

The Designs of RAID with XOR Engines on Disks for Mass Storage Systems

Tai-Sheng Chang, Sangyup Shim* and David H.C Du

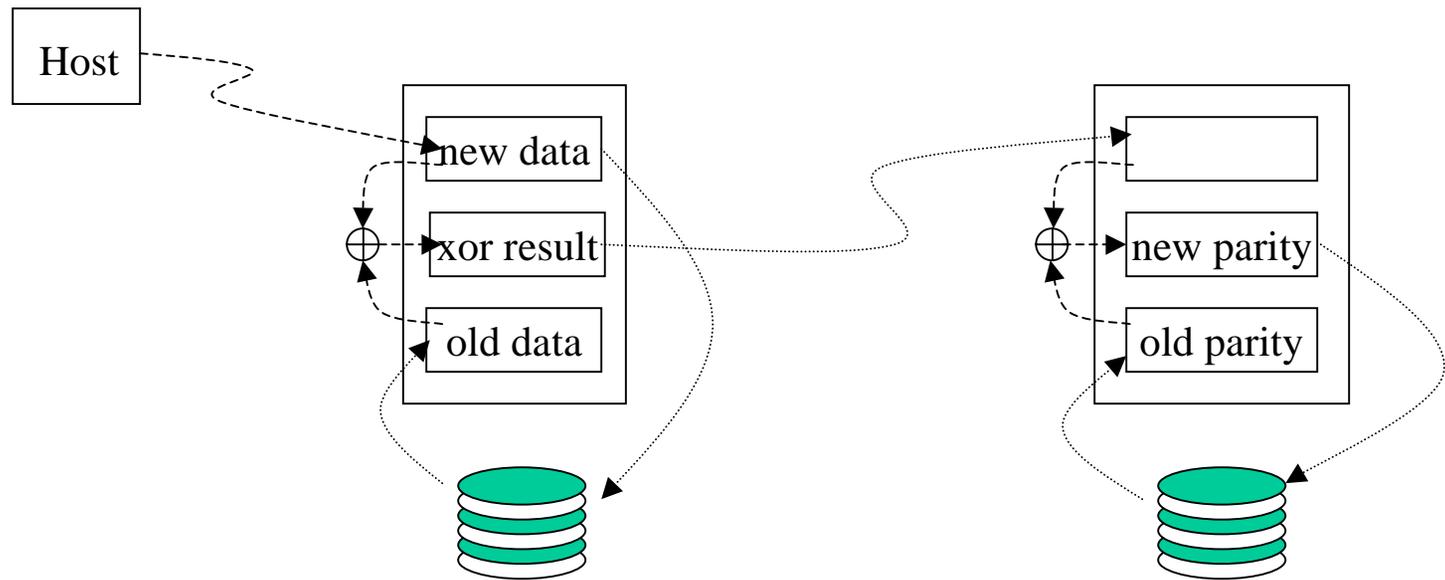
{tchang, du}@cs.umn.edu

**Department of Computer Science
University of Minnesota**

***: sishim@email.sjsu.edu**

**CISE Department,
San Jose State University, California**

Disk-Based XOR Operations



1. Read old data
2. XOR new and old data
3. Write new data

1. Read old parity
2. Compute new parity
3. Write new parity

Why Disk-Based XOR

- **Scalability**
 - Distributed XOR computations
 - larger number of disks with FC-AL
- **Lower cost than the RAID controller solution**
 - No need for XOR computation on Interface
- **Better performance than software RAIDs**
 - No need of CPU for parity update
- **Clustering with shared storage**
 - Multiple hosts with data sharing requires extra communication among host in traditional RAID

Plans and Status

- **Performance study**
 - *Sangyup Shim, Tai-Sheng Chang and David H.C. Du, “Efficient Implementation of RAID-5 Using Disk Based Read Modify Write”*,
- **Improve performance**
 - Investigate the impacts of disk scheduling on the performance and buffer requirement:
 - *Tai-Sheng Chang and David H.C. Du, “Impacts of Disk Scheduling on Disk-Based XOR RAID systems”, in preparation.*
- **Error Recovery**
- **Cluster of PC/workstations with data sharing**