

Spring 2017 Seminar Series

WHY NEW BUGS LIVE IN THE CLOUD? (AND HOW TO EXTERMINATE THEM)

Tuesday March 21, 2017

1:00 PM - 4:00 PM - HEC 450

As more data and computation move from local to cloud environments, datacenter distributed systems have become a dominant backbone for many modern applications. However, the complexity of cloud-scale hardware and software ecosystems has outpaced existing testing, debugging, and verification tools.

I will describe three new classes of bugs in large-scale datacenter distributed systems: (1) distributed concurrency bugs, caused by non-deterministic timings of distributed events such as message arrivals as well as multiple crashes and reboots; (2) limpware-induced performance bugs, design bugs that surface in the presence of "limping" hardware and cause cascades of performance failures; and (3) scalability bugs, latent bugs that are scale dependent, typically only surface in large-scale deployments (100+ nodes) but not necessarily in small/medium-scale deployments.

These findings are based on the results of our large-scale studies of 3000+ bugs in datacenter distributed systems and 500+ outage instances of Internet/cloud services. I will present some of our work in understanding and combating these three classes of bugs, including semantic-aware model checking (SAMC), taxonomy of distributed concurrency bugs (TaxDC), path-based speculative execution (PBSE), and scalability check (SCK).

Haryadi Gunawi

The University of Chicago



Haryadi Gunawi is a Neubauer Family Assistant Professor in the Department of Computer Science at the University of Chicago where he leads the UCARE research group (UChicago systems research on Availability, Reliability, and Efficiency). He received his Ph.D. in Computer Science from the University of Wisconsin, Madison in 2009. He was a postdoctoral fellow at the University of California, Berkeley from 2010 to 2012. His current research focuses on cloud computing reliability and new storage technology. He has won numerous awards including NSF CAREER award, NSF Computing Innovation Fellowship, Google Faculty Research Award, NetApp Faculty Fellowships, and Honorable Mention for the 2009 ACM Doctoral Dissertation Award.

His research focus is in improving dependability of storage and cloud computing systems in the context of (1) performance stability, wherein he is interested in building storage and distributed systems that are robust to "limping" hardware, (2) reliability, wherein he is interested in combating non-deterministic concurrency bugs in cloud-scale distributed systems, and (3) scalability, wherein he is interested in developing approaches to find latent scalability bugs that only appear in large-scale deployments.

