Guest Editorial: FAST'10

We are pleased to present a selection of papers from the 9^{th} *USENIX Conference on File and Storage Technologies (FAST'10)* as a special issue of ACM *Transactions on Storage*. FAST continues to serve as the premier venue for the publication of storage systems research, and this year's conference program displayed many innovative results in performance, reliability, and the adoption of emerging technologies. FAST received 89 submissions, of which 21 were selected and presented to more than 400 attendees in San Jose, CA, from February 23 to 16, 2010. We have selected six of those papers for this special issue, including the two Best Paper Award winners.

FAST has a tradition of publishing papers that enhance our understanding of how storage systems operate in the real world. In past years, this has included workload characterization, failure analysis of disk drives, and write errors in RAID systems. Although difficult to execute, these "study papers" provide high value to the research community. A new theme emerged in this year's program, with several papers studying latent sector errors in disk drives. "Understanding Latent Sector Errors and How to Protect Against Them" by Schroeder et al. provides a detailed analysis of latent sector errors across more than thirty server systems. The study reveals temporal and spatial dependencies in the occurrence of errors. The authors then demonstrate how systems must account for these dependencies to avoid data loss and present several techniques for doing so. On another topic, in "Optimizing Energy and Performance for Server-Class File System Workloads," Sehgal et al. provide a remarkably detailed study of file system workloads to characterize the trade-offs between power usage and performance and show how these trade-offs inform the configuration and design of file systems.

FAST's two award papers extend the capabilities of the file system. "Membrane: Operating System Support for Restartable File Systems" by Sundararaman et al. furthers work on failure isolation in operating systems by extending it to file systems. The authors demonstrate how to restart file systems without interrupting the execution of the operating system. They overcome the challenge of preserving or recreating file system state across failures. "quFiles: The Right File at the Right Time" by Veeraraghavan et al. presents a new file system abstraction that allows for multiple different physical representations of the same logical data. The "quFile" provides a transparent way to access the appropriate physical data, depending upon factors like security, power, and available resources.

The program also encouraged us to embrace design principles and new technologies fully. In "I/O Deduplication: Utilizing Content Similarity to Improve I/O Performance," Koller and Rangaswami show that content-addressable storage is about more than eliminating I/O and disk compression. They enhance

DOI 10.1145/1837915.1837916 http://doi.acm.org/10.1145/1837915.1837916

the performance of deduplication by using content similarity to make cache replacement and replication decisions in conjunction with content-addressable caching and dynamic replica selection. Solid state storage will fundamentally transform storage performance, power, and reliability, but systems have been slow to redesign for flash. "DFS: A File System for Virtualized Flash Storage" by Josephson et al. shows the insufficiency of adding flash to existing systems. The authors conclude that flash is not disk and that traditional file system design principles fail to realize the potential performance of flash storage. They go on to present novel file system structures and a new storage abstraction layer for flash.

We hope that you enjoy these papers as much as we have. We look forward to seeing you at FAST 2011!

RANDAL BURNS
Johns Hopkins University
KIMBERLY KEETON
Hewlett-Packard Labs
Guest Editors