MEETING BRIEFS

Math Attendees Find There's Life After Fermat Proof

The uncertain status of the recent proof of Fermat's Last Theorem (*Science*, 24 December 1993, p. 1967) was the hot topic in the hallways at the joint meetings of the American Mathematical Society and the Mathematical Association of America, held 12-15 January in Cincinnati. But a few other topics did manage to get discussed. Among them: How gambling pays off in computing the volume of high-dimensional shapes; how a mathematical description of water dripping down a window pane might lead to the design of digital pulses that could race through the optical fibers of the future; and how one can be fooled infinitely often by another theorem of Fermat's.

A Game of Darts and a Drunkard's Stagger

Computing the "volume" of objects with more than three dimensions might sound an arcane pursuit, but it's actually the first step in many real-world problems, from high-energy physics to multivariate statistics. Conceptually, the calculation is straightforward. But as a practical matter, the computation can be so labor-intensive that it exceeds the capacity of the fastest computers, even if one is only looking for an approximate answer. The reason is that the amount of work it takes to compute the size of an n-dimensional object generally grows exponentially with n. At the Cincinnatti meeting, however, László Lovász of Yale University and the Eötvös Loránd University in Budapest described a labor-saving strategy: gambling.

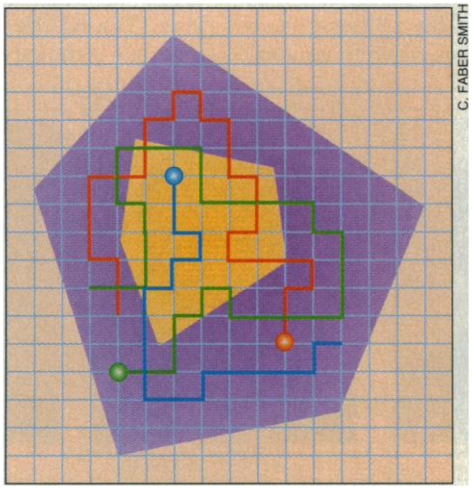

C. FABER SMITH

Drunkard's walks. Confined within a high-dimensional shape, they lead to random points in a short-cut for estimating a smaller shape's size.

The idea of gambling—trying lots of numbers at random—to speed up a calculation actually goes back to the 1940s, when the Polish-American mathematician Stanislaw Ulam introduced the "Monte Carlo method," a technique widely used in statistical physics. Applied to determining the size of an object, the idea resembles throwing darts: If the darts are thrown at random, the fraction that hits a particular region of the dart board—say a picture of your ex-boyfriend—is equal to the fraction of the board occupied by the region. Likewise, if an n-dimensional region sits inside an object of known size—say a "cube" with sides of length 1—then its volume can be approximated by picking random points inside the cube and counting the fraction that land within the region.

This method is now beginning to offer the promised speedup, but it's been a long time coming. The basic dart board method, in fact, offers no speedup at all, since high-dimensional targets are often so much smaller than the objects they sit in. For example, a unit "sphere" sitting inside the unit cube in n dimensions occupies less than $1/2^n$ of the cube's volume if n is greater than 12. The result is that for, say, a 20-dimensional sphere, the dart board—the cube—is so much larger than the target that the first billion or more darts are not likely to hit the sphere at all. It would take budget-deficit-like numbers of darts to obtain even a single decimal point's worth of accuracy.

But there are ways to make the approximation method work. A breakthrough came in 1989, when Martin Dyer at the University of Leeds in England and Alan Frieze and Ravi Kannan at Carnegie Mellon University found a random algorithm for approximating volume whose computational demands grew with the 27th power of the dimension, rather than exponentially. To ordinary folks, that won't sound like much of a shortcut, but the technique opened the way to further improvements. The basic idea is to place the n-dimensional shape in question within a nested set of n-dimensional shapes, each a little larger than the one inside it, working up to the unit cube. The dart board for figuring the volume of each shape is the next larger shape; by combining all the computations, you arrive at the volume of the first, smallest shape as a fraction of the unit cube.

Because of the nesting, each "target" occupies a substantial fraction of the board around it, and it turns out that the resulting speedup more than offsets the need to repeat the computation many times while working up through the series of dart boards. But this strategy raised a new problem: making sure all the random darts hit the right dart board during each successive computation. Dyer, Frieze, and Kannan's answer resembles a drunkard's stagger: taking a single dart that's already hit the board and moving it by small steps in random directions, making sure that it does not move off the board.

Though Dyer, Frieze, and Kannan's original random algorithm was still far too slow for practical application, a series of refinements in the last 4 years have boosted its efficiency. Most recently, Kannan, Lovász, and Miklós Simonovits of the Hungarian Academy of Science in Budapest have reduced the computational demands to the 5th power of the dimension, and, if a certain conjecture proves correct, to the 4th power.

The latest bounds bring the algorithm much closer to practicality. Increasingly powerful computers will do their part as well, but they can't take the place of efficient algorithms, says Lovasz, because bigger machines will only tempt users to tackle bigger problems. In fact, he says, "the faster computers become, the more need there is to use sophisticated methods."

A Compact Model Heads for the Data Highway

Bumper-to-bumper traffic normally proceeds at a crawl, if it manages to move at all. But in principle there's no reason cars can't cruise along inches apart at 70 miles per hour. Standard light waves on an optical fiber, however, must maintain a minimum distance between consecutive pulses, or else the information they carry will be garbled. But new mathematics holds out hope that future fibers could be designed to carry digital messages bumper-to-bumper at the speed of light. The benefit, of course, is that much more information could flow over the fibers of, say, Al Gore's national data highway.

The key to this mathematical scheme for crowding more traffic onto the data highway is generating optical signals that keep out of each other's way. Ordinary light pulses smear out with time; more durable signals called solitons have long, weak tails that can interfere with each other. But signals described by a new class of equations start off compact and stay that way, theorist Mac Hyman of Los Alamos National Laboratory said in an invited talk. As a result, these waves, dubbed "compactons," might open the way to packing information more tightly onto optical fibers.

Because of the optics of current fibers, most of the signals now traveling the information highway are described by traditional, linear wave equations. Such waves are subject to an effect called dispersion, which causes the signals to break apart, because higher-frequency components outrace their lower-frequency cousins. That effect stretches the signals, making it necessary to space them out along the fiber and periodically interrupt the pulses to restore them to their original form.

In the 1980s, researchers began experimenting with optical fibers on which the governing equations are nonlinear. These equtions have solutions that form "solitary waves," or solitons: highly localized waves that travel without suffering dispersion. But in the equations that have been considered up to now, each solitary wave has a pair of infinitely long "tails," one in front and one behind the peak of the wave. If a wave creeps up on another's tail, they interact in complicated ways that can garble the signal.

Now Hyman and Philip Rosenau at the Technion in Haifa, Israel, have developed a new class of tailless solitons. These "compact" waves, or compactons, behave much like cars: "They have absolutely no knowledge of each other until they touch," explains Hyman.

The key advance occurred to Rosenau and Hyman when they were thinking about how to describe the way water drops smear out as they drip down a window pane—a physical phenomenon in which dispersion plays a key role. But once they got the equations, Hyman says, they dropped the original problem and began exploring the mathematics. The result was several classes of wave equations in which adding an extra nonlinearity to the "dispersion term," which governs the stability of the wave's shape, changes traditional solitary waves with tails into compact waves.

That nonlinearity might prove to have practical value if researchers could find —or manufacture—materials that are consistent with the new equations. In the meantime, Rosenau and Hyman are experimenting with these waves mathematically by, among other things, colliding them deliberately. When solitons collide, the two waves neither fuse nor shatter; instead, they seem to pass straight through each other. The same is true of compactons, but they also leave behind an apparently endless string of tiny ripples.

The ripples "surprised the heck out of me," Hyman told *Science*. Understanding what causes them, and more generally how nonlinear dispersion affects the propagation of waves, is likely to keep theorists busy for years to come. Jokes Hyman: "I'm kind of throwing up my hands and hoping the graduate students dig into it."

The Little Theorem That Could—Doesn't

Fermat's Last Theorem gets all the press, but number theorists actually get more mileage out of another finding of the French mathematician, which they call Fermat's "little" theorem. Unlike Fermat's Last Theorem, which, in spite of the headlines, is mainly a mathematical curiosity, Fermat's little theorem is a workhorse. Among other things, it's a quick way to test whether a large number is prime or composite, a problem that often comes up in cryptography, for example. That workhorse, however, has just lost some of its power, with a finding by three mathematicians at the University of Georgia.

Number theorists have known for nearly 100 years that some composite numbers slip past Fermat's little theorem. There was a chance that the little theorem might retain much of its value for prime testing if the number of these impostors was limited: Anyone using the theorem as a prime test could simply check the result against the list of known impostors. No such luck, as Carl Pomerance, one of the three theorists, told his colleagues at the Cincinnati meeting. There are, in fact, infinitely many impostors.

Fermat's little theorem asserts that if *n* is a prime number, then it will be a factor of $a^n - a$, with *a* being any integer. For example, 5, being prime, divides $2^5 - 2 = 30$. On the other hand, 6 does not divide $2^6 - 2 = 62$, showing that 6 is composite. As it turns out, it usually doesn't take a lot of tries with different values of *a* to ferret out a composite; most of the time, if *n* is composite, it fails to divide $2^n - 2$—and if it passes that test, it's likely to fail the test with $3^n - 3$. But some composites, such as 561, 1105, and 1729, "pass" Fermat's little theorem for *all* numbers *a*.

The first to recognize these phonies was the American mathematician Robert Carmichael, around 1910; since then, they've been called "Carmichael numbers." Now it seems that Carmichael has more namesakes than he could have known. Using techniques from analytic number theory and some relatively recent results in abstract algebra, William "Red" Alford, Andrew Granville, and Pomerance built on a scheme for conjuring up Carmichael numbers put forward in 1956 by the Hungarian mathematician Paul Erdös. The trick, Erdös proposed, is to identify large sets of prime numbers that can be combined to form Carmichael numbers, then count up all the possibilities.

Erdös' argument suggested that such sets ought to be relatively easy to find. In practice it takes ingenuity, but in one computer run, Alford produced an example that gave rise to 2^{128} Carmichael numbers. However, to prove that there is an infinity of these ersatz primes, the theorists had to render certain parts of Erdös' argument rigorous. Their main result is that there are at least $x^{2/7}$ Carmichael numbers less than any (suitably large) number *x*. That's enough to prove there are infinitely many Carmichael numbers, but it probably underestimates their frequency. According to Pomerance, the exponent 2/7 should be replaceable by any power less than 1. The upshot: If Fermat's little theorem tells you a number is composite, you can rest assured it is, but if it tries to tell you the number is prime, take the answer with a grain of salt.

–Barry Cipra

A Do-It-Yourself Fermat Proof

Unlike his famous "last" theorem, Fermat's little theorem is easy to prove. One proof of the assertion, that any prime *n* will be a factor of $a^n - a$, with *a* being any integer, boils down to putting colored pebbles around a circle. For specificity, here's how the argument goes to show that 11 divides $3^{11} - 3$.

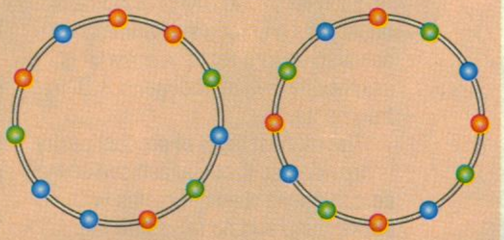

At each of 11 points around a circle, place a pebble of one of three colors (*left*). There are 3^{11} different ways to do this, of which 3 patterns use only one color each. The remaining $3^{11} - 3$ patterns, however, are not really all different: Each one is related to 10 others, which are just rotations of it. These patterns, therefore, clump together in groups of 11, which means that 11 must divide $3^{11} - 3$.

When the number of points is increased to 12, a composite number, some of the $3^{12} - 3$ patterns belong to groups of less than 12 (there are only three different rotations of the pattern at right, for example). Thus, 12 need not divide $3^{12} - 3$ (and it doesn't).

This argument—easily verified with a calculator—works the same way for any number *a* in place of 3 and any prime number *n* in place of 11: The prime number will pass the test. Unfortunately, not all composite numbers fail, and the number that don't has just been shown to be infinite (see main story).

–B.C.