"End-to-End Routing Behavior in the Internet," by Vern Paxson
A retrospective review

J. Kurose, C. Partridge, and E. W. Zegura

The 2006 ACM SIGCOMM Test of Time Award has been given to Vern Paxson for his paper, "End-to-End Routing Behavior in the Internet," published in the 1996 proceedings of the ACM SIGCOMM Conference. The award "recognizes a paper published 10 to 12 years in the past … that is deemed to be an outstanding paper whose contents are still a vibrant and useful contribution today." In this review, we try to explain why we picked this paper for the award. (In that light, we should note there were a number of outstanding papers that were strong contenders for the award).

A Time of Change in Measurement

One of the reasons that the paper remains vital and vibrant today is that it marks a moment of change in network measurement.

Network measurement is as old as networking itself. In 1969, when the ARPANET was being built, Len Kleinrock at UCLA was commissioned to put together a measurement center to analyze the performance of the network. Through the 1970s and 1980s, there was a tradition of network measurement, both by users and by the network providers. Equally important was a tradition of sharing the results. So if you were curious about, for instance, path stability, you could typically ask BBN (who ran the ARPANET) or MERIT (who ran NSFNET) and either get an answer or access to their raw measurement data.

By the early 1990s, measurement inside the network was becoming increasingly hard. A combination of privacy concerns and the rise of competing Internet Service Providers (ISPs) who viewed measurements as proprietary, meant that data about the Internet's (rapidly growing) core was increasingly hard to get. This change did not mean the end of Internet measurement: indeed, just the year before Paxson's paper, Jeff Mogul had published a brilliant paper using HTTP measurements to show the benefits of persistent connections [1]. But it appeared that research was becoming restricted to measurements (like those in Mogul's study) that could be completed without access to data on how the middle of the network behaved.

It was in this environment that Paxson's paper appeared. Paxson showed that, using proper statistical techniques (notably Wolff's elegant PASTA principle), one could gather considerable information about the behavior of the network core using measurement stations solely at the edge of the network. So, wonder of wonders, a group of PCs scattered at edge sites around the Internet and collecting carefully designed measurements could give us a lot of information about how both the edges and the middle of the network operated.

As the implications of Paxson's paper spread, we saw a revitalization of the field of network measurement. It was a new kind of network measurement, combining the collection of data with more sophisticated set of statistical techniques. The need to use more sophisticated statistical techniques had started a few years earlier: the famous self-similarity paper [2] already had forced a number of researchers to learn new analysis techniques to study their measurements. Paxson's paper showed that the statistical techniques also enabled us to capture new types of measurements.
Many people tie the resurgence of interest in network measurement, and, indeed, the creation of the Internet Measurement Conference (IMC) to the work this paper inspired.

The Paper Itself

One of the paradoxes of research is that not all important papers are actually worth reading. Sometimes the result is better explained by someone else. (Or, for instance, the written word was not the innovator's best way to communicate: reputedly Einstein did far better as a speaker than a writer). But Paxson's paper is a good read for a number of reasons.

First, it starts out right. The related research is short, but demonstrates the author is fully in command of the literature going back to 1978. And the experimental methodology is clearly spelled out, such that the experiment is repeatable by someone else. (Many methodologies, when examined even casually, fail to reveal enough about the experiment that one can have confidence it is repeatable.)

Then the results themselves are both valuable and fun. A number of routing pathologies are identified and there is a thorough discussion of routing stability. Finally, the paper looks as routing symmetry (is the path in both

directions the same?) and was the first to show just how prevalent asymmetry (previously assumed to be rare), actually was.

In summary, it is an important paper and rewarding reading. The combination makes it this year's winner of the SIGCOMM Test of Time Award.

References

1. J. Mogul, "The case for persistent-connection HTTP," Proc. ACM SIGCOMM '95.

2. W.E. Leland, M.S. Taqqu, W. Willinger, D.V. Wilson, "On the self-similar nature of Ethernet traffic," Proc. ACM SIGOMM '93.