

January 23, 2017 | By Doug Ramsey

UC San Diego Center at NSDI 2017: Innovating in Networked Systems

Researchers affiliated with the Center for Networked Systems (CNS) at the University of California San Diego have been selected to present some of their most up-to-date research at the 14th USENIX Symposium on Networked Systems Design and Implementation (NSDI 2017).

NSDI focuses on the design principles, implementation and practical evaluation of networked and distributed systems. The annual conference will take place March 27-29, 2017, in Boston, MA, and four papers with co-authors from CNS and the Computer Science and Engineering (CSE) department of the Jacobs School of Engineering have been accepted for submission to the prestigious meeting.



Computer science professor and CNS co-director George Porter co-authored two of the research center's scientific papers to be presented at NSDI 2017.

CNS co-director George Porter co-authored two of the papers. “NSDI is one of the most important conferences for us, because just like CNS, the symposium brings together researchers from across the networking and systems community,” said Porter. “Our papers accepted to the 2017 symposium are in line with NSDI’s stated goal of pushing architectural boundaries of network services, and promoting the research dialogue on networked systems.”

vCorfu

CSE Ph.D. student Michael Wei and CSE professor Steven Swanson have co-authored with VMware Research (where Wei is currently a researcher) and Princeton University a paper on “vCorfu: Large-Scale Data Stores over a Shared Log.”

vCorfu is a strongly consistent, cloud- scale object store built over a shared log. It augments the traditional replication scheme of a shared log to provide fast reads, and vCorfu leverages a new technique – composable state machine replication – to compose large state machines from smaller ones. “This enables the use of state machine replication to be used efficiently in huge data stores,” said Wei. “We will show that vCorfu outperforms Cassandra, the popular, state-of-the-art NoSQL database for cloud apps It does so while also providing strong consistency in opacity and read-own-writes, efficient transactions, and global snapshots at the scale of the cloud.”

vCorfu is available as an open-source project on Github at github.com/CorfuDB.



CNS-affiliated CSE Ph.D. students (clockwise from top left) Michael Wei and Arjun Roy, second-year M.S. student Karthikeyan Vasuki Balasubramaniam, and recent CSE alumnus Rahul Bhalariao (M.S. '16) co-authored papers to be presented at the conference in March.

Datacenter Fault Detection

CSE Ph.D. student Arjun Roy expects to complete his doctorate in 2017, and he collaborated with his advisor, CSE professor Alex C. Snoeren, on the paper to be presented at NSDI on “Passive Realtime Datacenter Fault Detection.” It reflects joint work with Facebook researchers Hongyi Zeng and Jasmeet Bagga, who are also co-authors on the paper. (The two Facebook engineers previously co-authored a paper at SIGCOMM 2015 with Roy and professors Snoeren and Porter on “Inside the Social Network’s (Datacenter) Network”.) Roy also did internships at Facebook in the summers of 2012, 2013 and 2014.

According to the paper’s abstract, “datacenters are characterized by their large scale, stringent reliability requirements, and significant application diversity. However, the realities of employing hardware with small but non-zero failure rates mean that datacenters are subject to significant numbers of failures, subsets of packets can be dropped or delayed without triggering a fault signal, so traditional fault detection techniques (involving end-host or router-based statistics) may not identify such errors.

In their paper, Roy and Snoeren describe how to expedite the process of detecting and localizing partial datacenter faults. It uses an end-host method generalizable to most datacenter applications. “We correlate transport-layer flow metrics and the delay incurred by

network-input/output system calls at end hosts with the path that traffic takes through the datacenter,” said Roy. “Then we apply statistical analysis techniques to identify outliers and localize the faulty link and/or switch or switches.

The paper will detail how the researchers evaluated their novel approach in a production datacenter (Facebook’s) carrying a workload servicing more than 100 million users.

ExCamera

In light of the massive explosion in video content on the Internet and for virtual reality, a team of two CSE Master’s students advised by professor George Porter has come up with a new approach to processing video with minimal delays. Second-year M.S. student Karthikeyan Vasuki Balasubramaniam (who is Porter’s teaching assistant this quarter in CSE 124 on Networked Services) and recent graduate Rahul Bhalerao (M.S. ’16) have had experience in industry (both at Amazon—Balasubramaniam as an intern at Amazon Prime, and Bhalerao currently working at Amazon Web Services).

The paper accepted to NSDI is entitled “Encoding, Fast and Slow: Low-Latency Video Processing Using Thousands of Tiny Threads.” In it, the researchers describe ExCamera, a system that can edit, transform and encode a video, including ultra-high-resolution 4K video (four times the resolution of high-definition TV) and stereoscopic virtual reality (VR) material, dozens of times faster than cutting-edge production systems at the largest providers.

The co-authors lay claim to two major contributions. First, “our coauthors at Stanford developed a novel encoding strategy focusing on fine-grained parallelism, which is rather unique in the encoding space,” explained Balasubramaniam.

Separately, noted Bhalerao, “ExCamera orchestrates encoding and other video-processing pipelines across the Amazon Web Services Lambda service. The system invokes thousands of threads in parallel, each handling only a fraction of a second of the video.” The UC San Diego was done in collaboration with researchers at Stanford University.

MegaSwitch

MegaSwitch is a multi-fiber ring optical fabric that exploits space-division multiplexing across multiple fibers non-blocking communications that can be rearranged to 30-plus racks and 6,000-plus servers. CNS’s George Porter co-authored the paper on “Enabling Widespread

Communications on Optical Fabric with MegaSwitch” with researchers at the Hong Kong University of Science and Technology, SUNY Buffalo, Yale University as well as Omnisense Photonics and CoAdna Photonics. (No UC San Diego students worked on the paper.)

According to Porter, “we were seeking an optical interconnect that can enable unconstrained communications within a computing cluster of thousands of servers.” Indeed, existing wired optical interconnects are not ideal for widespread communications in production clusters, and recent efforts to reduce the time it takes to reconfigure the optical circuit from milliseconds to microseconds only partially mitigated the problem (by rapidly time-sharing optical circuits across more nodes).

“We were still limited by the total number of parallel circuits available simultaneously,” explained Porter. “However, we wanted to evaluate the potential of WDM to scale to a large number of endpoints.”

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>.