Mix Well, Then Apply: Math Meeting in D.C.

Mathematicians are notorious coffee drinkers—so it's no wonder that the second International Conference on Industrial and Applied Mathematics (ICIAM), which met in Washington, D.C., from 8 to 12 July, budgeted $40,000 for coffee. More than 2000 applied mathematicians attended the 5-day java-fest. Between coffee breaks, they talked about flexible manufacturing, biorestoration, the curse of dimensionality, and a host of other topics. By the time they left, they had sat through more than 1000 talks—below, some highlights.

Goodbye Assembly Line

A key symbol of industrial production has long been the conveyor belt: identical objects moving smoothly along, while workers perform repetitive tasks. But in response to international competitive pressures, manufacturers are turning to systems that can produce variations on a product, rather than carbon copies. The new buzzword is flexibility. And its symbol is not the conveyor belt, but a roadmap of alternative production routes—a roadmap in which mathematics is playing a larger and larger role.

Such flexible manufacturing systems lead to mathematical problems industrial designers haven't faced before. Martin Grötschel, an applied mathematician at the Technische Universität and at the Konrad Zuse Zentrum in Berlin, described some of these problems—and some solutions—in an invited address to a large crowd at ICIAM.

"You want to build an optimal system that responds to all types of orders and can flexibly manufacture many different goods," Grötschel said in an interview with *Science*. "These systems are very complex, and that's what we're talking about: how to really do this in an efficient manner, how to build up a system, and how to run it."

Take scheduling. With identical goods on an assembly line, there is no scheduling problem: You just figure out which nut has to be tightened before what screw, and then set up the line accordingly. But when differing items—say three slightly different types of 35 millimeter cameras—are being produced, scheduling becomes an issue. Camera X may require work on machines A, B, C, and D; camera Y may need work on machines A, B, E, and F; and camera Z may need to stop at machines C, D, E, and F. How do you route things so that there isn't a big traffic jam of cameras at one machine?

Another example—a real-life one that Grötschel and colleagues tackled for a computer company in Germany—is the drilling of holes on printed circuit boards. The drilling is done by a laser beam, and the goal is to waste as little time as possible in moving the laser from one location to another. Grötschel's group was confronted by a version of the "traveling salesman problem," a classic math puzzle in which a salesman has to visit a number of cities and wants to do so with a minimum of time on the road. The puzzle is easy to solve, even by hand, when the numbers are fairly small. Computers can cope with somewhat larger versions. But when the numbers get just a bit bigger, the array of possible solutions gets so large that it quickly outstrips the capacity of any computer.

And that's where hardcore math comes into play. By analyzing the combinatorial problems that arise, applied mathematicians can often find solutions that are close to optimal, or at least an improvement over previous approaches. Grötschel's group developed algorithms that, while not exact, shave nearly 50% off the travel time of the designers' original solution.

Flexible systems are attractive, but they pose problems if they become too elaborate. "The drive to more and more complex systems may not lead to any improvement," Grotschel says, "because there's nobody who can write a program that can optimize or control the whole system—it's just too complicated."

Curse Foiled—Again

Last year, Henryk Wozniakowski, a computer scientist at Columbia University and the University of Warsaw, broke the "curse of dimensionality" for multivariate integration, a computational procedure with applications ranging from theoretical physics to financial forecasting (*Science*, 11 January, p. 165). Speaking to colleagues at a session on continuous complexity theory at ICIAM, WOzniakowski reported that he has now broken the same curse for a related problem—one that has even broader applicability.

The curse of dimensionality creates monsters—exponential monsters, that is. Certain computations involving multivariate functions seem doomed to grow into scary exponential Frankensteins as the number of variables increases. For integration, where the task is to approximate the $(d+1)$-dimensional "area" under a continuous, d-dimensional "curve" by sampling the function at a finite set of points, the curse says that the number of sample points required to guarantee an average error less than e is proportional to $1/e^d$. That's a terrifyingly huge number, even when only 4-place accuracy is asked for a function of a modest 10 variables.

Wozniakowski managed to break this curse by coming up with the mathematical equivalent of a silver dagger through its heart: an algorithm demanding a number of sample points proportional to only $1/e$. That stacks up well, at least for small e's, against a randomized approach to integration known as the Monte Carlo method, which has been in use since the 1940s.

Now Wozniakowski has produced a new and improved silver dagger that applies to the "approximation problem," whose goal is to approximate the entire graph of a multivariate function by sampling it at a finite set of points. Such approximation is a key operation in many scientific and engineering applications, including seismic exploration, medical imaging, and computer vision—all of which attempt to build up a complete picture from a limited number of measurements. Wozniakowski's algorithm turns a previously intractable computation into something that is at least theoretically within reach.

The precise problem is to sample the values of a continuous function on the d-dimensional "unit cube" (meaning there are d variables, each ranging in value between 0 and 1) in such a way that the function's value at any other point of the cube can be approximated by combining values of the sampled points—much as a curve can be approximated by a series of straight lines along its length.

The reason for looking at the "average" rather than "worst-case" error is that computer scientists have long known that it's impossible to guarantee a given level of accuracy uniformly for all continuous functions. So instead they consider what happens to an "average" function, using a technical definition of "average" that does for continuous functions what the familiar bell-shaped curve does for numbers.

The straightforward approach—sampling the function at equally spaced points on a d-dimensional grid—is no good (except in the one-dimensional case): It succumbs directly to the curse of dimensionality. Wozniakowski has now broken the curse by using something resembling a grid with many of its points removed. By doing so, he finds that the number of points required for an average function to be approximated with error less than e is roughly proportional to $1/e^2$.

That's not quite as good as the $1/e$ result for multivariate integration, but Wozniakowski has shown it's the best result possible. Moreover, notes collaborator Joe Traub, also of Columbia, Wozniakowski's algorithm has no theoretical competitor, since for approximation, Monte Carlo methods don't help.

Wozniakowski and Traub hope to break the curse of dimensionality for more problems, including problems in optimization and differential equations. Their ultimate goal is to define the group of problems that are computationally "tractable" in the average-case setting. At the same time, they would like to make Wozniakowski's theoretical breakthroughs more practical. Right now the results are only asymptotic, meaning that the proportionality to $1/e$ (for integration) and $1/e^2$ (for approximation) holds only for small values of e. "We want to get non-asymptotic results," says Wozniakowski. "They'll be much more important for applications."

Microbial Math

The contamination of groundwater aquifers by toxic chemicals is widely recognized as a major problem in the United States. An example: At the Hanford Site in south central Washington state, underground disposal of carbon tetrachloride over 18 years has left a 5-square-kilometer plume of contaminated groundwater that may cost $60 billion to clean up. And that's only one case: Toxic chemicals routinely leak from landfills and rusted storage tanks everywhere.

But in many cases, the gray cloud of toxic waste has an unanticipated microbial silver lining: organisms already present in the environment are capable of cleaning things up, or at least containing the damage. Indeed, in recent years the use of microbes as toxic scavengers—in situ biorestoration as this approach is called—has emerged as a promising treatment technology.

Yet there's a big if: The microorganisms must be triggered properly by the introduction of dissolved oxygen or other nutrients into the system, and providing just the right nutrient triggers for this weapon isn't simple. The process is physically and chemically complex, involving transport and interaction of substrate, dissolved oxygen, and microorganisms, as well as the movement of water within the aquifer. To help unravel some of these complexities, researchers at Rice University are turning to mathematical models. Although the models are in the early stages now, Mary Wheeler and her colleagues hope that ultimately they will enable environmental scientists to characterize contaminant migration, aid in making regulatory decisions, and predict and evaluate the performance of restoration projects.

Mary F. Wheeler

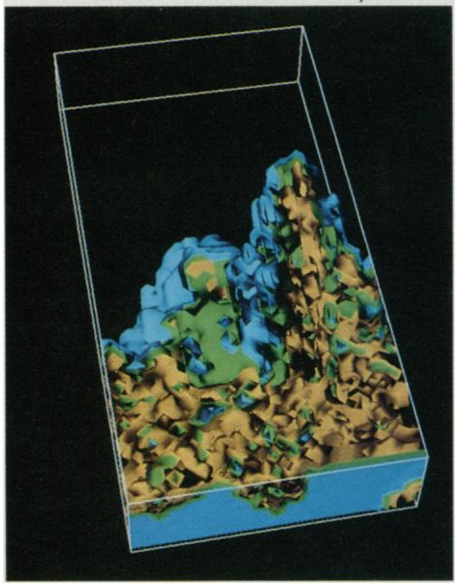

Odor eaters. *The distribution of microbial mass in a model of in situ biorestoration, as modeled by Mary Wheeler of Rice University and her colleagues. The lowest concentration is in blue, the highest in yellow.*

In an invited address at the Washington meeting, Wheeler described the groundwater model, which consists of a system of nonlinear partial differential equations. There's no hope of finding their exact solutions—only numerical approximations. One thing that makes the problem "very hard," Wheeler notes, is the fact that things in the model happen on different time scales: Chemical reactions are very fast, the flow of fluids is slower, and the model needs to consider even longer-term effects—that may extend over hundreds of years.

There are "lots of computational questions to be settled," Wheeler says, adding that two things may help her in that effort. New computer architectures for parallel computation should make it possible to create more realistic, three-dimensional models. At the same time, interactive graphics and other computer visualization tools should help investigators see more directly what the models are trying to show them.

The effort could be well worth it, Wheeler argued, because models have one great virtue: They could prevent those directing cleanup efforts from making expensive mistakes in real time. Field studies and experiments cost millions of dollars, Wheeler notes, "and if you make a mistake, it can be even more costly. Doing it on the computer is very cheap." ■ BARRY CIPRA

A Most Improbable Planet

If you were looking for a planet circling some distant star, about the last place you'd expect to find one would be in the harsh environs of an ultradense neutron star. But a group from the Nuffield Radio Astronomy Laboratories at Jodrell Bank, in England, believes it has stumbled upon a planet in just such an improbable place.

Their report in this week's *Nature* could be the first solid evidence of a planet outside our own solar system. Caution is in order, however, for there have been many putative glimpses of extrasolar planets in the past, and most have not stood up to further scrutiny.

"We weren't looking for anything like this," says Andrew Lyne, who was monitoring the radio signals from some 40 pulsars—rapidly spinning neutron stars—with his coworkers Matthew Bailes and Setnam Shemar when they found the key evidence. Pulsars, like all neutron stars, are thought to form when massive progenitor stars explode as supernovae. Even if the star had a solar system, Lyne says with some understatement, "people had assumed that any planets would have a hard time surviving."

Still, Lyne and his colleagues think that a planet might be the cause of puzzling oscillations they observed in the signal from one pulsar 30,000 light-years from Earth. Ordinarily, pulsars give off an uncannily regular, clocklike sequence of radio pulses. But in this case, he says, "we couldn't make sense of the arrival times we were seeing." Finally a pattern emerged: Over a 6-month cycle the interval between pulses—roughly a third of a second—got gradually shorter by 8 milliseconds, then longer again.

The explanation they favor: an unseen mass orbiting the neutron star every 6 months, pulling it to and fro by a distance of 8 light-milliseconds. The workers calculate that a planet about 10 times as massive as Earth, orbiting at about the same distance as Venus, would do the job.

If its existence is confirmed—and David Black, director of the Lunar and Planetary Science Institute in Houston, told *Science* the claim is "going to warrant confirming"—that still leaves the problem of explaining how it got there. Maybe, Black suggests, the pulsar formed in "a kinder, gentler supernova." Or maybe the explosion destroyed the star's original planets but left debris that condensed to form new ones. Either way, Black says, the astronomers may have found "a true celestial freak." ■ TIM APPENZELLER