Turbulence, Fractals, and CCDs

While turbulence has captured people's imagination for millennia, the beginnings of our current understanding date from the 1930s and 1940s.

by Paul E. Dimotakis

This was one of the first images ever to capture a high level of detail within a turbulent flow—detail enough to convince Benoit Mandelbrot (MS '48, Eng '49) that turbulence was an example of the class of mathematical creatures he called fractals. The picture was made by injecting a jet of water carrying a fluorescent dye into a tank of standing water. A laser then sliced through the flow, lighting up only the dye molecules within that slice. From Dimotakis, Miake-Lye, and Papantoniou, *Physics of Fluids*, 1983.

Two and a half thousand years ago, the philosopher Heracleitus sat on the banks of a small river near the ancient Greek town of Ephesus, in Asia Minor, tossing little sticks into the water. As he watched them float irregularly downstream on the turbulent river, he remarked, "Twice into the same river you could not enter." Despite reasonably steady initial conditions (the spring is the same) and boundary conditions (the banks are the same), the turbulent flow in the river is never the same twice. Heracleitus put in a nutshell the problem that bedevils researchers in turbulence to this day—how can you analyze something that changes randomly and uncontrollably from moment to moment?

Now, turbulence isn't always a curse. It is often a blessing—without turbulence, we wouldn't have much animal life on this planet. When we exhale, for example, our breath comes out as a little jet of gas that mixes with the surrounding air. Then, when we inhale, only a very small part of the exhaled carbon dioxide comes back in. Without turbulence, we would reinhale most of it, although, as my 11-year-old son Manolis noted, not for long. And turbulent vortex rings shed from our heart valves are crucial in helping them close. It doesn't take a large change in the flow through the valve to alter its dynamics and cause life-threatening difficulties, as the work of Professor of Aeronautics Mory Gharib (PhD '83) and others is helping us appreciate. Any creatures that didn't master the dynamics of turbulence in their breathing and internal circulation, as well as in other important turbulent-flow phenomena (such as swimming and flying) would have rapidly gone extinct.

More broadly, we rely on turbulent mixing to sustain and drive all kinds of things, including many flow and combustion devices in which chemical reactions occur. Consider a jet engine, for example. Our ability to fly at high speeds is limited, in part, by our ability to mix fuel and air quickly and efficiently at flow speeds that are high compared to the speed of sound, i.e., at high Mach numbers. The inherent unsteadiness that leads to and sustains turbulence tends to diminish as the Mach number increases. Flows that would be strongly turbulent at low Mach numbers often aren't at high Mach numbers, and less mixing results. But at the same time that we're trying to maximize mixing within the engine, we need to minimize mixing (and thus heat transfer) in the flow along the engine's interior surfaces, so that they don't melt. Partly as a consequence of such considerations (and many others—flight, especially commercial flight, is a complex interplay between economic as well as aerodynamic forces), we've been flying at the same speed for the last 30 years or so—except for the Concorde, which is not economically viable because of its high fuel consumption for its size. That's a remarkable statistic, considering commercial aviation's enormous progress in so many other ways. So, if you ask whether it will always take this long to fly across the Pacific, or to Eastern Europe, the answer partly depends on learning how to both promote and limit turbulent mixing.

While turbulence has captured people's imagination for millennia, the beginnings

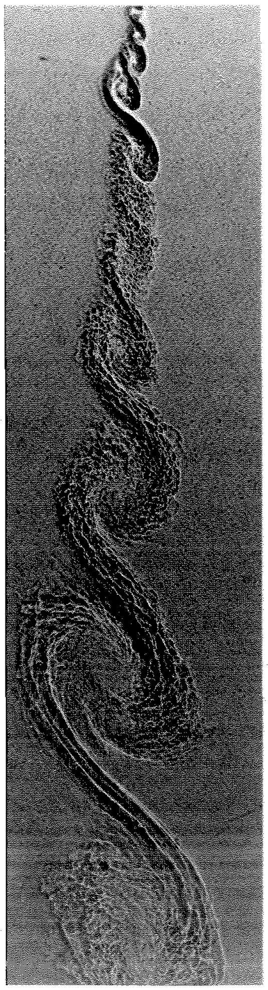

This shadowgraph, and others like it, provided the first evidence of large-scale order in turbulent flows. Here, a stream of nitrogen at four atmospheres pressure (left) traveling at 1,000 centimeters per second blows by a helium-argon mixture (right) with the same density and pressure but traveling only 380 centimeters per second. The zone where they mix is made visible by their different refractive indices, in exactly the same way that you see heat shimmers when looking across a blacktop parking lot in August. From John H. Konrad's PhD thesis, 1976.

of our current understanding date from the 1930s and 1940s, when Ludwig Prandtl in Germany, Theodore von Kármán at Caltech, G. I. Taylor in England, A. N. Kolmogorov in the then Soviet Union, and others elsewhere proposed that descriptions based on local averages and other statistical tools such as spectral analysis could provide useful information about the nature of turbulence, and that it was possible to describe, and even predict, some aspects of turbulent flows. Because turbulence is chaotic, irregular, and nondeterministic, statistical treatments appeared to be the only possible way to describe it. These methods often work well, in fact, but it's difficult to extract much information from them about many properties of turbulence—such as drag, entrainment, and mixing—that are important to engineers.

Then, in the late '60s to early '70s, largely as a result of experiments initiated at Caltech by Garry Brown, then a research fellow and later a professor of aeronautics (and now head of Mechanical Engineering at Princeton), and Anatol Roshko (MS '47, PhD '52), von Kármán Professor of Aeronautics, Emeritus, the picture changed. Brown and Roshko found that, despite its obvious disorder, turbulent flow is organized to a fair extent, primarily at its largest scales, as shown in this photo (left) by one of Roshko's students. The dynamical properties that engineers were struggling to understand depended on the behavior of these large-scale structures, which were present even in intensely turbulent flows. These discoveries provided hope that it would be possible to describe the dynamics of turbulence in relatively simpler terms than previously thought necessary. Heracleitus' epigram didn't sound so daunting anymore. You still couldn't step twice in the same river, but at least you could describe the river better.

From the days of von Kármán, much of the progress in turbulence has rested on the use of flow-visualization techniques. It's difficult to see a complicated, nonperiodic geometrical pattern in a sequence of numbers, when such a pattern may be obvious from a casual glance at a photograph. Our brain has an uncanny ability to decipher complexity and discover order in visual data. A two-year-old can look at a drawing and tell you whether it's a cat or a dog; that distinction cannot be made easily using the largest conventional computers.

Unfortunately, such visual data tended to be "soft" back then, because extracting quantitative information from pictures was difficult. The data were recorded on photographic film, a few measurements were made from the pictures, and a limited statistical analysis was laboriously done by hand. The "hard" mathematical treatments that most researchers were interested in mostly relied on point measurements. You'd put an instrument, or an array of instruments, in the flow and get a series of readings as the flow moved past the array. Only a few numbers—mean values, or, at most, spectra from wave analyzers—were recovered. The rich continuum of spatial and temporal properties of turbulence could not easily be discerned in such data.

In retrospect, the evidence of large-scale order in turbulence can be seen in the old point-mea-

The laser-induced fluorescence apparatus (left) is essentially a high-tech aquarium on a stand. The jet, tagged with a fluorescent dye, shoots straight down from the plenum, whose lower surface is immersed in the reservoir water. A rotating mirror of adjustable height sweeps the laser rapidly through the flow perpendicularly to its direction of travel, illuminating a cross section of the flow (right). The CCD camera then records the frozen slice of turbulence through the tank's glass bottom. With slightly different optics, the system can also take slices along the flow's axis, as in the picture on page 30.

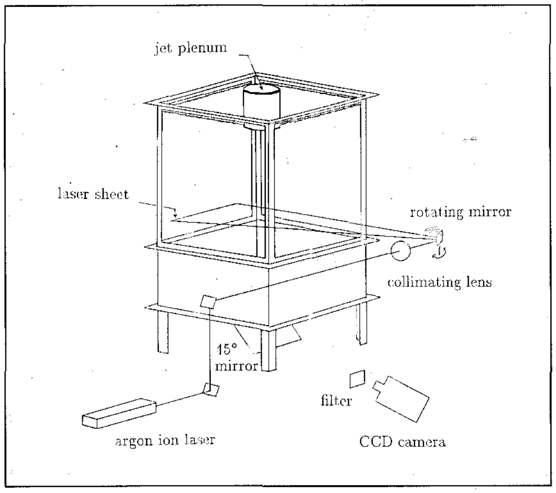

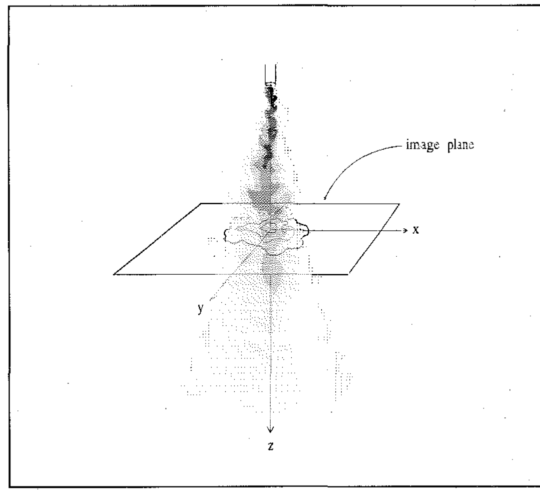

surement data, but it was so contrary to expectations that it was overlooked. In the late '40s, for example, Hans W. Liepmann, now the von Kármán Professor of Aeronautics, Emeritus, but then a young Caltech professor, was analyzing point-velocity data from a hot-wire array and found strong evidence that the points near the edge of a turbulent flow were only turbulent intermittently. Brown and Roshko's experiments some 20 years later showed why: the probe was periodically being engulfed by the largest vortices—the ones you can see in the photo on the opposite page—in the same way that a piling just above the tide line gets immersed in the swash from each breaking wave. Liepmann also noticed that these vortices tended to pair up. However, von Kármán pooh-poohed the results, and there the matter stood for two decades.

Today, with the advent of CCD (charge-coupled device) cameras, and the image-compression and data-handling technology developed by Caltech's Jet Propulsion Laboratory (JPL) and elsewhere to send us breathtaking images from planets we will not be able to visit ourselves in the foreseeable future, we can record two-dimensional information at a million or more points simultaneously, with an accuracy that matches yesterday's best point-measuring instruments. A 1,000-x-1,000-pixel CCD array is equivalent to placing one million measuring instruments in the flow, all recording at the same time without disturbing the flow or getting in each other's way.

Back in the mid-1970s, our lab at Caltech was the first to develop laser-induced fluorescence techniques for fluid mechanics. Coupled with digital CCD imaging, these methods have provided quantitative, multidimensional (field, as opposed to point) flow measurements. (We used the first [linear] CCD arrays at about the same time.) Much of our work has focused on turbulence generated by shooting a jet of water, tagged with a fluorescent dye, into a reservoir of quiescent, untagged water. A laser selectively excites the dye, which fluoresces with an intensity proportional to its concentration. A CCD camera then records the fluorescence, which shows how the jet fluid mixes with the entrained reservoir fluid. By rapidly sweeping a laser across the jet (or pulsing a sheet of laser light) we can, in effect, freeze any slice of the flow at an instant of time.

With this dense, quantitative turbulent-flow data, we can begin to ask questions about the complex geometry that turbulence generates. Geometry, to most people, brings to mind triangles and circles, spheres and cubes—the simple, regular shapes that fascinated the ancient Greeks. Well, turbulence isn't so kind. It generates irregular shapes that aren't amenable to the analyses that Thales of Miletus (near Ephesus); Pythagoras, a short swim away from Ephesus on the island of Samos; and many others developed, and that were so eloquently documented by Euclid in Alexandria in the third century B.C. How do we describe the geometric characteristics of the interface between the jet fluid and the entrained reservoir fluid in the photo on page 22, for example? How can we measure that interface's surface-to-volume ratio (a way of quantifying mixing), and determine whether it increases

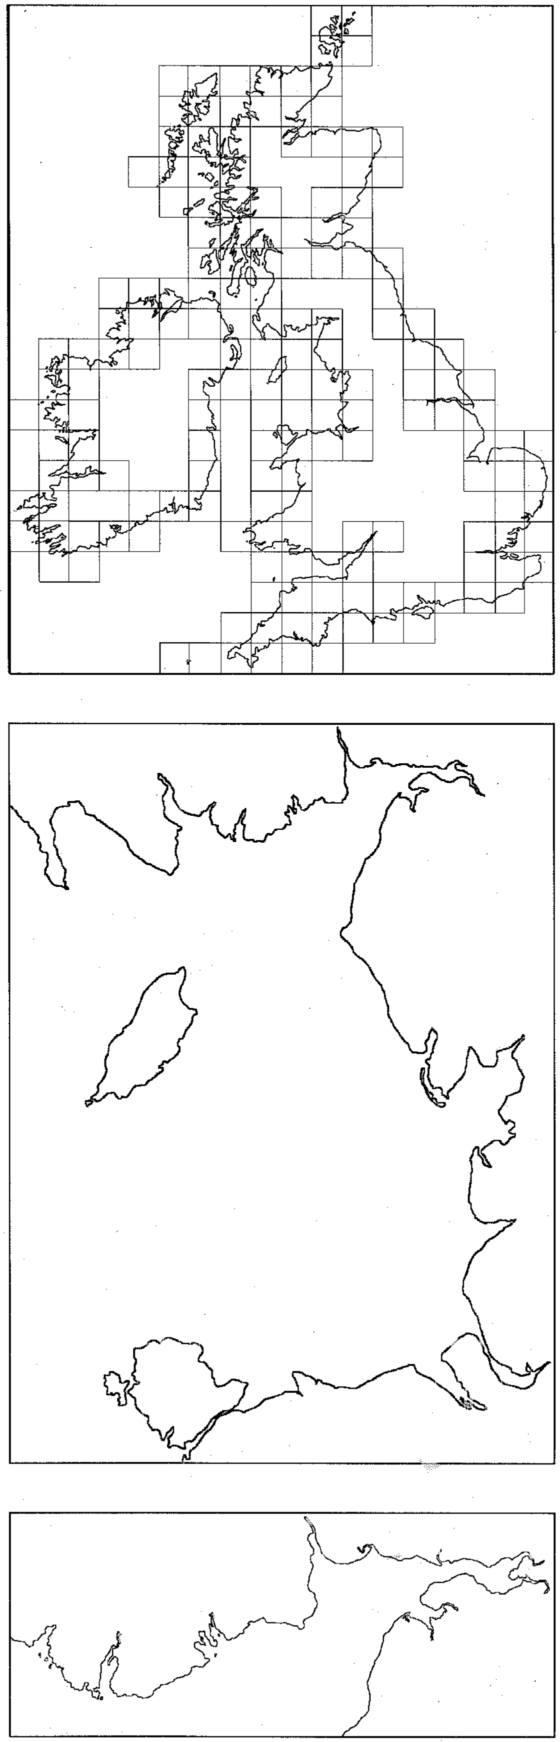

The coastline of Britain remains crinkly, whether you're looking at the entire thing (top), or just the part along the Irish Sea (middle), or just Solway Firth (bottom). If you tiled the entire map like a bathroom floor and then counted how many of those tiles covered some piece of the shoreline, you'd have a measure of how long it was. And if you made the tiles smaller and smaller, the measured length of the coastline would get longer and longer.

or decreases as the flow velocity increases?

A very exciting development took place about 20 years ago, when Benoit Mandelbrot (MS '48, Eng '49) coined the term "fractal" to describe the geometry of irregular objects, and suggested that special tools previously limited to a relatively arcane branch of mathematics could be applied to study such a geometry. His idea was that a fractal object looks equally complex, no matter at which scale you choose to examine it. You can zoom in on a small piece, or pull back and look at the whole thing, and it will look similarly complex. A coastline, for example, looks convoluted, whether you're looking at the entire coast of Britain, or just the part along the Irish Sea, or just Solway Firth, or just the harbor at Kirkcudbright, or just a piece of the rocky strand at low tide.

The ideas behind fractal mathematics had been put forth in bits and pieces by many people, but were first applied to natural phenomena by Lewis Fry Richardson, who, in a paper published posthumously in the early '60s, actually did a fractal analysis (he didn't call it that, of course) of Britain's coastline. Imagine that you have a map of Britain, and you're trying to describe how crinkly the coastline is. There are several possible approaches, but a good one is to draw a so-called bounding rectangle that just barely contains the coastline. You then fill the rectangle with contiguous, nonoverlapping tiles and count how many of them cover some segment, however small, of the coast, including all the nearby islands. As we make the tiles smaller and smaller, we need more and more of them to cover the

Plotting the logarithm of the number of tiles needed to cover the coastline, N(λ), versus the logarithm of the tile size, λ, should give a straight line whose slope, D, is the fractal dimension. But the line isn't exactly straight, so D goes from 2 when the tiles are large down to 1 when the tiles are tiny.

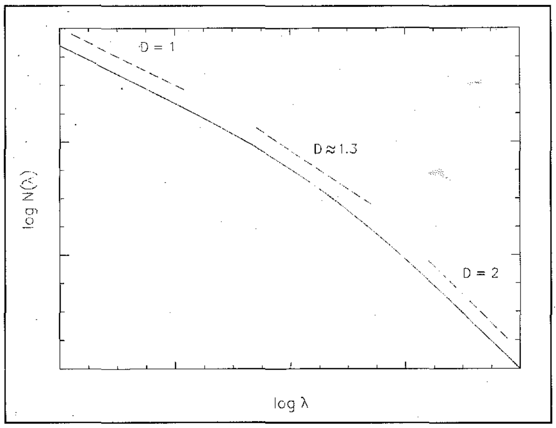

same coastline. We can plot the logarithm of the number of coastline-covering tiles, N(λ), for a given tile size, λ, versus the logarithm of λ. According to Richardson and Mandelbrot, we should get a straight line with a negative slope, i.e.,

$$\log N(\lambda) = -D \log \lambda + \text{constant}$$

In this expression, the negative slope, D, is a constant, which means that the number of coastline-covering tiles is:

$$N(\lambda) \propto \lambda^{-D}$$

This is a power-law relation, because N(λ) depends on a variable (λ), raised to a constant power (–D). The exponent, D, is called the fractal dimension. If the coastline is straight, then D = 1, corresponding to a one-dimensional object, i.e., a line. If the coastline is all scrunched up and visits nearly every point in the interior of the bounding rectangle that contains the tiles, then D is closer to 2, and the coastline approaches the solidity of a two-dimensional object. For the west coast of Britain, D is about 1.3, which means that the British coastline is not as baroque as, say, the fjords of Norway, for which D is about 1.5.

But, if you look at the log-log plot of N(λ) versus λ more closely, you see that the line isn't exactly straight. At first, each time we cut the tile size in half, the number of tiles that contain a part of the coastline (however small) is squared. D still equals 2, in other words. This is because if you're cutting very large tiles, each smaller piece will still cover some stretch of shore. The subdivided covering tiles still fill the entire bounding rectangle, as for a two-dimensional object. This is called the embedding dimension, because our fractal island (of dimension 1.3) lives in a two-dimensional space. At the other end of the curve, for very small tiles, the plot's slope approaches 1. This is because you now need as many tiles as the "arc length" divided by the tile size, λ. The arc length is simply a number—the length of the coastline as it's drawn on the map. This geometric figure is represented down to a particular resolution, and thus has a corresponding arc length, even though the actual object—the coastline itself—effectively does not. And since the real object is approximated on the map by a line, which is a one-dimensional entity, we call this the topological dimension.

The fractalists say we understand the slope = 2 region and the slope = 1 region. So we'll ignore those extremes and study the region in between, where we hope D has some fixed intermediate value. For the most part, fractals discussed to date have been of this power-law variety and describe objects whose complexity is the same regardless of scale. Mandelbrot, in fact, adopted this attribute as the defining property of fractals.

Our first laser-induced fluorescence photos were taken in a small fish tank in 1976, as part of a research project with Rick Miake-Lye (BS '78), at about the same time that Mandelbrot was formulating his proposals. Mandelbrot visited Caltech in the late '70s, and I showed him our pictures. He was very excited and asked for a slide of the picture on page 22. He presented it at the physics colloquium he gave later that day as a clear demonstration of fractal behavior in turbulence. He even referred to it as such in the subsequent edition of his book on fractals. But we had already spent some time trying to do a fractal analysis of that picture, and tried again after he left, and could not get a power law.

I hesitated to publish this counter-result, however, because it was just one picture. Perhaps no power-law behavior emerged because our statistical sample simply wasn't big enough. Every picture is different—Heracleitus's insight: the river is not the same twice—so several pictures would have had to be analyzed, and average tile counts computed for each value of λ. However, our analysis methods back then were primitive and very time-consuming. I projected the slide on a lecture-room wall, and tried to measure N(λ) by counting the number of times a λ-length string was needed to get from one end of a contour line to the other—the technique Mandelbrot had recommended.

At this point, I should explain what, exactly, we were measuring. What is the "coastline" of a turbulent jet? There are actually many lines one

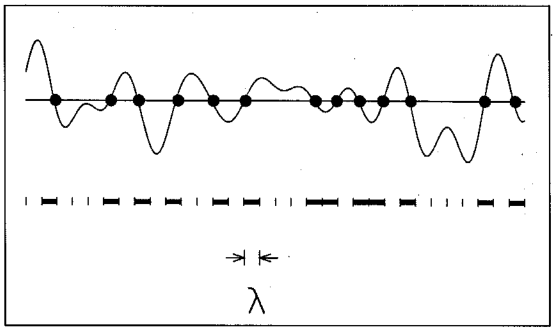

A quick lesson in chaos management: You start with some raw data (top)—a succession of random peaks and valleys—and you pick an elevation, such as the horizontal line. The set of all points at that elevation (the heavy dots) is called a "level set" and preserves the random qualities of the original data. You can now use line segments as tiles of size λ (bottom), and count how many segments it takes to cover all the data points in the level set.

can measure. Since the dye fluoresces in proportion to its concentration, given a laser sheet of uniform intensity all points of equal brightness in the image will correspond to the same concentration of jet fluid. Connecting all points of equal concentration on the image gives a set of contour lines—analogous to the elevation contours on a topographic map—called isoconcentration lines, or isocontours. Isocontours are also called "level sets," because every point in the set is at the same level—the same elevation in a topographic map, or, in this case, the same concentration. We approximated these isocontours photographically by varying the exposure while making a set of high-contrast prints (or slides) of each image.

So collecting enough data by hand to provide decent statistics would have been unthinkable, but getting the images into a computer for analysis was also difficult. We graduated from measuring isocontours off the wall to using a film scanner at JPL to get the image in digital form. However, even with the scanner, it was still so laborious to analyze a single picture that doing many of them was not in the cards. Also, despite the novelty of Mandelbrot's fractal ideas, it wasn't clear whether this approach was leading anywhere, and I didn't dare ask students to spend much time on it.

But even after we finally learned how to do a computerized analysis, we still weren't home free. There were problems with the way we were estimating the fractal dimension. We called our method the "stretched-string" algorithm because that was how we had done it on the wall. Just as a rock climber negotiates a tricky face by securing a safety rope to pitons at closely spaced intervals, the computer belayed an imaginary string from a point on the isocontour that had been reached by stretching the string from the previous point reached in the same way, and continued to do so until a complete circuit of the contour had been made. Unfortunately, the coverage counts thus derived were not unique, because there were often many choices of where to place the other end of the string, which led to different counts. And the stretched-string algorithm could cheerfully yield fractal dimensions that exceeded 2 for two-dimensional data, which we regarded as nonsensical.

We really had to wait until we could use digital imaging, initially in the form of a linear CCD array oriented perpendicularly to the jet's axis, to acquire good data in bulk. In those days, you couldn't buy a digital camera. You beat the bushes until you found a manufacturer who'd sell or give you a noisy CCD chip. Then you built all the electronics around it, converted the voltages to numbers with expensive and difficult-to-use analog-to-digital converters, and recorded the numbers any way you could. We had to develop, from scratch, the electronics to acquire and store digital data for subsequent computer processing; a technology we've been refining ever since. Our setup was first used for fractals in 1985, as part of an Aeronautics 104 class project by grad students Sheldon Green (MS '85, PhD '88) and Giancarlo Losi (MS '85, PhD '90). Able now to record digital records in the computer from the start, we could gather enough data to obtain reliable statistics. Our results continued to be inconclusive, however, and I still did not wish to venture a publication. The stretched-string algorithm remained troublesome, among other issues.

By the late '80s, we had done enough thinking and doodling that we decided we could design an experiment to settle at least some of these issues once and for all. Grad student Paul Miller (MS '87, PhD '91) and I began using laser-induced fluorescence to make long, digital records of the jet fluid's concentration, as a function of time, at a fixed point on the jet's axis. These plots looked like a slice through a very jagged mountain range—peaks and valleys in succession. We then selected all the points in time where the jet-fluid concentration crossed a fixed threshold—the one-dimensional analog of an isocontour—and "tiled" them with line segments, again counting the number of (one-dimensional) tiles required to cover the threshold-crossings as a function of tile size. And, having abandoned the stretched-string algorithm in favor of contiguous, nonoverlapping tiles, we also revisited the linear-array data from

Below: This plot of fractal dimension, D, versus the logarithm of the tile size, λ, was calculated from the one-dimensional temporal data. If the data had a power-law fractal region, it would appear as a horizontal plateau (or at least a kink) in the curve. After Dimotakis, *Nonlinear Science Today,* 1991.

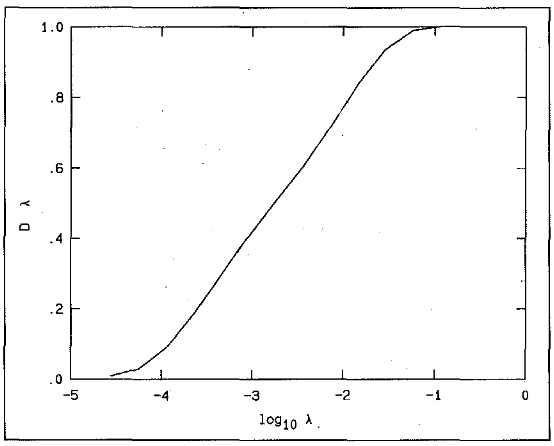

Below: Noisy data can give your level set a lot of spurious members. In this example, the real signal (heavy line) crosses the chosen elevation (dashed line) twice and twice only. But when this data is cloaked in noise (light line), scores of new points appear. To make matters worse, the new points don't just surround the real level crossings, but can also appear where the original signal approaches, but does not cross, the chosen elevation. Fortunately, during World War II, Norbert Wiener proposed a very clever scheme (classified at the time, declassified later, but not well understood until later still, when it was described by someone else) for recovering a signal from noisy data and computing its level crossings quite reliably. After Miller and Dimotakis, *Physics of Fluids A*, 1991.

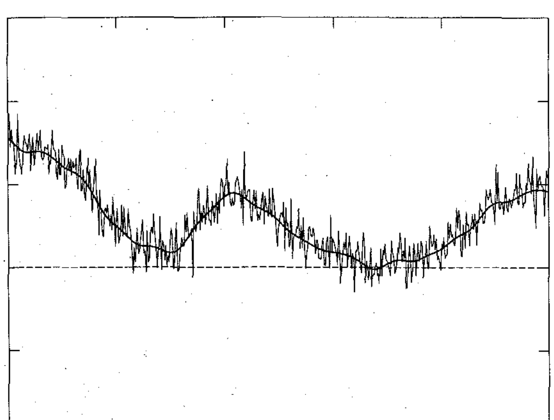

Sheldon and Giancarlo's Ae 104 experiment—still luxuriously spinning on a disk in our computer network—and calculated one- and two-dimensional tiling statistics for them too.

Our new, one-dimensional, temporal data produced statistically reliable tile-coverage counts with a logarithmic slope that smoothly increased from 0 to 1. Similarly, the space-time data from the Ae 104 experiment yielded a slope that smoothly increased from 1 to 2. In neither case did the fractal dimension (the slope, D) pause at any particular value. No power-law region appeared. I should note that, by then, many investigators had reported having found a constant fractal dimension in all kinds of flows, and the same fractal dimension to boot. So ours was a very controversial result. We were comfortable, however, with the care we had expended in our analyses to eliminate the influence of noise on the data and to understand the subtleties of the various algorithms from determining N(λ). And as our earlier data had prepared us for a curve, we did not attempt to fit a straight line to the log N(λ) versus log λ plot, which would of course have given us a constant dimension. It took 14 months to get the paper, which was eventually published in 1991, through the reviewing cycles: answering all the reviewers' queries and objections, documenting that we'd addressed and eliminated all sources of error, and, incidentally, doubling the paper's length in the process.

There could no longer be any question, at least in our minds, that turbulence generated level sets of smoothly varying fractal dimension—i.e., a continuous dependence of D on tile size—whose values were bounded by the topological dimension from below and the embedding dimension from above. There was no choice but to extend Mandelbrot's inspired proposals—insistence on uniform geometric complexity, regardless of scale, had to be abandoned if fractals were going to be useful in describing turbulence. Mandelbrot's original (power-law) fractals had to be regarded as an important special case of a broader mathematical framework, but a special case nevertheless.

Our paper caused a lot of confusion. Was Pasadena turbulence again different, as had been alleged in the early '70s, when large-scale behavior was discovered? Was it because we had primarily relied on temporal data, even though the scant spatial data we had analyzed were also in accord? So Haris Catrakis (BS '91, MS '91, PhD '96) and I forged ahead to see if our conclusions would survive the test of time and the results of improved experiments.

By then, CCD technology had progressed to

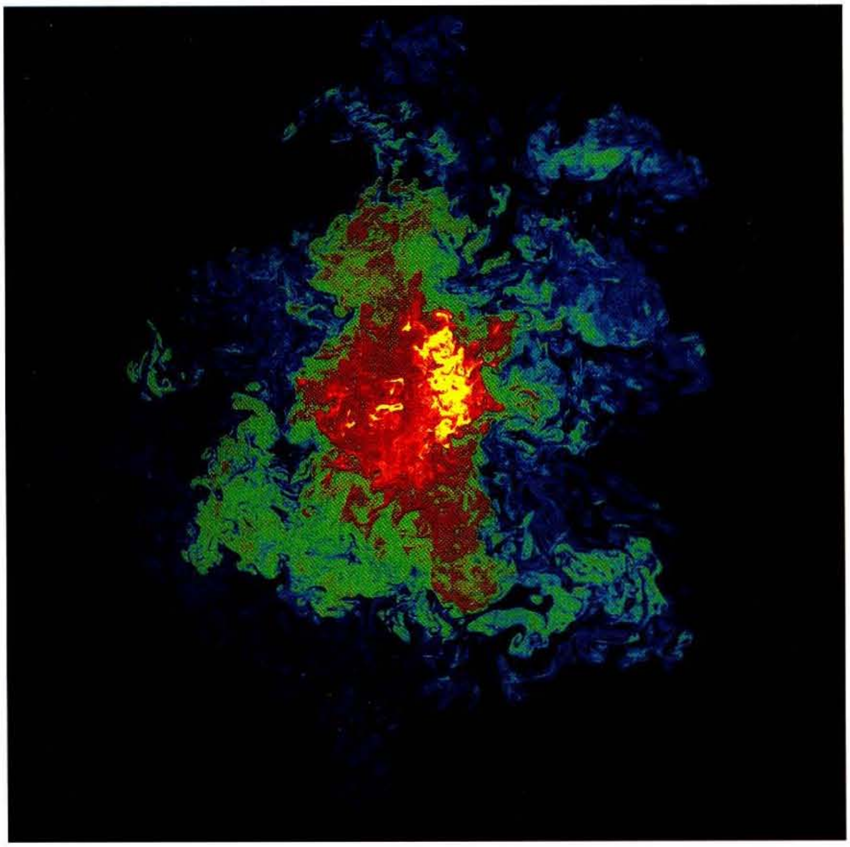

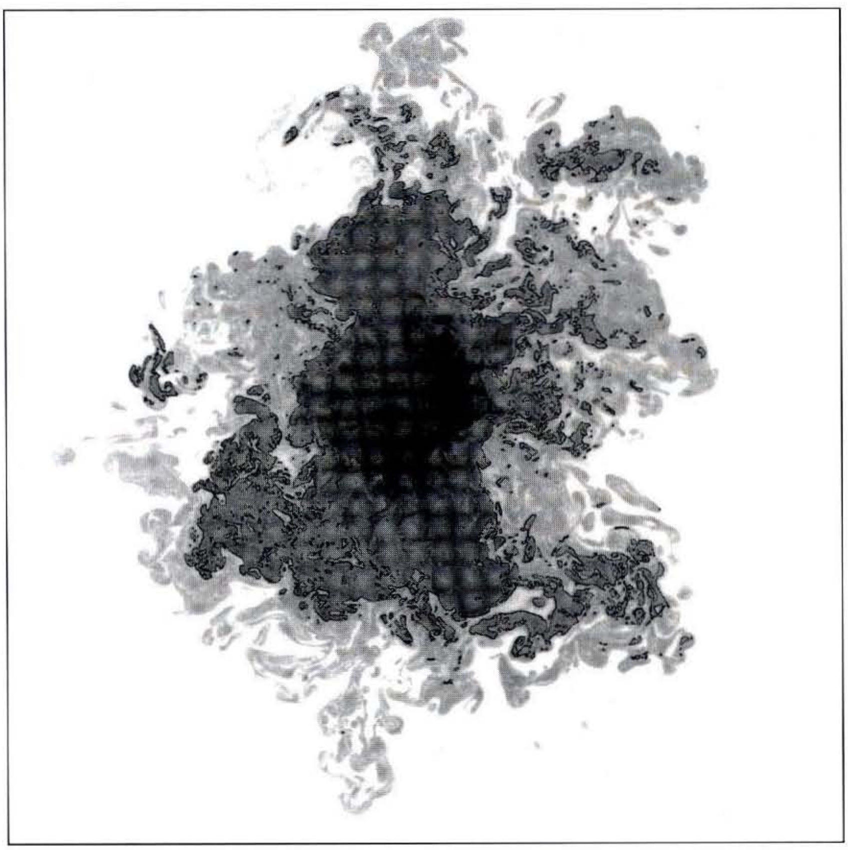

Top: A two-dimensional slice of a turbulent jet, taken with the apparatus described on page 25. The colors represent the jet fluid's concentration. Blue is the most dilute, while green, orange, red, and yellow are progressively more concentrated. Bottom: The heavy line is an isocontour calculated from the same image. Several "islands"—secondary isocontours outside the main one—and "lakes"—secondary isocontours inside the main one—can be seen. From Catrakis and Dimotakis, *Journal of Fluid Mechanics*, 1996.

where we could record two-dimensional spatial data (i.e., images) that were almost as good as previous point measurements. (A typical image appears at left.) These two-dimensional, laser-induced fluorescence slices, oriented perpendicular to the jet axis, allowed us to give the irregular geometry of two-dimensional isocontours the same rigorous statistical treatment that was only possible in one dimension before. And more powerful data-acquisition, storage, and processing systems allowed us to analyze several isocontours for each Reynolds number from images recorded over a range of Reynolds numbers.

The Reynolds number measures the relative importance of viscous diffusion in a flow. If the Reynolds number is low, viscous effects are important, and viscous damping prevents flow fluctuations and turbulence. For example, flowing honey has a very low Reynolds number and is hard to make turbulent. But the Reynolds number increases with flow speed, so water, for example, can easily be at a high enough Reynolds number to be turbulent. If you fill your bathroom sink very slowly, the water necks down as it leaves the faucet and you get a nice, smooth, laminar flow. If you turn the water up, the flow suddenly becomes unsteady. The Reynolds number has crossed a critical value above which the small, inevitable fluctuations in the flow in the pipe supplying the faucet are amplified and sustained by the flow's kinetic energy. Viscous damping is no longer sufficient to keep the flow calm. This doesn't mean that flows above some critical Reynolds number are always turbulent, only that turbulence requires a minimum

Reynolds number to be sustained.

Haris has made many important contributions in the course of his PhD research, but to make a long story short, the new images yielded the same fractal behavior as our previous data—a continuously varying D that spanned its possible range of values (in this case, from 1 to 2). This held true throughout the Reynolds-number and isocontour ranges we investigated. In summary, our experiments suggest that turbulence generates structures that are not equally complex at all scales. Instead, turbulence is more complex at larger scales (larger D), and less complex at smaller scales (smaller D), with a continuously variable D(λ) required to describe it. Haris has also found reports of such geometric behavior from other fields, for example in the analysis of alveolar tissue from rabbit lungs (the alveoli are the little sacs where gases exchange into and out of the bloodstream, so a high surface-to-volume ratio is obviously desirable), and in cloud-shape distributions. We've decided to call this variable D(λ) the "scale-dependent fractal dimension," and geometrical figures that display it "scale-dependent fractals," to distinguish them from the original "power-law" fractals.

Turbulence takes a large eddy, strains it, splits it into smaller eddies, and then again into smaller ones yet. It also merges eddies to make larger ones, and merges those again to make bigger ones yet, producing a very rich distribution of shapes and sizes.

So what have we learned from all this? We've recently realized that if you know the D(λ) curve, you can work backward and compute the distribution of spatial scales in the flow. In particular, you can say what the distribution of nearest distances to an isocontour is. Paul Miller had discovered this, in an inverse way, in the one-dimensional temporal data—a result we included in the 1991 paper. He used a random-number generator to sprinkle points on a line with a statistical distribution of his choosing to see what distribution gave a D(λ) that looked like the turbulent-jet data. He found that if the point spacings were log-normally distributed (that is, if the logarithms of the distances between successive pairs of points had a Gaussian distribution—the classic bell-shaped curve) then a fractal analysis of those spacings gave a D(λ) that very closely matched our one-dimensional data.

Haris and I have extended that to higher dimensions. Of course, what you mean by "spacing" in two dimensions must be defined, because it can be measured in many ways. In this context, we measure it as the (distribution of) sizes of the largest tiles, randomly placed, that do not touch the level set—the isocontour—at any point. Similarly, in three dimensions, one would be placing a box so it doesn't touch an isosurface, the level set in that case. We've also studied the size distribution of isocontour islands and lakes in our two-dimensional data. In this case, since we're talking about distinct objects, we can define "sizes" more naturally as the square roots of the individual areas—if the island or lake were square, its size would be the length of its edge. We found that this size distribution is also log-normal. Nature often generates log-normal distributions whenever it merges or subdivides things, and turbulence, in a way, makes islands and lakes by the fusion and fission of eddies. Turbulence takes a large eddy, strains it, splits it into smaller eddies, and then again into smaller ones yet. It also merges eddies to make larger ones, and merges those again to make bigger ones yet, producing a very rich distribution of shapes and sizes. The largest eddies are bounded by the full spatial extent of the turbulent region, while the smallest ones have sizes dictated by viscosity and, in the case of concentration data, diffusion.

What does this mean in the real world? These geometrical properties are important in describing the non-premixed combustion of hydrocarbons, for example. If we ignite aviation fuel in a jet engine, the burning rate is typically not limited by the rate of the chemical reaction of fuel and oxygen in air. The limiting factor is the rate at which the fuel mixes with the air and finds the oxygen it needs to burn, which is almost entirely determined by the characteristics of the turbulence that brings the two reactants together. The burning is confined to the unsteady, three-dimensional surface on which the mixture of fuel and oxidizer is at the stoichiometric ratio—the ratio at which the two will completely consume each other, with no leftover fuel or oxygen. This constant-concentration surface is also a level set, like the ones we've been studying in our water jets. Knowing the statistics of the distance distribution from a point to that isosurface tells us how far the fuel has to diffuse to meet the oxygen, or vice versa. Premixed combustion, as in an internal-combustion engine—in which fuel and air are mixed ahead of time and ignited later on—occurs on an equally complex combustion surface (of more-or-less constant temperature), which can also be treated as a level set.

One has to be careful when extrapolating our water-tank results to air, however, because of the differing diffusion properties of the two fluids. From a molecular viewpoint, a gas is mostly empty space. If you mix two gases together, the molecules go zipping past one another and carry their momentum some distance between collisions—there's a long mean free path. The diffusivity of mass (the molecules) and the diffusivity of their momentum is very nearly the same. But in a liquid, the molecules are, effectively, in contact with one another. There's no such thing as free flight. You can transport momentum without transporting the molecules themselves, as happens in those toys with a row of hanging steel balls—you hit the ball at one end and the ball at the other end takes off. Momentum has been transported with almost no transport of mass. To transport a molecule some distance in a liquid, many other molecules have to get out of the way. This doesn't happen easily, so the diffusion of molecules in a liquid is about a thousand times slower (slower still for a large molecule) than the diffusion of momentum. For a chemical reaction to proceed, it is individual molecules that must mix, not their momentum—we have to get the acid to meet the base, or the fuel to meet the oxygen.

In the midst of all this fuss about such images, we should not forget that turbulence is not two-dimensional. It's a three-dimensional process that evolves as a function of time, so it's really a four-dimensional phenomenon. To look at this, we would need to capture three-dimensional data—ideally, as a function of time. One way to do this is to slice the flow very thinly, very quickly—before it changes—and assemble the slices into a three-dimensional image. Kelley Scott (BS '84) tried to do this as a SURF (Summer Undergraduate Research Fellowship) project over ten years ago, but the technology was not there. Werner Dahm (PhD '85) finally did so at the University of Michigan after leaving here, but he used low spatial resolution and low flow speeds (low Reynolds numbers); it's easy to appreciate that the data rates required for decent spatial and temporal resolution quickly push this approach beyond the reach of present-day technology.

But in the last few months, we've managed an early peek at what such data might look like. Dan Lang (BS '76, MS '77, PhD '85), the group's expert in electronics and many other things, has integrated a new CCD into a digital camera and high-speed data-acquisition system that records at a resolution of 1,024-x-1,024 pixels with an excellent signal-to-noise ratio. Designed by Jim Janesick, Andy Collins, and other digital-imaging wizards at JPL, this CCD is a spare late-technology chip whose sibling will be on the upcoming Cassini mission to Saturn. Haris, Dan, and I have used the system to assemble some of the first three-dimensional, high-resolution data sets ever taken of turbulence. We recorded successive image slices in a flow for which a rate of one frame per second was almost fast enough. (By contrast, TV cameras, which have much lower resolution and signal-to-noise ratio, generate 30 frames per second.) We got our first peek at the data by computing some isosurfaces from a small portion of that space-time data, using about 60,000 polygonal facets to render each isosurface. Now David Laidlaw (MS '92, PhD '95), a postdoc in computer science, has rendered a much larger portion of the image data using up to 3,000,000 polygonal facets per surface. One of his images appears on the cover of this magazine.

And we continue to push forward—we can now achieve a framing rate 10 times higher than that first effort, yet with the same spatial resolution and signal-to-noise ratio. Or, by "binning," i.e., summing 2-x-2-pixel regions on the CCD before reading them out, we can read 512-x-512 images at 20 frames per second. To study fully developed turbulent flow, however, we estimate we would need a resolution of no less than 1,000-x-1,000 pixels per frame, read out at something like 1,000 frames per second, at Reynolds numbers large enough for bona fide turbulence. This translates into a minimum of a billion pixels per second, or, if we digitize the data stream at 12 bits per pixel, a data rate of 1.5 gigabytes per second. Dan is currently developing this kiloframe-per-second system.

We're also improving related laser-imaging technology so that we can study gas-phase turbulent mixing. To date, we've used water as our fluid because it's a thousand times denser than air, so we can get much more fluorescence signal per unit volume. We hope that, in a few years, we'll be able to compensate for the three-order-of-magnitude signal loss we'd encounter if we probed gas-phase flows. In fact, we're almost there now, after recent advances in gas-phase imaging made by Senior Research Fellow Dominique Fourguette. Since the same flow equations describe gases and liquids, comparing liquid- and gas-phase flows directly will be tremendously valuable. The only difference would be in the diffusivity of mass, so we'll be able to take a very complex phenomenon and change only one dial, leaving everything else the same. Any differences will then be attributable to turning that one dial. That's extremely valuable in science, so we're very excited by the prospect.

And finally, there's the issue of supersonic turbulence, whose nature is largely terra incognita. We need to make progress there if we're to fly faster than we do now. The usual theories of turbulence, even on the level of von Kármán, Taylor, and Kolmogorov, don't apply beyond the speed of sound, because the basic assumptions on which they rely are no longer valid. To study supersonic turbulence, one needs some way of recording the instantaneous velocity field—the simultaneous velocity of every point in the flow at one instant in time. This has not been possible, to date. A

Mapping Jovian wind speeds by ICV. A grid superimposed on Voyager 2 images of part of Jupiter's southern hemisphere (top) is distorted one rotation later (middle), with the motion of each square's central dot indicated by a line segment (bottom). After Tokumaru and Dimotakis, *Experiments in Fluids*, 1995.

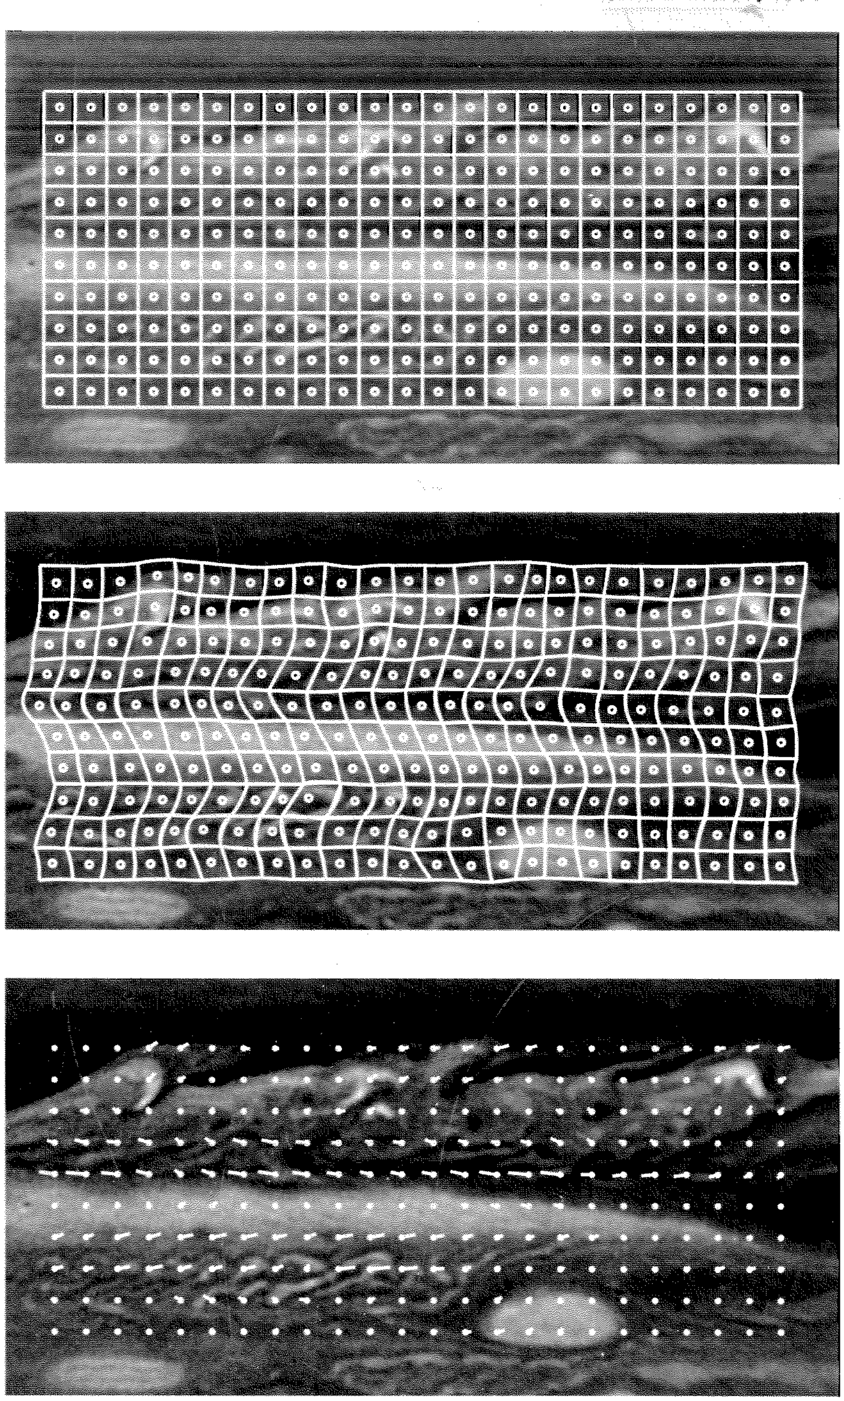

few years ago, Phil Tokumaru (MS '86, PhD '91) and I took some first steps in that direction by developing a method to deduce the velocity field from flow images recorded in quick succession. We call this method Image Correlation Velocimetry (ICV for short), and it looks for the mapping, or displacement field, that turns one image into the next one. The velocity field is then the displacement field divided by the time between successive images. Grad student Galen Gornowicz is now helping improve the method, which we've tested on a few toy flows: mapping the velocity field around an accelerating wing section whose performance is known, for example, and measuring wind speeds on Jupiter from a pair of images obtained from JPL. Our method works for the simple laboratory flows we've tested it on thus far, and we've been told by our friends in planetary science that our Jovian wind speeds are right. To use this method, however, one needs high-signal-to-noise-ratio images that are close enough in time to be reasonably well correlated. To do this in a supersonic flow, we have to solve the gas-phase-imaging problems mentioned above *and* record images as close as a few microseconds apart. So our JPL friends have designed and helped us fabricate a CCD that can record two high-signal-to-noise-ratio images with the requisite microsecond-scale spacing, which can then be read out and digitized at ordinary framing rates. We call this device the "Mach-CCD," after the flow speeds it is intended to decipher. We're bench testing it now.

There's a need for improved digital imaging in many fields. Chris Martin, professor of physics, and astronomy grad student Brian Kern have built a system that records optical phase-front distortions—the twinkle in starlight. Dan and I used our Cassini CCD camera system, in parallel with their system, to record 10- and 20-frame-per-second sequences of high-quality, short-exposure images on the 200-inch Hale Telescope at Palomar. In this collaboration, the hope is to understand how atmospheric turbulence causes optical distortion and test ideas for correcting it. Excited by these prospects, we've decided to up the specs of the kiloframe-per-second system that Dan is developing so that it will be able to run at that data rate for longer times, for both turbulence and astronomical applications. Scott Fraser, the Rosen Professor of Biology, and Professor of Physics Jerry Pine are interested in applying this capability to biological imaging, and we look forward to working with them. In these and other areas, the ability to follow high-resolution, high-signal-to-noise, two-dimensional data in millisecond or smaller time intervals would put

A fireball that the eye perceives as volume-filling in fact consists of complicated three-dimensional isosurfaces in constant random motion. But you've got to look fast to freeze them—this is a 1/1000th-second exposure.

So that's where we are. As with much of science, progress often awaits the development of a new technology, which allows a better view of nature, which improves our understanding, which begets new questions, which in turn await a new technology, which …

many important phenomena within direct reach of quantitative scrutiny.

So that's where we are. As with much of science, progress often awaits the development of a new technology, which allows a better view of nature, which improves our understanding, which begets new questions, which in turn await a new technology, which … So it is with our quest for a better description of turbulence. In this round, first there was the excitement of fractals, because they promised a description of complex geometry. Then came the disappointment when we realized that we couldn't test the idea because we couldn't record and analyze adequate data to check it. Then the technology arose to do so, followed by the disappointment when we found that turbulence wasn't a power-law fractal. And now there's the excitement of realizing that the mathematics of fractals can be extended to accommodate the behavior that our experiments have revealed. Fractal language gives us the proper tools to talk about turbulence, if you're not bent on fitting straight lines to things that are curved. The new scale-dependent fractal dimension contains a lot more information and is better able to describe turbulent mixing and combustion. But valuable as that is, it isn't enough. We need local velocity-field information along with the isosurface-geometry data, a need that has spurred the development of Image Correlation Velocimetry. Soon we'll be able to derive the local velocity field from the same set of images that will give us the isoscalar geometry—two birds with one stone!

Every now and then we make a little bit of progress in understanding turbulence, and then there's a long wait until the next step. We think this research will lead to a big jump in our understanding, and we're excited. This jump took the confluence of a new idea—Mandelbrot's proposal to apply the notion of fractals to turbulence—and two technologies—the advent of digital imaging, which generates large amounts of high-quality data directly in computer-manipulable form, and an astonishing increase in computing power. All three components are advancing today, and we're relying on their continued progress for the next jump. Will it be the last? Victory over turbulence has been declared on a semiregular basis, every time based on different means. And every time, turbulence has risen, undefeated, to mock us. The ancient Greek gods may well have left this piece of the classical world as their legacy to remind us of the perils and pitfalls of hubris. □

Paul E. Dimotakis arrived at Caltech as a freshman in 1964, and has been here ever since, earning his BS in physics in 1968, his MS in nuclear engineering in 1969, and his PhD in applied physics in 1973. He then rose through the professorial ranks, becoming the Northrup Professor of Aeronautics and Professor of Applied Physics in 1995. Besides turbulence, his research encompasses such other fluid-mechanical phenomena as cavitation and gas dynamics.

As a consultant, he has participated in the development of pilotless drone aircraft, high-power chemical lasers, the stealth fighter, the space shuttle, sealed computer hard disks, and helped with the fluid-mechanics design for the "Leap-Frog Fountain" at Disney's Epcot Center in Florida. An avid sailor, he was a member of the America^3 team and contributed to the sail design for their successful defense of the Americas Cup in 1992.

The work on turbulence, fractals, and digital imaging described in this article has been supported over the years by the Air Force Office of Scientific Research.