2007 Sigcomm Test-of-Time Award Recipient

"Toward an Active Network Architecture"

David. L. Tennenhouse and David J. Wetherall

ACM Computer Communication Review, 26(2) (April 1996), pp. 5-17

The ACM SIGCOMM Test of Time Award http://www.sigcomm.org/tot/ recognizes a paper published in the past 10 to 12 years in ACM Sigcomm 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.

The recipient of the 2007 Test-of-Time award is "Toward an Active Network Architecture," published in the April 1996 issue of ACM SIGCOMM Computer Communications Review. Rather than designing and evaluating a solution to a specific technical problem, the paper articulates a sweeping vision of a future Internet that allows users to "program" the network elements – an "active" network. The paper argues that we need a systematic way to interpose computation inside the network, rather than burying our heads in the sand about the proliferation of "middleboxes" like firewalls and Web proxies. Active networks could be built, the paper argues, by capitalizing on the advances from the early to mid 1990s in mobile code and component-based software engineering.

The paper, and the many others that followed, launched the networking research community into a passionate debate in the mid to late 1990s about the merits of active networks, relative to traditional "passive" networks that provide basic end-to-end communication. The US Defense Advanced Research Projects Agency launched a research program on active networking, and the EU launched the FAIN (Future Active IP Networks) program. The paper was an early example of many things that seem commonplace now – the importance of in-network functionality (known as "active services" in the 1990's, and now including NATs, SIP gateways, WAN accelerators, and routers performing deep-packet inspection), the close relationship between distributed systems and networking research (as evidenced by vibrant conferences like the Symposium on Networked Systems Design and Implementation), and the interest in revisiting the architectural underpinnings of the Internet.

The paper proposed two approaches to active networks – integrated and discrete – with different notions of programmability and control. In the integrated approach, every message (or "capsule") contains a program that the network elements execute on its behalf. The users are in charge. Capsules require a new programming environment for writing, and executing, software over a distributed collection of nodes. Not surprisingly, when researchers critique active networks, they point to the risks and overheads of user packets carrying code. The paper identifies the key challenges to realizing capsules: safety (restricting the actions that the mobile code can perform) and resource allocation at the network elements. These were "active" areas of research (if you'll excuse the pun) in the mid-to-late 1990s.

The discrete approach is a modest view of active networks, where the network elements select a program to run based on bits in the packet header. In this case, network operators are in charge. Different programs can be executed for different users and applications. Limiting programmability to the network operator reduces the risk of unsafe code, and allows large programs to run without the overhead of carrying the code in the data packets. The discrete approach is deployed today in systems like PlanetLab, where the nodes support multiple "slices" that each run customized software, such as a research experiment or a long-running service. The GENI (Global Environment for Network Innovations) facility envisioned by the U.S. National Science Foundation has a similar notion of slice-level programmability.

The paper "Towards an Active Network Architecture" is surprisingly relevant today, years after the heated debates on active networks have subsided. The paper is widely read in graduate networking courses, allowing new generations of students to ask fundamental questions about whether networks should be programmable and, if so, what functionality if valuable to provide within the network, and who should program these applications. We hope that the Test-of-Time award will spur yet others to read this outstanding paper and consider the lessons that it (as well that of research that followed on its heels) teaches, as researchers investigate the design of a next generation network architecture.

David Wetherall is currently with Univ. of Washington and INTEL Research (djw@cs.washington.edu), and David Tennenhouse is with New Venture Partners (dtennenhouse@ieee.org).

We hope you enjoy reading (or rereading) the paper as much as we did!

2007 Test-of-Time Committee:

Christophe Diot
Thomson

Jim Kurose
University of Massachusetts (Chair)

Jennifer Rexford
Princeton University