Selecting the 2008 SIGCOMM Test-of-Time Award Winner(s)

Vern Paxson
UC Berkeley / ICSI
vern@icsi.berkeley.edu

Nick McKeown
Stanford University
nickm@stanford.edu

Jennifer Rexford
Princeton University
jrex@cs.princeton.edu

SIGCOMM's Test-of-Time Award was established in 2006 to recognize "*papers published 10 to 12 years in the past in Computer Communication Review or any SIGCOMM sponsored or co-sponsored conference that is deemed to be an outstanding paper whose contents are still a vibrant and useful contribution today.*" Last year the three of us were selected to form the committee to pick the 2008 Awardee(s), with Vern Paxson chairing. In a departure from previous years, during which a single paper received the award, for 2008 the award went to a set of three papers:

- *Internet routing instability*, Craig Labovitz, G. Robert Malan, and Farnam Jahanian, Proc. SIGCOMM 1997.
- *The macroscopic behavior of the TCP congestion avoidance algorithm*, Matthew Mathis, Jeffrey Semke, Jamshid Mahdavi, and Teunis Ott, CCR 27(3), 1997.
- *Modeling TCP throughput: A simple model and its empirical validation*, Jitendra Padhye, Victor Firoiu, Don Towsley, and Jim Kurose, Proc. SIGCOMM 1998.

In this brief note we discuss the rationale that went into the selection process leading to the winning papers. First, there was no shortage of good candidates. These were drawn from looking through proceedings of SIGCOMM-sponsored conferences held between 1996 and 1998, as well as all of the issues of CCR and IEEE/ACM Transactions on Networking during that interval. (Note, we did not receive any external nominations.) We winnowed this list down to 24 candidates which we then each independently evaluated in terms of our "short list." From that process, both the work on measuring Internet routing instabilities and that on TCP throughput modeling struck all three of us as clear landmark efforts with lasting impact on the research field (Google Scholar lists more then 3,600 citations across the set).

First, the work by Labovitz et al tackled a pressing emergent problem plaguing real networks: large volumes of BGP routing messages indicating changing conditions, and thus leading to recomputation of routes, in the absence of a commensurate degree of actual instability in connectivity. Their investigation of this problem demonstrated how to analyze network routing dynamics at truly global scale, in the process discovering that deficiencies in vendor implementations of BGP were greatly exacerbating instability issues. This work resulted in vendors incorporating implementation changes leading to dramatic improvements in BGP message traffic stability. More fundamentally, the work was an early, extremely successful example of researchers working with operators and measurement data taken directly from operational networks that opened the door to numerous subsequent efforts tackling a wide range of important practical issues facing the Internet. The paper also drew the research community's attention to the many challenges facing the Internet's interdomain routing system, leading to a wide variety of research that continues to this day.

The joint work by Mathis et al and Padhye et al was of a quite different flavor. Together, these two papers changed the mental models researchers employ when reasoning about Internet transport performance. Prior to the appearance of the 1997 paper, it was already widely appreciated by network researchers that TCP's dynamics are under many circumstances very heavily dominated by its congestion-control mechanisms (as devised by Jacobson in the mid 1980s). However, the mechanisms are complex enough—involving multiple state variables and modes, and coupled with underlying window flow control and timers for both retransmissions and generation of acknowledgments—that it was quite difficult to build up intuition for how TCP performs based on the network conditions a given connection encounters. The 1997 paper by Mathis et al developed simple analytic formulae describing how the most dominant components of TCP performance scaled with the basic parameters of packet size, round-trip time, and packet loss rate—a relationship now taught to undergraduates as a basic element of thinking about Internet performance. The subsequent 1998 paper by Padhye et al then extended this framework to incorporate additional TCP mechanisms, such as timeout-based retransmission, and evaluated the refined model over an extensive set of *in situ* measurements. In particular, it addressed many of the lingering concerns that the notion of such modeling might lack sufficient nuance to apply to important situations (e.g., connections encountering enough loss that they repeatedly engage in timeout retransmissions). More broadly, the two papers served to take TCP dynamics out of the realm of ad-hoc reasoning and instead provided apt, tractable abstractions to use when examining performance in a variety of settings. Ultimately, these models then also led to notions of how to provide a degree of "fairness" when designing congestion control for non-TCP transport protocols, such as the IETF's *TCP Friendly Rate Control* and *TCP-Friendly Multicast Congestion Control* specifications.

The final question we as the selection committee faced concerned which of these two research efforts to recognize as "**the**" most outstanding and still vibrant and useful contribution. Initially, we took selecting a single contribution as a requirement, since the previous two awards went to single papers. However, the course of discussion found us circling one another, each shifting from one paper to another as we refined our views and considered the perspectives that our colleagues offered. In the end it was clear to us that both efforts reflected singular achievements of great lasting influence—and also of a quite different nature from one another—for which we could not reasonably rank one above the other. Similarly, for the work on TCP modeling, the two papers struck us as a pair that together had an additive impact well exceeding what we might try to assign to either one by itself. In consultation with the SIGCOMM Awards Chair, we therefore selected the three papers to together share the award. Reflecting on this choice now, it continues to strike us as appropriate, and we would advocate that future Test of Time selections should likewise consider recognizing multiple papers.