

Case Studies from the Real World: The Importance of Measurement and Analysis in Building Better Systems

Bianca Schroeder
University of Toronto
Department of Computer Science
Toronto, Canada
bianca@cs.toronto.edu

ABSTRACT

At the core of the "Big Data" revolution lie frameworks and systems that allow for the massively parallel processing of large amounts of data. Ironically, while they have been designed for processing large amounts of data, these systems are at the same time major producers of data: to support the administration and management of these huge-scale systems, they are configured to generate detailed log and monitoring data, periodically capturing the system state across all nodes, components and jobs in the system. While such logging information is used routinely by sysadmins for ad-hoc trouble-shooting and problem diagnosis, we point out that there is a tremendous value in analyzing such data from a research point of view. In this talk, we will go over several case studies that demonstrate how measuring and analyzing measurement data from production systems can provide new insights into how systems work and fail, and how these new insights can help in designing better systems.



BIO

Bianca is an associate professor and Canada Research Chair in the Computer Science Department at the University of Toronto. Before joining UofT, she spent 2 years as a post-doc at Carnegie Mellon University working with Garth Gibson. She received her doctorate from the Computer Science Department at Carnegie Mellon University under the direction of Mor Harchol-Balter. She is an Alfred P. Sloan Research Fellow, the recipient of the Outstanding Young Canadian Computer Science Prize of the Canadian Association for Computer Science, an Ontario Early Researcher Award, an NSERC Accelerator Award, a two-time winner of the IBM PhD fellowship and her work has won four best paper awards and one best presentation award. She has co-chaired the TPCs of Usenix FAST'14, ACM Sigmetrics'14 and IEEE NAS'11, and is an associate editor for IEEE TDSC. Her work on hard drive reliability and her work on DRAM reliability have been featured in articles at a number of news sites, including Computerworld, Wired, Slashdot, PCWorld, StorageMojo and eWEEK.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

ICPE'16 March 12-18, 2016, Delft, Netherlands

© 2016 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-4080-9/16/03.

DOI: <http://dx.doi.org/10.1145/2851553.2858660>