



SCHOOL OF INFORMATICS AND COMPUTING  
**Computer Science Colloquium Series**



# **Yue Cheng**

**Virginia Tech**

**Tuesday, March 7, 2017**  
**3:00 PM**  
**Oak Room, IMU**

## **Breaking the Monolith: Rethinking Storage System Design**

**Abstract:** Distributed object-based storage systems serve the growing and disparate needs of almost all of large-scale web services in use today. This is driving the proliferation of a wide variety of distributed object-based stores. However, extant, monolithic, scale-out storage system designs present unique challenges in adapting to the ever-changing storage requirements in both efficiency, e.g., performance and resource efficiency, and flexibility, e.g., ease-of-use and programmability.

My research takes two crucial steps on this difficult road to optimize and design better object-based stores. In this talk, I will first show that an approach to storage system design based on a simple core principle, resource partitioning, can yield systems with significantly improved performance and resource efficiency under dynamic, skewed, and multi-tenancy workloads. I will first show how to effectively exploit fine- and coarse-grained resource partitioning in MBal, a distributed in-memory object caching system, that offers a holistic solution wherein the load balancing model tracks hotspots and applies different strategies based on the severity of the imbalance. Then, I will discuss a fundamental challenge faced by all practitioners and developers working on scalable storage: how to implement a fast and reliable scale-out storage system with minimal engineering effort? I will present how modular design rather than the extant monolithic approaches can ease the burden of designing new storage systems, especially by enabling an innovative decoupling of the control and data plane in distributed storage. I will conclude with a brief discussion of my vision for future scalable storage systems.

**Biography:** Yue Cheng is a Ph.D. candidate in the Department of Computer Science at Virginia Tech. His research interests include distributed systems, cloud storage, and Internet of Things. His work has been published in premier venues in computer systems and high-performance computing, including USENIX ATC, ACM EuroSys, and ACM HPDC. He has worked and collaborated with leading storage researchers at IBM Research and Dell EMC. He received his B.Eng. in Computer Science from Beijing University of Posts and Telecommunications. Yue loves traveling and all kinds of outdoor activities.



**SCHOOL OF INFORMATICS  
AND COMPUTING**

**INDIANA UNIVERSITY  
Bloomington**