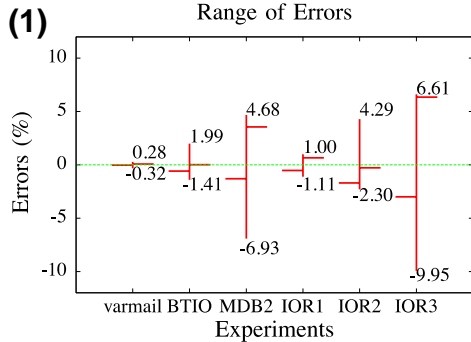

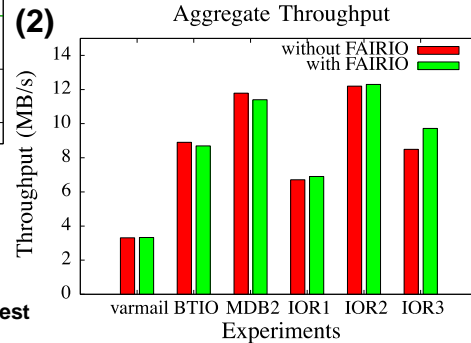


I/O Coordination to Improve HEC System Performance



I/O Performance Delivered by RAID FAIRIO



- (1) Error in Delivered Differentiated I/O Service
 (2) Aggregate Throughput

6 workloads, each comprised of 4 I/O request classes (MDB2 = MADbench2)

Novel Ideas

- Differentiated I/O service in a consolidated storage environment, i.e., I/O workload performance isolation and I/O service proportional to designated workload weights. (This work uses a patent-pending innovation that resulted from a previous DoE-funded project, DAiSES, Dynamic Adaptation in Support of Extreme Scale.)
- Coordination of productive and defensive I/O of multiple workloads, based on mathematical models that identify checkpoint intervals that optimize either execution time or number of defensive I/O operations, to best utilize system resources.

Impact and Champions

IMPACT:

- Provide differentiated I/O service with little affect on aggregate throughput.
- Enhance system performance by increasing meaningful utilization of allocated computing resources (in particular, I/O).
- Extend scalability of checkpoint/restart fault management.
- Reduce I/O system stress and resultant failures.

CHAMPIONS:

- Texas Advanced Computing Center: Jim Browne and John Hammond
- Argonne National Laboratory: Phil Carns (Darshan)

PATRICIA J. TELLER AND SARALA ARUNAGIRI, THE UNIVERSITY OF TEXAS AT EL PASO, DEPARTMENT OF COMPUTER SCIENCE

Milestones/Dates/Status (Jan 2009-Sept 2013)

	Scheduled	Actual
• FAIRIO algorithm framework	Dec 2009	Dec 2009
• Initial RAID FAIRIO Algorithm	Jun 2010	Jun 2010
• I/O characterization of periodic-checkpointing RAXML and NAMD	Dec 2010	Jun 2011
• RAID FAIRIO publication	Feb 2011	Oct 2011
Journal paper being submitted	Nov 2011	Best Paper
• Validation of model to optimize # of defensive I/O operations	Jun 2011	
• Migration to DiskSIM 4.0 and expansion of FAIRIO experiments	Jun 2012	
• Further analysis of failure data for model parameters	Dec 2012	
• FAIRIO at file system level	Dec 2012	
• Model-based I/O coordination	Jun 2013	



U.S. DEPARTMENT OF ENERGY

Office of Science



30 September 2011