Parlez-vous Wavelets?

Ingrid Daubechies. (Photo courtesy of AT&T Bell Laboratories.)

Mathematicians are like the French," the German poet Goethe once remarked. "They take whatever you tell them and translate it into their own language—and from then on it is something entirely different."

Goethe's observation is as true now as ever. But times may be changing. In the last ten years, mathematicians and researchers in diverse areas of science, engineering, and even art have discovered and begun to develop a theoretical language they can all understand. This new common language is sparking new collaborations. Many mathematicians are now crossing over into such applied areas as signal processing, medical imaging, and speech synthesis. At the same time, much deep but abstract-sounding mathematics is becoming accessible to researchers in fields from geophysics to electrical engineering.

The new language is wavelet theory. Those who speak it describe wavelets as powerful new tools for analyzing data. Wavelet theory serves as a kind of numerical zoom lens, able to focus tightly on interesting patches of data—but without losing sight of the mathematical forest while attending to the trees, twigs, buds, and grains of pollen.

"Never before in anything on which I've worked have I had contacts with people from so many different fields," says Ingrid Daubechies, a mathematician at AT&T Bell Laboratories and a leading authority on wavelet theory. Because there are so many aspects to the subject, "you have all these ideas brewing together—it's very fertile for everybody concerned," Daubechies adds. "It's a very nice laboratory for showing that applications can have interest for pure mathematics, and vice versa."

Mathematically, wavelets are an offshoot of the theory of Fourier analysis. Introduced by the French mathematician Joseph Fourier in his essay *Théorie analytique de la chaleur* (analytic theory of heat), published in 1822, Fourier analysis seeks—with great success—to understand complicated phenomena by breaking them into mathematically simple components. The fundamental idea is to take a function and express it as a sum of trigonometric sine and cosine waves of various frequencies and amplitudes. The familiar and well-understood trigonometric functions are easy to analyze. By combining information about a function's sine and cosine components, properties of the function itself are easily deduced—at least in principle.

Fourier analysis is among mathematics' most widely used theories. It is especially suited to analyzing periodic phenomena, periodicity being the most prominent property of sines and cosines. But even so, the theory has its limitations and its pitfalls. The main problem is that finding detailed information about a function requires looking at a huge number of its infinitely many Fourier components. For example, a transient "blip," obvious in a graph, is impossible to recognize from its effect on a single component. The reason, in essence, is that each sine and cosine wave undulates infinitely in both directions; thus a single wave can't help locate anything. Indeed, the sharper the blip, the more Fourier components are needed to describe it.

Wavelet theory takes a different approach. Instead of working with the infinitely undulating sine and cosine waves, wavelet analysis relies on translations and dilations of a suitably chosen "mother wavelet" that is concentrated in a finite interval. Almost any function can serve as the mother wavelet; this makes wavelet theory

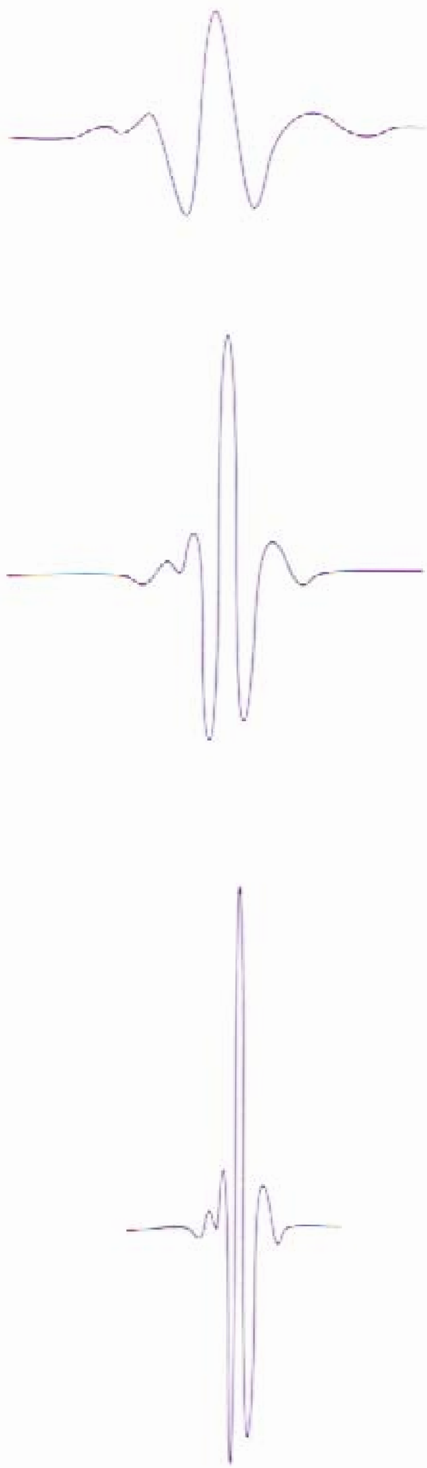

Figure 1. *A "mother" wavelet (top) and two "daughters." (Figure courtesy of Ingrid Daubechies.)*

more flexible than traditional Fourier analysis. "Daughter" wavelets are formed by translating, or shifting, the mother wavelet by unit steps and by contracting or expanding it by powers of two (see Figure 1). One then expresses other functions as combinations of wavelets, just as Fourier analysis represents functions by combining sines and cosines.

The fact that the mother wavelet is concentrated in a finite interval gives wavelet theory its zoom-in capability: An interesting blip in a function can be analyzed by looking only at those wavelets that overlap with it: finer details are resolved by looking at increasingly contracted copies of the mother wavelet in the vicinity of the blip.

Many of the ideas underlying wavelet theory have been around for decades, but the subject itself got off the ground only recently. The story starts in the early 1980s in France, when wavelets were introduced by geophysicist Jean Morlet and mathematical physicist Alexander Grossmann. In 1985, mathematician Yves Meyer constructed a family of wavelets with two highly desirable mathematical properties, called smoothness and orthogonality. (Interestingly, J. O. Stromberg at the University of Tromso in Norway had constructed such a family several years earlier, but the connection with the nascent theory of wavelets was not realized until after Meyer's work.)

The following year, Meyer and Stephane Mallat gave the subject a solid foundation with a theory of "multiresolution analysis." Then in 1987, Daubechies constructed a family of wavelets that, in addition to being smooth and orthogonal, were identically zero outside a finite interval. Daubechies's construction opened up the field. "Compactly supported" wavelets are now easy to come by, and are among the most commonly used in applications.

And applications are abundant. Wavelets are being tested for use in everything from digital image enhancement—making blurry pictures sharp—to new methods in numerical analysis (itself widely used in scientific computing). "They're a very versatile tool," says Daubechies. Not all the applications will pan out, but many will, and some already have. "There are some very nice success stories," Daubechies adds.

One such story may have far-reaching effects, especially for the next generation of criminals. The Federal Bureau of Investigation has adopted a wavelet-based standard for computerizing its fingerprint files. The FBI has around 200 million fingerprint cards on file, according to Peter Higgins, deputy assistant director of the Bureau's Criminal Justice Information Services division, and 30,000 to 40,000 identification requests pour in every day. At present, the FBI's fingerprint files consume about an acre of office space. The goal, says Higgins, is to digitize the files, store them electronically, and "put [them] in something that would fit in a 20 × 20-foot room."

It sounds easy: after all, entire encyclopedias now fit on a compact disk with room to spare. But that's words. Images are something else. At a resolution of 500 pixels per inch, a standard fingerprint card contains nearly 10 megabytes of data. Transmitting that much information over a modem—something the police would like to be able to do—takes hours at today's transmission rates. For a dozen cards, it's quicker to use Federal Express.

What's needed is some way to compress the data on a fingerprint card without distorting the picture. That's where wavelets come in. By treating the fingerprint image as a two-dimensional function, it's possible to represent it with a combination of wavelets. With a suitably chosen family of wavelets, only a relative handful are needed to represent a fingerprint, and the contribution of each wavelet can be rounded off, or "quantized," which reduces the amount of data that needs to be stored or transmitted.

The wavelet standard for fingerprints was developed by Tom Hopper at the FBI and Jonathan Bradley and Chris Brislawn at Los Alamos National Laboratory. The standard allows many kinds of wavelets to be used—in effect, each electronic fingerprint "card" will include formulas for its particular wavelets, as well as the wavelet representation of the fingerprint itself. So far, one family of wavelets has been approved for use. It compresses fingerprint data by a factor of approximately 20 to 1—reducing 10 megabytes to a much more manageable 500 kilobytes—yet gives images that pass the FBI's automated recognition tests. Indeed, the reconstructed fingerprints look almost exactly like the originals (see Figure 3).

Bradley and Brislawn have also applied wavelet techniques to another kind of data compression: managing the numerical geysers that gush out of supercomputers when running such things as global climate models. "High-performance computers are reaching the point where their ability to churn out data is surpassing our capacity for storing and analyzing it," says Brislawn. In the approach he and Bradley have developed, the computer decomposes the solution (for example, a color-coded map of global ocean temperatures) into wavelets; the user can then control the output by specifying how much detail—that is, how many of the wavelet components—he or she wants to see. One challenge is to figure out how much you can compress the output without sacrificing quantitative capabilities of a model, such as long-term statistical predictions of climatic conditions. Brislawn notes, "This looks like a tough question that we won't be able to answer until we get a better idea of what the models are capable of predicting."

Other researchers are studying the use of wavelets not as post-processing tools, as Bradley and Brislawn are doing, but directly in scientific computation itself. Gregory Beylkin at the University of Colorado has been studying applications

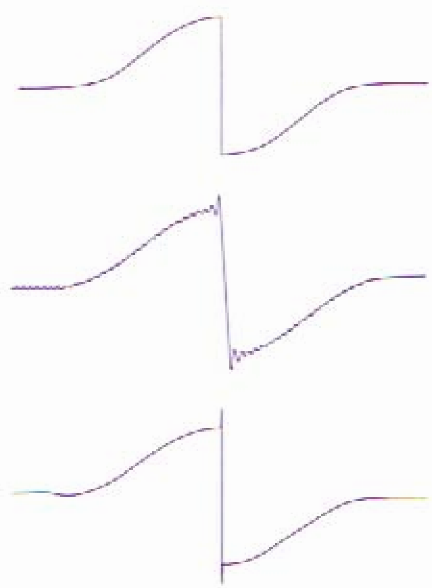

Figure 2. *A Fourier (middle) and wavelet (bottom) reconstruction of a function (top) with a sharp discontinuity. The Fourier reconstruction uses 65 nonzero coefficients, the wavelet reconstruction only 18. (In both cases, the discontinuity causes an overshoot, known as a Gibbs phenomenon, but it is much more localized in the wavelet reconstruction.)*

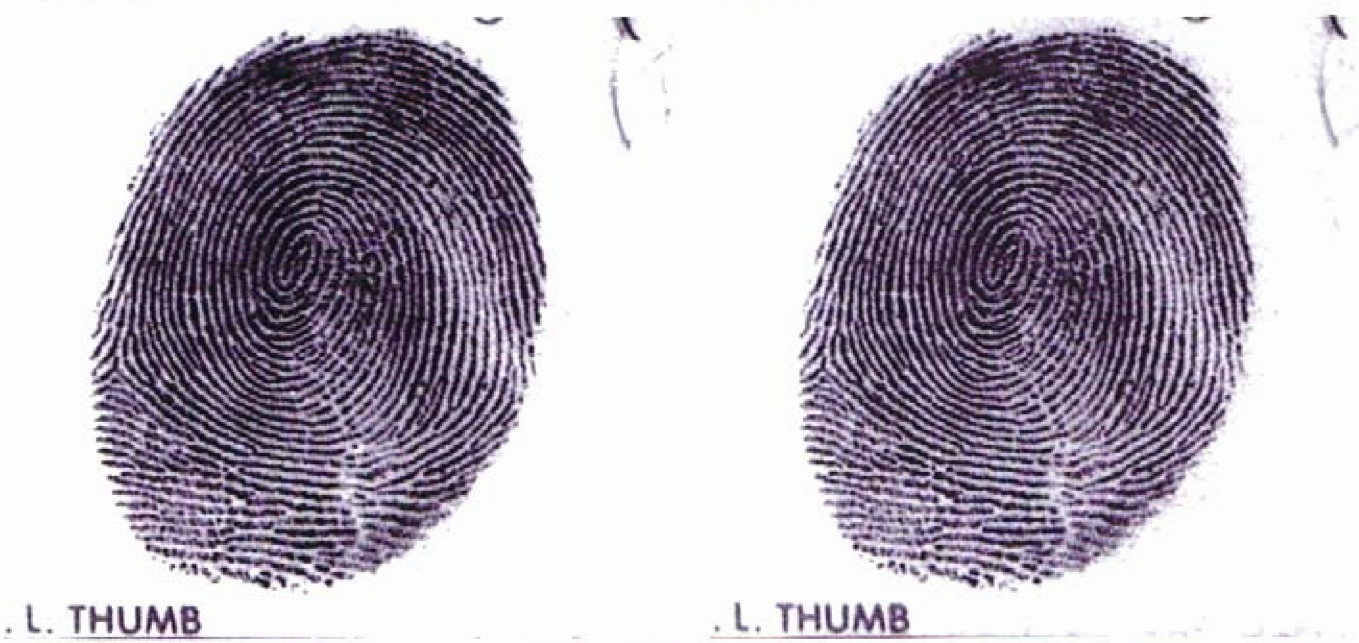

Figure 3. *(Left) Original 768 × 768 8-bit Gray-scale fingerprint image. (Right) Fingerprint image compressed 26.0:1. (Photos courtesy of Chris Brislawn, Los Alamos National Laboratory.)*

Wavelet theory is making some of the hard-won insights of mathematicians working in the abstract reaches of analysis accessible to researchers in many fields.

of wavelets in numerical analysis. Many problems—such as solving a system of partial differential equations that describes the flow of oil underground—boil down to working with huge matrices, or square arrays of numbers. Such matrices are easier to work with if many of their entries are zero. Beylkin has shown that wavelet analysis can reduce a wide class of matrices to the desired form.

Wavelets are especially suited to analyzing sound. Indeed, there's a strong resemblance between wavelets and musical notes. The mother wavelet can be likened to a particular note—say a quarter note at middle C—played at a particular time. Its translates represent the same quarter note at middle C played at other times, while its contractions and expansions are eighth- and half-note C's, played at higher and lower octaves. Ronald Coifman at Yale University and Victor Wickerhauser at Washington University in St. Louis have developed a technique they call adapted waveform analysis, in which a catalog of waveforms is automatically searched for the wavelets best suited to a particular problem. Among the applications is removing noise from recorded sound.

Coifman and his colleagues recently cleaned up an old piano recording of Johannes Brahms playing one of his own Hungarian Dances. Over the years, the recording had acquired several layers of noise. Brahms's live performance was recorded in 1889 on a wax cylinder, which later partially melted. The damaged cylinder was re-recorded on a 78 rpm disk; the version Coifman began with had been recorded from a radio broadcast of the 78. By then the music, competing with pops, hiss, and static, was all but inaudible. Wavelet techniques made it possible to remove enough noise to hear Brahms playing.

Wavelets are also helping researchers clean house in theoretical statistics. "As soon as we were exposed to wavelets, we made the equivalent of about ten years' progress in months," says David Donoho, a Stanford University statistician who has led the way in applying the new theory. Donoho and his colleague Iain Johnstone have developed a "wavelet shrinkage" method for removing numerical noise from data. Their method, which they've shown to be optimal from several technical vantage points, first decomposes data into wavelets and then shrinks each wavelet component according to a rule that eliminates small components altogether.

Donoho expects the insights wavelets supply to set a new agenda for theoretical statistics. Having solved many of the technical problems theorists had long struggled with, "we're in a better position to say what the right questions are for statistical theory to focus on," he says.

Indeed, that may be wavelets' most important legacy. Wavelet theory is making some of the hard-won insights of mathematicians working in the abstract reaches of analysis accessible to researchers in many fields. "Wavelets teach you a way to think about problems so that a lot of ideas in abstract harmonic analysis become natural," Donoho says. The theory of wavelets does more than simply decompose and reconstitute complicated mathematical functions. In Donoho's view, "It's a tool to restructure your thoughts."

Daubechies agrees. "A lot of things are starting to come together," she says. At the same time, she adds, "it's clear that we still need new advances in order to fulfill all the promises that we think are there."