c) 200 seconds after adding parahydrogen. The PASADENA effect has disappeared.

Having both positive and negative lines within a multiplet can lead to problems. If the line widths become greater than their splitting distance, they overlap and cancel, reducing the overall signal. Michael Pravica, a senior working under the SURF (Summer Undergraduate Research Fellowship) program this summer, has demonstrated a variation wherein all lines in a multiplet add together constructively. The researchers have christened this variant the ALTADENA effect. It only remains to figure out what the acronym stands for.

Weitekamp summarized his work thus: "It's an interesting project because it combines quantum mechanics, where everybody learns and forgets about ortho- and parahydrogen, with real chemistry — catalysis."□ — *DS*

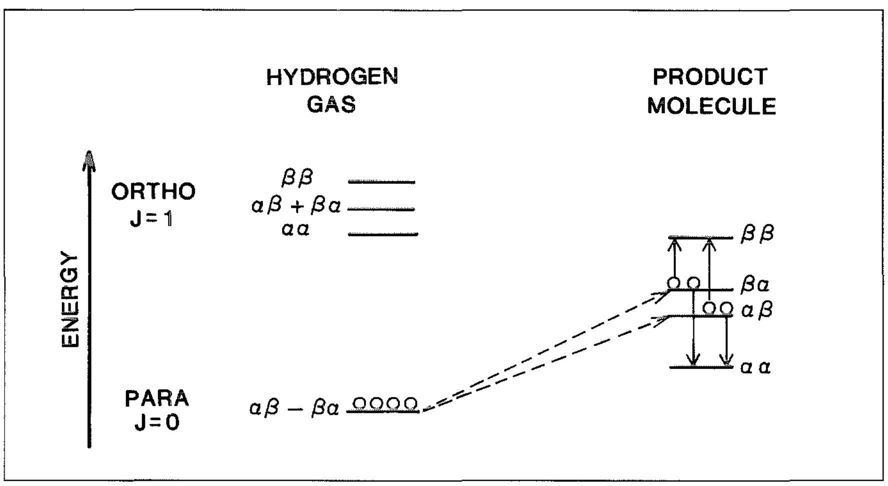

The dashed arrows show how pure parahydrogen adding across an asymmetric double bond can only reach two of the product's four available energy states. The product can make transitions to the other energy states as shown by the solid arrows. Up arrows are absorptions; down arrows, emissions.