Benchmarking Tape Systems

Ethan Miller
University of Maryland Baltimore County
elm@umbc.edu
Theodore Johnson
AT&T Research
johnsont@research.att.com

Introduction

- Tape systems have a wide range of performance characteristics
  - Transfer rate
  - Seek time (short & long)
  - Rewind time
- Mass storage systems must understand tape performance to optimize transfers
- Benchmarks can supply useful data for models of storage

Overview

- Motivation
- Tape drive taxonomy
- Benchmarks used
- Results
- Implications for storage designers
- Conclusions

Motivation

- Tape drive performance important for:
  - Hierarchical storage systems
  - Tertiary storage in databases
- Detailed performance information crucial for:
  - Modeling storage systems
  - Optimizing access to removable media
- Many characteristics(i.e., seek time) non-obvious

Tape Drive Taxonomy

- Tape drive technology
  - Helical-scan
  - Linear / serpentine
- Tape packaging
- Directory
- Partitioning
- Block size

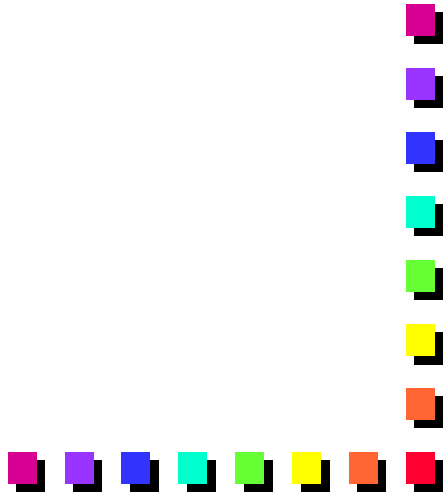

Serpentine vs. Helical Scan

- Serpentine
  - Tracks run the length of the medium
  - Forward and reverse tracks
  - Similar to audio cassette
- Helical scan
  - Tracks run diagonally on medium
  - Forward tracks only
  - Similar to VHS VCR

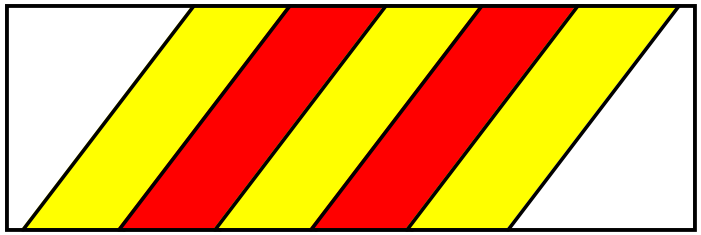

Sample Tape Drives

- IBM 3590
  - Serpentine, cartridge (directory at start of tape)
  - Variable block size, no partitions
  - High-speed: 8.9 MB/s transfer
- Ampex DST 310
  - Helical-scan, cassette (multiple landing zones)
  - Fixed block size, partitions for data management
  - High speed: 14.2 MB/s (large transfers)

Benchmarks

- Mount & unmount times
- Seek times
  - Measure time until a block is readable
  - Compare true seeks to reading intervening data
  - Use various starting points to provide better model
- Transfer rate

Seek Benchmarks

- Long seek from start of tape
- Long seek from middle of tape
  - Pick representative starting locations
  - Find unusual seek time behavior
- Short seek from middle of tape
  - Give seek command; follow by reading a block
  - Compare to simply reading intervening blocks

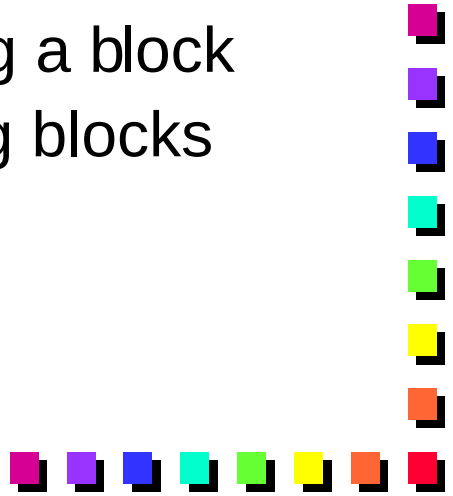

Transfer Rates

- As expected, less expensive drives were slower
- No clear winner between helical-scan and serpentine

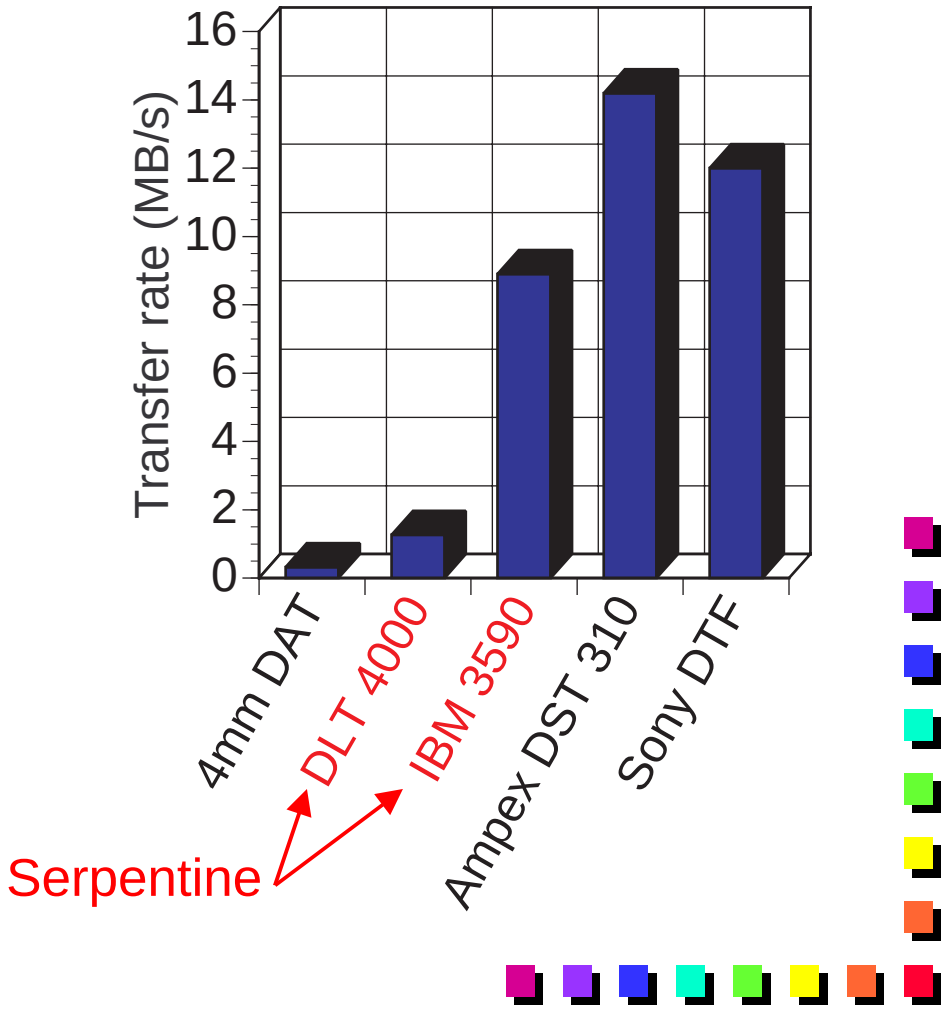

Seek & Rewind (Helical)

- Seek time varies linearly with destination
  - Seek and rewind have similar cost functions
  - No variation from linear delay
- Seek & rewind overhead
  - Relatively large
  - Similar for seek & rewind

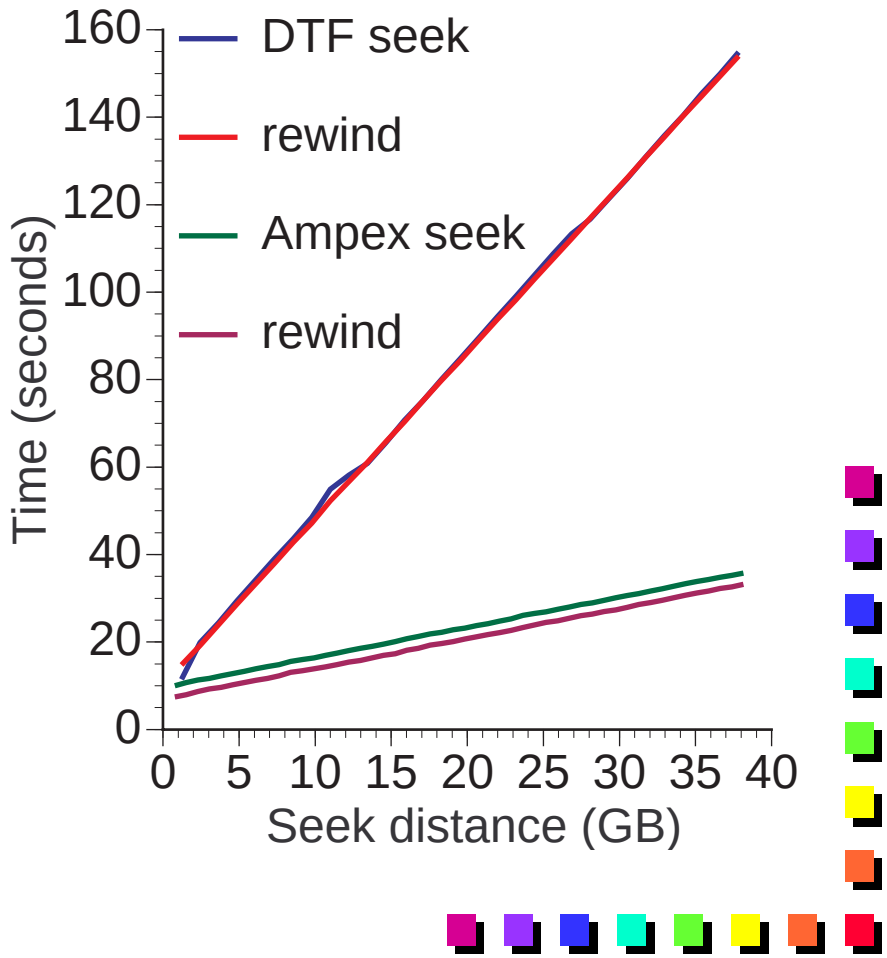

Seek & Rewind (Serpentine)

- Seek time is complex stepped function
  - Two speeds: fast seek & read
  - End-of-track behavior may vary by track
  - Choosing optimal file placement is non-trivial
- Rewind time is linear with distance
- Pattern repeats for other tracks

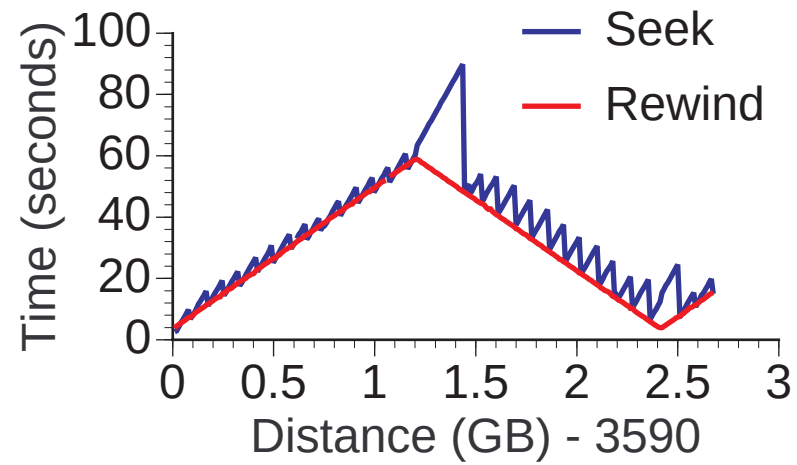

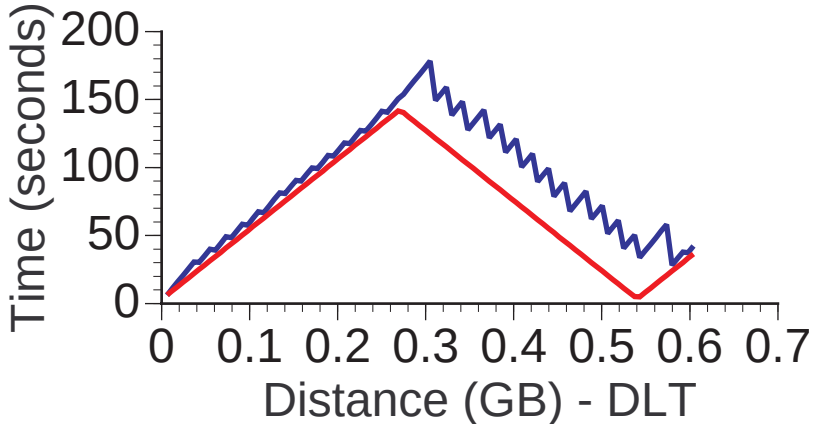

Short Seeks vs. Reads (Helical)

- Seek times
  - Linear time
  - Large fixed overhead
- Read times
  - More complex time function
  - Small fixed overhead
- Results
  - Avoid short seeks (< 100 MB)
  - Instead, read intervening blocks

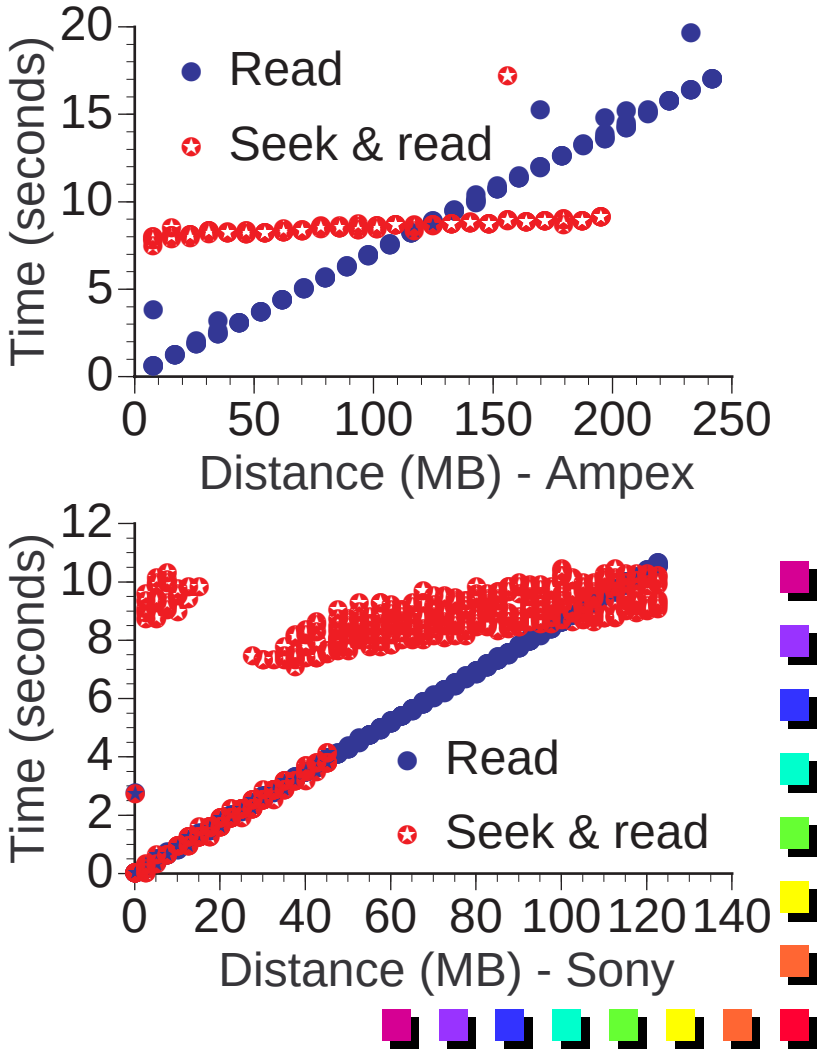

Short Seeks vs. Reads (Serp.)

- Low overhead seeks
  - No difference between short seeks & short reads
  - Short seeks take almost no time
- Seeks never worse than reads
- Always use seeks - no need to make a choice

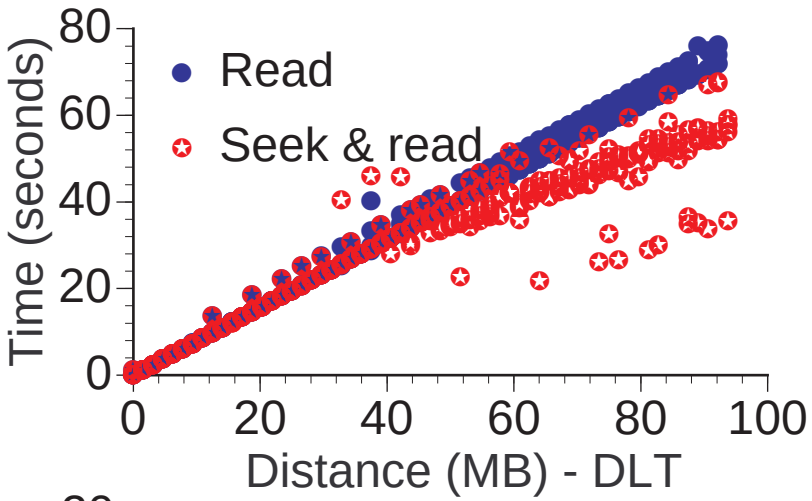

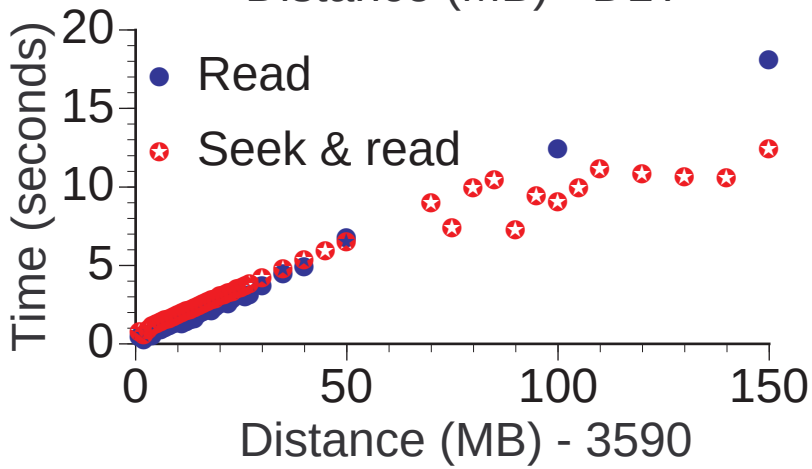

Implications for Mass Storage

- Serpentine tape drives
  - Lower seek overhead
  - Better at short seeks
  - Allow optimizations by choosing file position on individual tracks
- Helical scan tape drives
  - Simpler performance model
  - Logical block numbering reflects true seek time between locations

Suggestions for Storage Systems

- Serpentine tape drives
  - Place large files at the start of a track
    - Reduces seek time to file start
    - Reduces response time & drive utilization
    - Wastes too much space?
  - Reorder reads by seek time, not logical block number on tape
- Helical scan tape drives
  - Use reads rather than short seeks
  - Incorporate this knowledge into tape system?

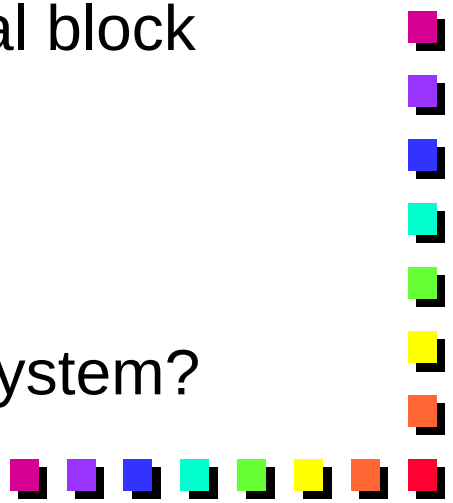

Lessons Learned

- Tape drive performance can have quirks
  - 3590 had firmware bug (since fixed) that affected performance at end-of-track
  - Some tapes report “ready” immediately and tack delay onto following request
  - Integrate knowledge of quirks into storage system
- Make sure performance “improvements” actually do so
- Tape performance can be complex, so repeat measurements several times

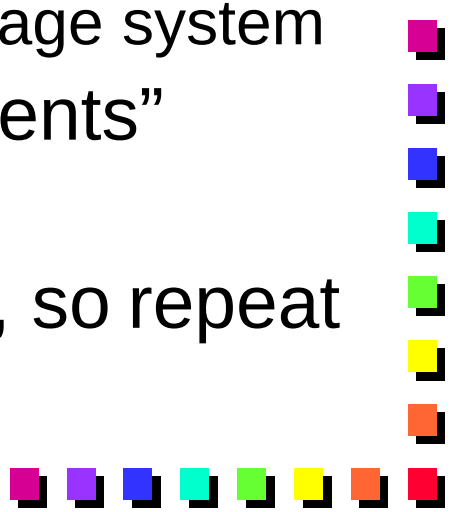

Future Work

- Benchmark new tape drives
- Create standard tape benchmarks?
  - Allow users to run them
  - Provide to vendors so they can supply more detailed performance information
- Build parameterized models for tape drive performance
  - Use in mass storage systems
  - Use in databases with tertiary storage

Conclusions

- Tape performance is more complex than it would appear at first glance
- Mass storage systems can use knowledge of performance model to improve performance
- Simple benchmarks can provide detailed information about tape performance