

March 30, 2017 | By Doug Ramsey

Computer Scientists Honored for ‘Tracing’ Research That Stood 10-Year Test of Time

Faculty from UC San Diego, Brown University, and UC Berkeley share in networked systems award

At the USENIX Symposium on Networked Systems Design and Implementation (NSDI) this week in Boston, Mass., a team of researchers accepted an award for the most influential paper among those presented a decade ago at the annual conference. The 2017 NSDI Test of Time Award was presented during a luncheon on March 26 to two former graduate students at UC Berkeley who co-authored the paper published at NSDI 2007, along with their three UC Berkeley advisors

Rodrigo Fonseca and George Porter are now professors of computer science, respectively, at Brown University and the University of California San Diego. They accepted the award for their paper*, “X-Trace: A Pervasive Network Tracing Framework,” along with one of their former advisors, professor Ion Stoica. (Other co-authors on the paper – UC Berkeley professors Randy H. Katz and Scott Shenker – did not attend the award ceremony.)

Porter and Fonseca were still at UC Berkeley when they worked on the original paper. “We wrote X-Trace while we were Ph.D. students,” recalled Porter. “It was really an honor to work with my colleagues on this project, which formed the basis of Rodrigo’s and my Ph.D. dissertations.” Stoica remains a professor of computer science in the Electrical Engineering and Computer Science department of UC Berkeley. (It’s not Stoica’s first Test of Time award: he received the SIGCOMM Test of Time Award in 2011.)



Pictured (l-r): MIT’s Hari Balakrishnan presented the Test of Time Award March 26 at NSDI 2017 in Boston to Brown’s Rodrigo Fonseca, UC San Diego’s George Porter, and UC Berkeley’s Ion Stoica.

Modern Internet systems often combine different applications, span different administrative domains, and function in the context of network mechanisms (tunnels, VPNs, overlays and so on). In their 2007 paper, the co-authors argued that “diagnosing these complex systems is a daunting challenge.” “Many diagnostic tools existed at the time, but none existed for reconstructing a comprehensive view of service behavior,” said Brown’s Fonseca.

X-Trace was not the first tracing framework, but it was influential given that it was effectively the first framework for end-to-end tracing to focus on generality and pervasiveness. “It was based on the observation that an increasing number of systems would be built from heterogeneous components, built and operated by different people,” explained Fonseca. “In contrast, existing tracing frameworks required a specific language, or were targeted to a particular system.”

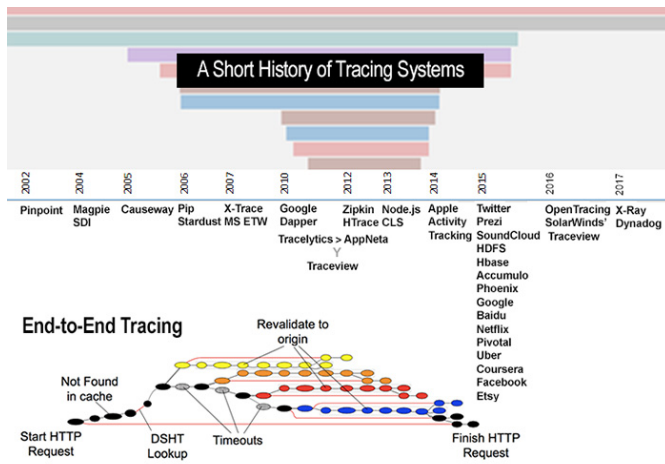
The researchers implemented X-Trace in protocols and software systems, and in their prize-winning paper, they set out to explain three different use scenarios: domain name system (DNS) resolution; a three-tiered photo-hosting website; and a service accessed through an overlay network.

Hari Balakrishnan, who co-chaired NSDI in 2007, broke the news of the Test of Time Award to the recipients. “We’re very pleased to share that your X-Trace paper from NSDI 2007 has been selected for an NSDI Test of Time Award,” he wrote. “The award honors a paper published ten years earlier at NSDI with retrospectively the most impact on research or practice.”

Indeed, the X-Trace paper has proved to be prescient – in both research *and* practice. “Today many Internet-scale backend systems are built using a ‘microservices’ approach, with hundreds of loosely connected components tied together to offer larger services,” noted Porter. “Debugging these systems effectively requires what X-Trace provided: the ability to correlate events in one component to events in other arbitrary components, even if they were many steps far removed from the first.”

The rapid adoption of tracing began with Google’s introduction of Dapper in 2010 (see graphic), which offered a similar primitive to X-Trace. Twitter’s Zipkin and Cloudera’s HTrace were open-source implementations of Dapper. Another current competitor in the market, called Traceview, also has X-Trace in its DNA after a series of startups and acquisitions dating back to 2010.

“By 2015 many companies such as Netflix, Baidu, Uber, Facebook and Etsy were deploying internal trace solutions very similar to our ideas presented in the X-Trace paper,” observed Fonseca. “And the interest persists in a rather recent initiative called OpenTracing, which is trying to standardize end-to-end tracing.”



History of tracing systems and echoes from X-Trace and Google Dapper approaches to end-to-end tracing. Graphic at bottom represents an HTTP request going through CoralCDN, a distributed content distribution network. The request is tried in parallel in four different web caches (the different colored paths) before succeeding at the last one. (Source: Rodrigo Fonseca)

The NSDI award is not Fonseca’s first for his work on tracing: he co-authored a paper on ‘pivot tracing’ that received a Best Paper award at the 2015 Symposium on Operating Systems Principles. That same year, Fonseca won an NSF CAREER Award for his work on ‘causal tracing’ to elucidate understanding of the performance of distributed systems. (Causal tracing covers a wide variety of tracing systems and frameworks, including X-Trace itself, as well as Dapper, Zipkin, HTrace, and many others.)

“It’s becoming increasingly difficult to understand how a system behaves, and, especially, how and why it fails,” said Fonseca. “Causal tracing is a technique that

captures the causality of events across all components, layers and machines, and it eases the task of understanding complex distributed systems.”

Now a co-director of UC San Diego’s Center for Networked Systems (CNS), George Porter’s research encompasses the fields of computer networking, data-intensive computing and computer systems, with a specific focus on data center networking. “I work to reduce the barrier to developing, deploying and managing applications that are able to process massive amounts of data,” said Porter. “At the same time, we aim to ensure that the resulting systems are practical, low-cost and energy efficient.”

Porter also received an NSF CAREER Award (in 2016) for work on a scalable multiplane data center network. He plans to demonstrate a hybrid electrical-optical network topology that will scale to hundreds of thousands of servers – at link rates reaching 1.6 terabits per second.

Meanwhile, the excitement surrounding tracing continues unabated. In 2017, for example, Amazon has released X-Ray, which offers distributed tracing for Amazon Web Services, and another company, Datadog, also released an end-to-end tracing product earlier this year.

MEDIA CONTACT

Doug Ramsey, , dramsey@ucsd.edu

UC San Diego's Studio Ten 300 offers radio and television connections for media interviews with our faculty, which can be coordinated via studio@ucsd.edu. To connect with a UC San Diego faculty expert on relevant issues and trending news stories, visit <https://ucsdnews.ucsd.edu/media-resources/faculty-experts>.