advanced research that NSF funds. I see it as a really good working partnership. …

On why NIH receives larger funding increases than other science agencies:

What it tells me is that we're not making the case [for the value of research] as strongly as we need to. It's relatively easy to make the connection between the involvement of NIH and a cure for a disease, or for sequencing a gene that explains Alzheimer's. What's not so easy to make is the case that it's the fundamental research that allows that gene to be sequenced and that will provide a 30% return on investment for the economy, and that there's a relationship between our standard of living and our investment in science and technology. We don't make that case very well, and I'm open to ideas on how to do it better.

NETWORKS

Fractals Reemerge in the New Math of the Internet

Traffic on the Internet has unpredictable bursts of activity over many time scales. This fractal behavior has implications for network engineering

Even a casual user of the Internet knows it is nothing like the phone system. Punch a number on your telephone, and the call will nearly always go through promptly, but click on a Web link or call up your e-mail for that urgent message, and you could be in for a long wait. This phenomenon reflects what may be the most fundamental difference between telephone service and the Internet: "It is no longer people doing the talking," says mathematician Rob Calderbank of AT&T Labs Research in Florham Park, New Jersey. Instead, it is computers talking to computers. As a result, he says, "the statistics of calling have completely changed."

Since the turn of the century, the telephone system has been built on the assumption that calls arrive at any link in the network in what's known as Poisson fashion: The likelihood of a call arriving at any given time is independent of earlier calls, for example, while call length varies only modestly. As a result, call volume fluctuates minute by minute, but over longer time scales the fluctuations smooth out. In contrast, AT&T mathematician Walter Willinger and his collaborators have shown that the machine chatter over the Internet is fractal. It has a wild, "bursty" quality that is similar at all time scales and can play havoc, Willinger says, with conventional traffic engineering.

When he first published that claim in 1993, other researchers regarded the idea as either wrong or meaningless. Willinger wasn't surprised by the reaction. Many mathematicians view the application of fractals to physical and social phenomena with some skepticism—after all, fractals have come and gone as fads in everything from hydrology and economics to biophysics. But as Willinger and network researcher Vern Paxson of the Lawrence Berkeley National Laboratory in California write in the September issue of the *Notices of the American Mathematical Society*, the fractal nature of local network traffic is now well established, and new studies have extended his analysis to the Internet as a whole. Indeed, the fractal mathematics of networks has become a fact of life for engineers designing everything from routers to switches. "Willinger's work has certainly changed the way we think about network traffic," says network researcher Scott Shenker at Xerox Palo Alto Research Center in California.

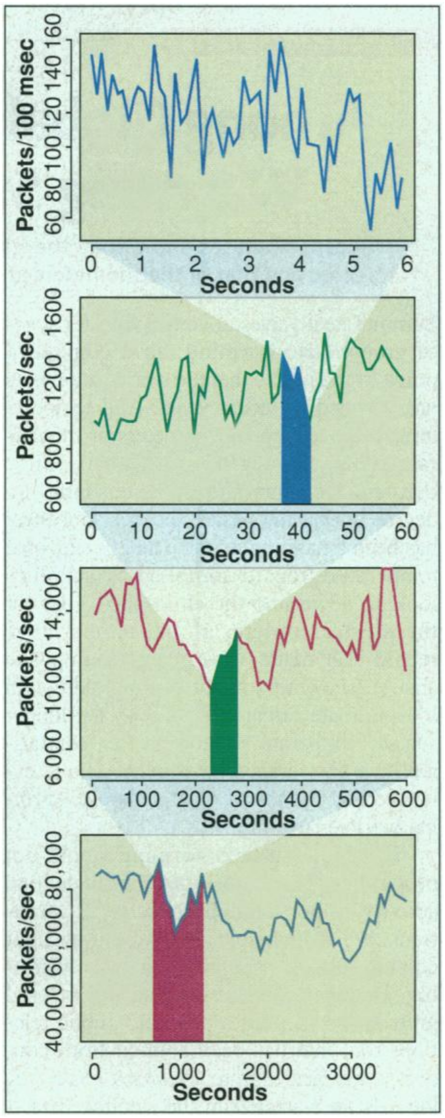

Self-similar. Internet traffic shows similar fluctuations over a range of time scales.

SOURCE: WILLINGER AND PAXSON, *NOTICES OF THE AMS*

The existing paradigm for what is known in the business as POTS, or Plain Old Telephone Service, dates back to the turn of the century and the work of a Danish mathematician named Agner Erlang, who derived a formula expressing the fraction of calls that have to wait because all lines are in use. "What he found empirically by going up to this little village telephone exchange and taking measurements with a stopwatch," says Calderbank, were the hallmarks of Poisson behavior. Call arrival times were random and independent of each other, and call durations clustered around an average value. Call frequency fell off rapidly at much longer durations.

The nature of communication—and its statistics—changed dramatically in the 1970s with the coming of what are called packet networks, beginning with Arpanet and ethernets and progressing to today's Internet. Not only did computers start doing most of the communicating, but the method of sending a message through the networks also changed. Whereas telephone networks hold open a continuous line for each call, packet networks break up a message into distinct information packets—maybe several hundred packets for a relatively small message—and send each one separately to its destination.

Despite these changes, says Paxson, researchers had such confidence in the Poisson paradigm that it continued to dominate their thinking. "People were writing papers," says Paxson, "and they would bend over backward to try to fit what they were seeing into the Poisson modeling framework, because it was so compelling."

In the late 1980s, however, Will Leland and Dan Wilson at Bell Communications Research (Bellcore) in Morristown, New Jersey, put together a hardware system that could, for the first time, accurately monitor and record the traffic flow on packet networks, much as Erlang had done for his local telephone exchange. Willinger, who was at Bellcore at the time and whose background was in probability theory, analyzed the resulting data in collaboration with Murad Taqqu, a Boston University mathematician. What they saw looked nothing like Poisson behavior.

"With some relatively very simple ways of visualizing and plotting this stuff," he says, "we could immediately see the very nice nature of the traffic: When you looked at the [variation in the] number of packets per millisecond or per second or per minute, it always looked the same." Such self-similarity is a fundamental characteristic of a fractal process, as is the bursty behavior Willinger and Taqqu also observed. "You can see areas where the traffic behaves quite nicely," Willinger says, "and then periods where it's extremely variable and goes up and down like crazy."

> The discovery of fractal behavior "has certainly changed the way we think about network traffic."
>
> —Scott Shenker

Willinger and his collaborators published these findings 5 years ago to mixed reviews. Paxson, for instance, says he was "deeply skeptical" when he first read the paper, then tried to disprove it and couldn't. Now he describes himself as a "missionary zealot." To convince the rest of the community, says Willinger, he had to explain the fractal behavior, not just describe it. He did so with Taqqu, Wilson, and Robert Sherman, another Bellcore researcher, in a paper published in 1995.

To get at the root of the fractal behavior, the researchers looked at traffic between source-destination pairs in a network. They found that characteristics of the traffic—the duration of busy periods, for instance, or the size of the transmitted files—had what's known as a heavy-tail distribution. Whereas telephone call durations, in a Poisson distribution, are tightly clustered around a mean value, heavy-tail distributions include large numbers of values that are arbitrarily far from the mean.

Telephone calls, for example, might have a mean duration of a few minutes, never lasting less than a few seconds and rarely extending beyond 15 minutes—the classic three standard deviations that encompass 99% of the distribution. But machines communicating on a network don't have the same habits as humans on a telephone. The researchers found, says Willinger, that "the busy or idle periods could last from milliseconds to seconds to minutes and even longer." The actual size of the documents being sent also varied by as much as six or seven orders of magnitude. When traffic sources have this heavy-tail behavior, says Paxson, "there are theorems that say that you're going to get fractal correlations in your traffic."

In their latest work, presented this month in Vancouver at a meeting of ACM SIGCOMM, a networking association, Willinger and his AT&T collaborators Anja Feldmann and Anna Gilbert found that the structure of packet networks themselves also contributes to the bursty nature of the traffic, at time scales of less than a few hundred milliseconds. At least one reason for the behavior, says Willinger, is the way the dominant network protocols break up each electronic message into hundreds or thousands of packets before sending them over the network.

As Willinger and Paxson describe in their September article, they and others have now documented fractal behavior for traffic on the Internet as well as on smaller networks. "Everyone buys its existence now," says Shenker. "The only debate is over how much it affects design issues." Willinger and others cite one instance in which it already has: in the design of buffers for Internet routers, which store packets during busy periods until they can be sent onward to their destination. "If you look at first-generation Internet switches," Willinger says, "buffer sizes were very small, maybe big enough for a couple hundred packets. Now, they're two or three orders of magnitude larger, because engineers realized very quickly that with the fractal nature of traffic, buffers have to accommodate much more variable traffic than was assumed in a Poisson world."

Other AT&T researchers are monitoring network traffic to extend Willinger's work, says Calderbank. "Today we're really where Erlang was around the turn of the century," he says. "Like Erlang, we are trying to understand the fundamental nature of data traffic by taking measurements. If we could understand the mechanisms at work, then we could do the engineering so that applications would run better."

–GARY TAUBES

GLOBAL CHANGE

Among Global Thermometers, Warming Still Wins Out

Recent analyses show that the gap between the satellite temperature record and that of thermometers at the surface is more apparent than real

Summer heat waves, together with forecasts of greenhouse warming, have convinced much of the public that the world as a whole has warmed in recent years. And temperatures recorded by thermometers at the surface show warming of about a half-degree during this century and a couple tenths of a degree during the past 2 decades. But there has been a nagging doubt: The 20-year-long temperature record compiled by satellites looking down into the atmosphere—by far the most complete, global temperature record ever made—has given the opposite answer, showing a slight cooling. Although most climate researchers rely on the longer surface temperature record, a few contrarians have seized upon the satellite data as evidence that the threat of greenhouse warming has been overblown.

The slight surface warming would not prove that greenhouse gases from human activity, rather than natural climate variations, are responsible. Nor would a slight cooling rule out a future greenhouse warming. But the apparent cooling has offered greenhouse skeptics a powerful public relations tool that has been applied from congressional hearings to *Reader's Digest*. If there is no warming in the satellite data—"our only truly global record of lower atmosphere temperature," as greenhouse skeptic Patrick Michaels of the University of Virginia, Charlottesville, puts it—then the surface data must be flawed and the threat of greenhouse warming much exaggerated.

Now, however, in the wake of new analyses of the satellite data, most researchers are more convinced than ever that the satellite cooling trend is not the show-stopper contrarians make it out to be. After considering the effects of El Niños and volcanoes and correcting for decay of the satellites' orbit, researchers are seeing not a cooling but a small warming. The error bars in these new analyses are larger than before, but the trend is close to that in surface records. Although the contrarians still aren't budging, leading satellite analyst John Christy of the University of Alabama, Huntsville—who has been reporting a cooling trend for a decade—agrees that the satellite data are compatible with a slight warming trend. The discrepancy between satellite temperatures and model predictions of moderate greenhouse warming "isn't that large," says Christy.

The satellites in question weren't designed to monitor global warming. Launched starting in 1979 for the National Oceanic and Atmospheric Administration (NOAA), they fly from pole to pole at altitudes of about 850 kilometers carrying instruments called Microwave Sounding Units (MSUs). These pick