Performance Models and Systems Optimization for Disk-Bound Applications

Mithuna Thottethodi[†], Rahul Shah[‡], Vijay Pai[†], T.N. Vijaykumar[†], Jeffrey S. Vitter[‡]

[†]School of Electrical and Computer Engineering
Purdue University

[‡]Department of Computer Science
Purdue University

Performance Model Enhancements

- Individual disk block access: Flat Model
- Research tasks: Refined Model (seek time, xfer rate)
  - Simple enough to reason about; design algorithms for
  - Accurate enough to model performance
  - Preliminary model within 10% of prototype disk arrays

Memory

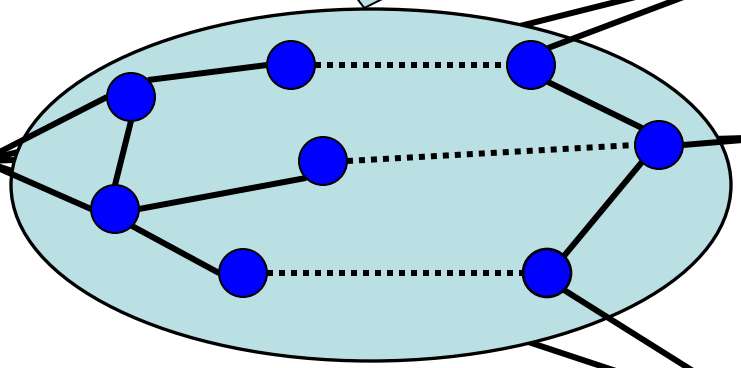

Disk 1

Disk 2

Disk D

Model-driven System Optimizations

- Application-aware Placement/Migration
  - Accessed concurrently place on separate disks
- Network-aware Placement/Migration
  - Place on separate disks AND disjoint network paths
- Seek-aware Placement/Migration
  - Amortize seek overheads over multiple block transfers

Tension among goals

- Application-aware Cost-effective System design
  - Match parallelism to application requirement
- Application-aware Disk Caching for Power/Thermal efficiency
  - Number of spun-up disks = expected concurrent accesses

Execution-driven Storage Simulator

- Goal: Capture dynamic and data-dependent access patterns for improving disk, multiprocessor, and interconnect systems
- Method: Direct Execution strategy
  - Decouple Functional version/ Timing version
  - Instrument File system to interface with disk/network timing model
  - Instrument communication to synchronize clocks and ensure causality (All communication is visible.)
  - Extend Durbhakula et al [HPCA ‘99]

Workloads

- External Memory Algorithms
  - Sort, matrix multiply
- File System
- Real-time Workloads
  - RT Visualization/Video-on-demand
- HECURA Panel Suggestions
  - Univ. Chicago FLASH
  - DARPA SSCA #3

Measures of Success

- Theoretical Advances
- Performance improvements on workloads of interest via
  - Better Algorithms
  - Better Systems
- Use of simulator in the community