



Fall 2018 Seminar Series

Risk-Aware Failure Identification to Improve the RAS in Erasure-coded Data Centers

WEDNESDAY, NOVEMBER 21, 2018

10:00 AM – HEC 450

Data reliability, availability, and serviceability (RAS) of erasure-coded data centers are highly affected by data repair induced by node failures. Compared to the recovery phase of the data repair, which is widely studied and well optimized, the failure identification phase of the data repair is less investigated. Moreover, in a traditional failure identification scheme, all chunks share the same identification time threshold, thus losing opportunities to further improve the RAS.

In this talk, I will outline my research in data storage and I/O systems. Specifically I will present RAFI, a novel risk aware failure identification scheme, where chunk failures in stripes experiencing different numbers of failed chunks are identified using different time thresholds. For those chunks in a high risk stripe, a shorter identification time is adopted, thus improving the overall data reliability and availability. For those chunks in a low risk stripe, a longer identification time is adopted, thus reducing the repair network traffic. Therefore, the RAS can be improved simultaneously.



Xubin He
Temple University

Dr. Xubin He is a Professor in the Department of Computer and Information Sciences at Temple University. He is also the Director of the Storage Technology and Architecture Research (STAR) lab. Dr. He received his PhD in Electrical and Computer Engineering from University of Rhode Island, USA in 2002 and both his MS and BS degrees in Computer Science from Huazhong University of Science and Technology, China, in 1997 and 1995, respectively. His research interests focus on data storage and I/O systems, including big data, cloud storage, Non-Volatile Storage, and scalability for large storage systems. He has published more than 100 refereed articles in prestigious journals such as IEEE Transactions on Parallel and Distributed Systems (TPDS), Journal of Parallel and Distributed Computing (JPDC), ACM Transactions on Storage, and IEEE Transactions on Dependable and Secure Computing (TDSC), and at various international conferences, including USENIX FAST, USENIX ATC, Eurosys, IEEE/IFIP DSN, IEEE INFOCOM, IEEE IPDPS, MSST, ICPP, MASCOTS, LCN, etc. He is the general co-chair for IEEE NAS'2009, program co-chair for IPCCC'2017, ICPADS'2016, MSST'2010, IEEE NAS'2008 and SNAPI'2007. Dr. He has served as a proposal review panelist for NSF and a committee member for many professional conferences in the field. Dr. He was a recipient of the ORAU Ralph E. Powe Junior Faculty Enhancement Award in 2004, the TTU Chapter Sigma Xi Research Award in 2010 and 2005, TTU ECE Most Outstanding Teaching Faculty Award in 2010, and VCU ECE Outstanding Research Faculty in 2015. Dr. He is a principal investigator of multiple research projects totaling more than \$3 million mainly sponsored by the U.S. National Science Foundation (NSF). He holds one U.S. patent. He is a senior member of the IEEE, a member of the IEEE Computer Society and USENIX.