

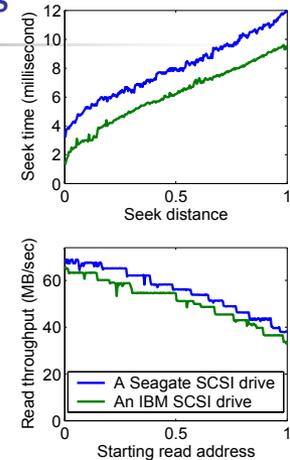
Storage Device Characterization

Kai Shen

Dept. of Computer Science, University of Rochester

Disk Characteristics

- Characterization of two SCSI disk drives
- Long seeks, but relatively fast sequential access
- Implication in system design:
 - infrequent bulk sequential access
 - put co-accessed data near each other
 -



2/12/2008

URCS 573 - Spring 2008

2

SCSI vs. ATA

- Beyond interface difference, mostly a business-oriented economical argument [Anderson et al. 2003]
 - some are willing to pay a lot more with a bit more reliability, controllability, and sometimes performance (typically at a loss of space capacity)
- Serial ATA
 - simplification: simpler cabling due to serial transfer
 - complication: allow multiple outstanding requests, possible support for command queuing

2/12/2008

URCS 573 - Spring 2008

3

Disks vs. Tapes

- Disks vs. tapes:
 - Performance characteristics: tapes are horrible for non-sequential access, and even for fine-grain sequential writes
 - Economical sense: tapes are much cheaper. Is it still true?
- Backup through disk replication

2/12/2008

URCS 573 - Spring 2008

4



Solid-State Disks

- Solid state: no (or almost no) mechanical moving parts
- RAM-based disk (with battery and backup disk)
 - Pro: latency bound by software processing (interrupts), throughput bound by bus bandwidth
 - Con: complex backup/restoration during power loss
 - Con: bulky (containing battery)
- Flash disks
 - Pro: truly solid state
 - Con: latency (~0.1ms) is longer than RAM-based disks
 - Con: throughput is less than even conventional disks
 - Con: writes are particularly slow due to redundancy coding
 - Con: writes wear out flash quickly



Five-Minute Rule

- Economical rule about whether to keep data in memory or on disk [Gray&Putzolu 1987]
 - one disk access per second costs about \$2000
 - 1KB memory costs \$5
 - breakeven economical point - 1KB data accessed once per 400sec
- Seems to still approximately hold for more recent disks [Gray&Graefe 1997]
- Absolutely not true for SSDs
 - cause for re-evaluating storage-related system designs (file system, databases, etc.) [Graefe 2007]