Comparing Performance of Solid State Devices and Mechanical Disks

Jiri Simsa

Milo Polte, Garth Gibson

PARALLEL DATA LABORATORY

Carnegie Mellon University

Motivation

- Performance gap [Pugh71] technology
  - Spans 6 orders of magnitude
  - Right mix of price, capacity, performance, ...
  - Cost-effectiveness - $ per performance

- Flash for consumer and enterprise storage
  - Flash believed to be the gap technology
  - What is Flash good at?

    - small random I/O

How well does Flash do?

- IOZone, high-end consumer devices

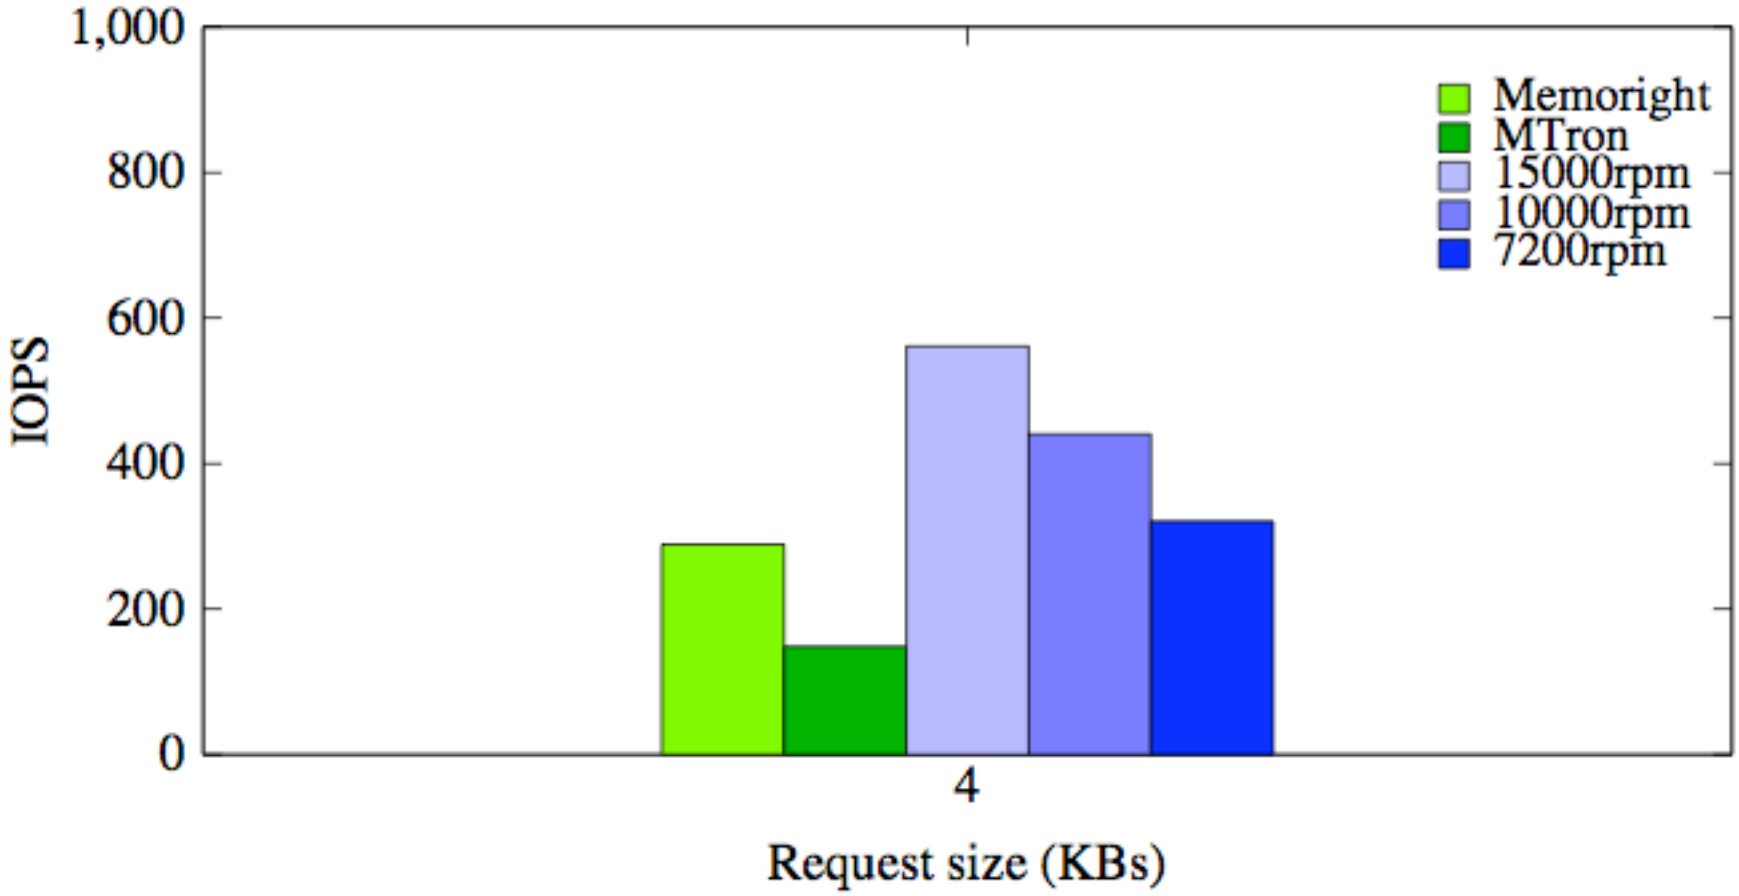

Outline

- Performance measurements
  - Experimental setup
  - High-end consumer devices
  - Enterprise devices
  - Interpreting the results
- Hybrid storage research questions

Experimental Setup

- Machines:
  - Intel CPUs, 1 - 4 GB DRAM, Linux 2.6
  - SATA, SCSI, and PCIe interconnects
  - Range of disks and Flash based Solid State Devices

- Workload:
  - single-threaded IOZone test
  - varying request sizes and access patterns
  - 4 GB file through an ext3 file system
  - OS and device write buffers on

Measured Solid State Devices

'08 high-end consumer and enterprise devices:

- Mtron Mobi 16GB @ $370
- Memoright GT 16GB @ $510
- Intel X25-M 80GB @ $730
- Intel X25-E 32GB @ $810
- FusionIO ioDrive 80GB @ $2400

Compared to 7200rpm, 10k rpm, and 15k rpm Seagate magnetic disks.

Sequential Reads

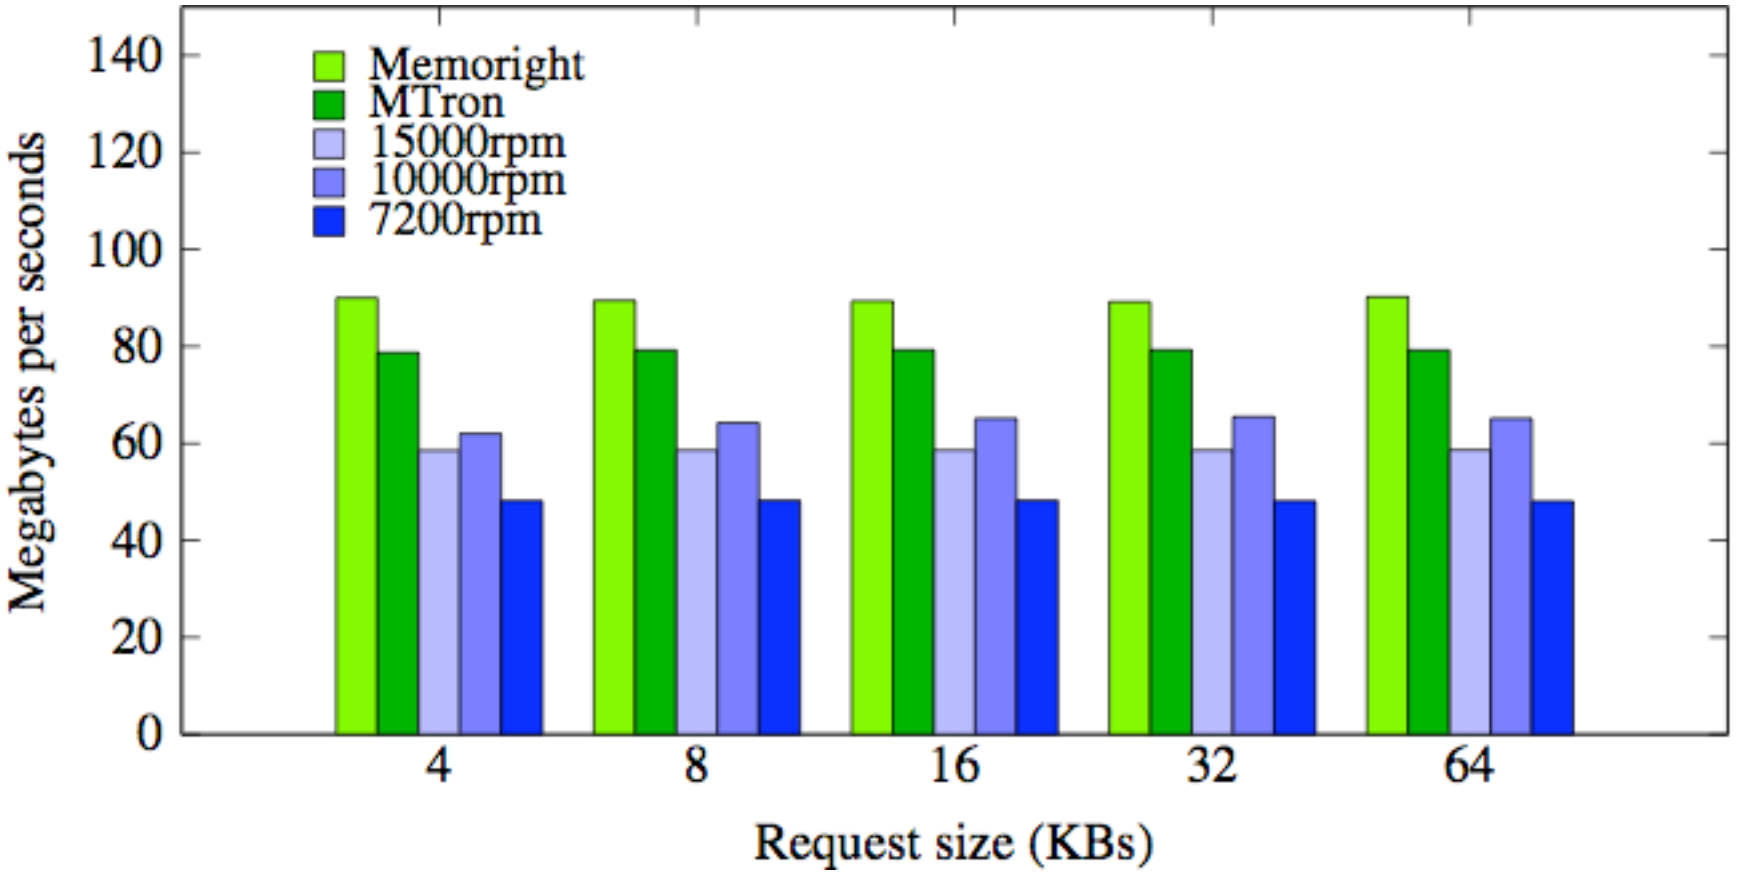

Sequential Reads

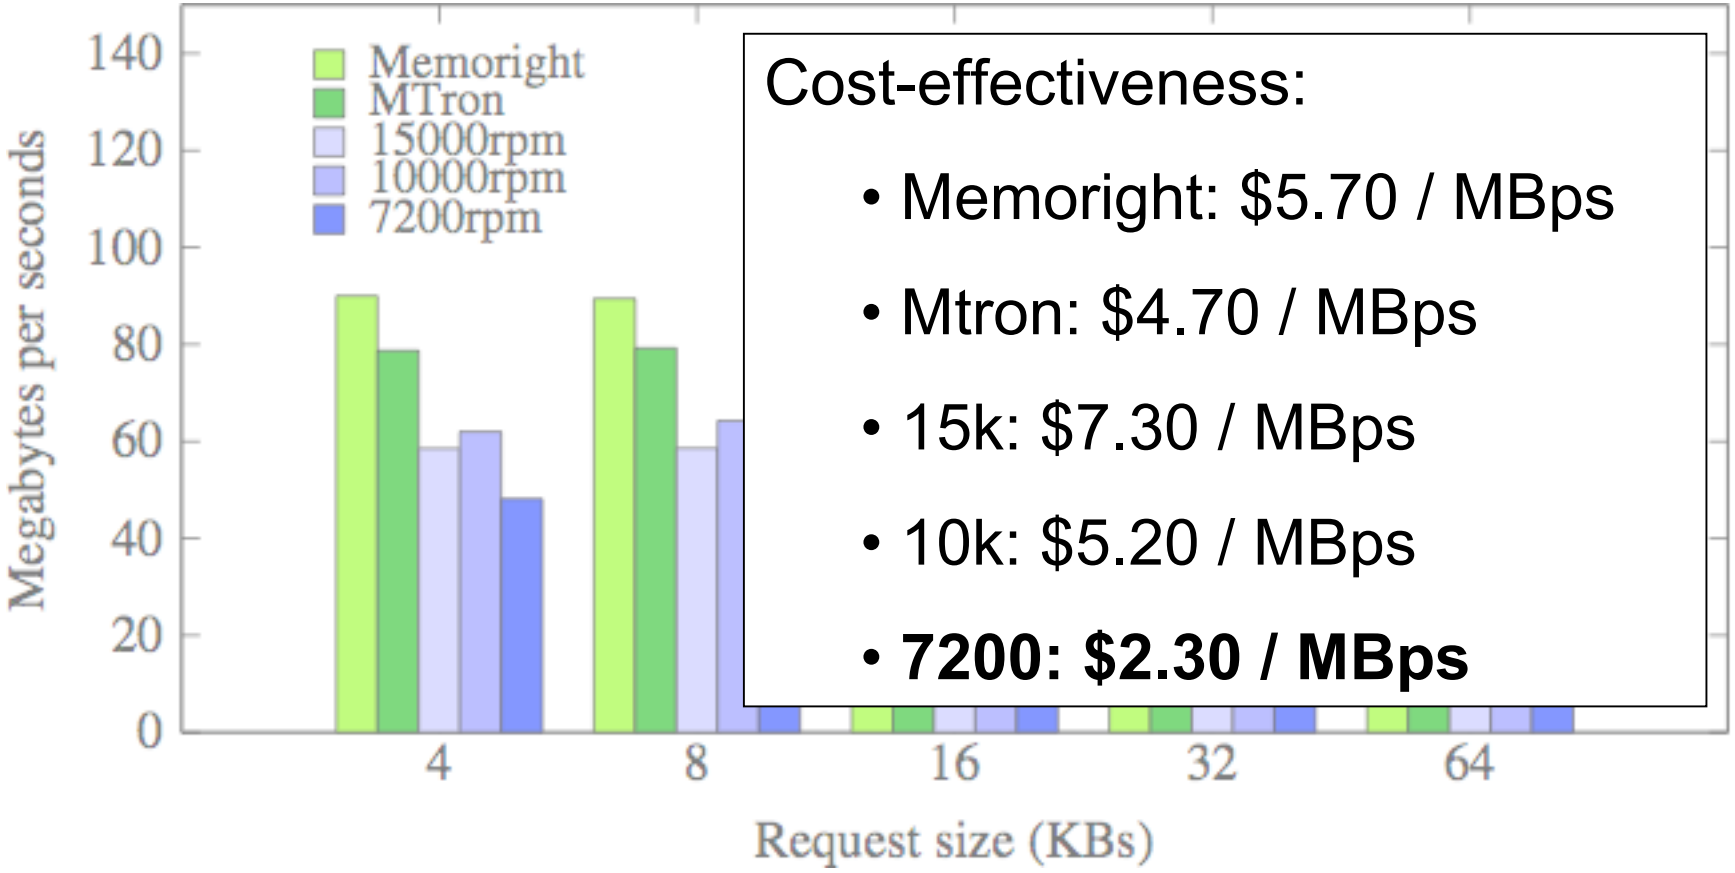

Cost-effectiveness:

- Memoright: $5.70 / MBps
- Mtron: $4.70 / MBps
- 15k: $7.30 / MBps
- 10k: $5.20 / MBps
- **7200: $2.30 / MBps**

Sequential Writes

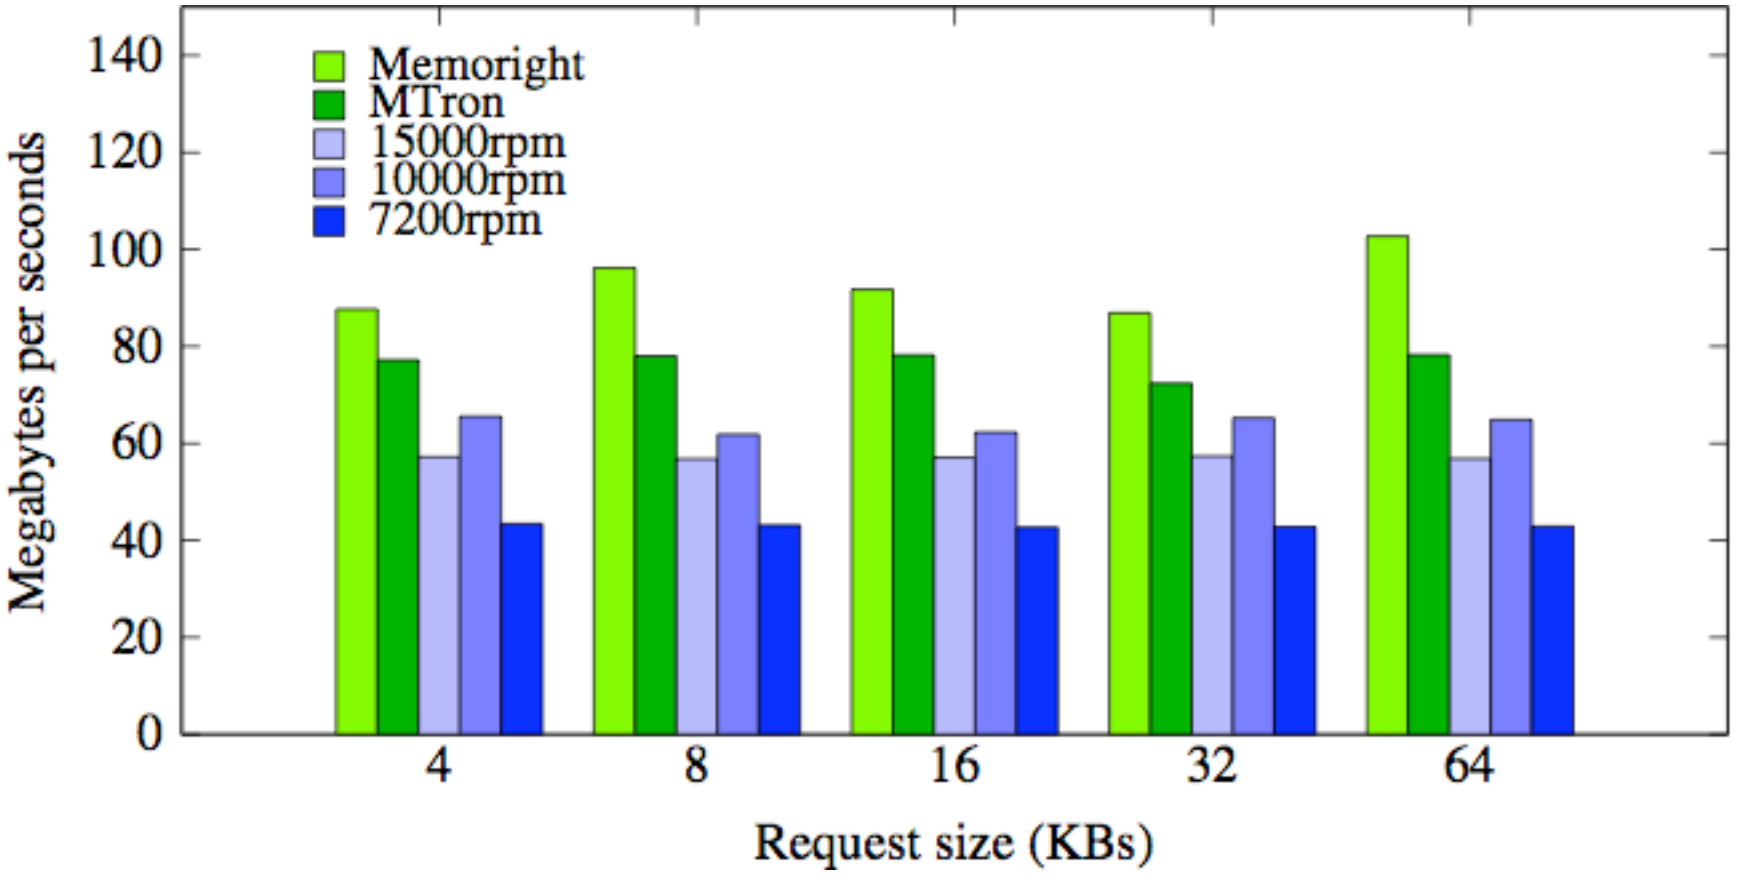

Sequential Writes

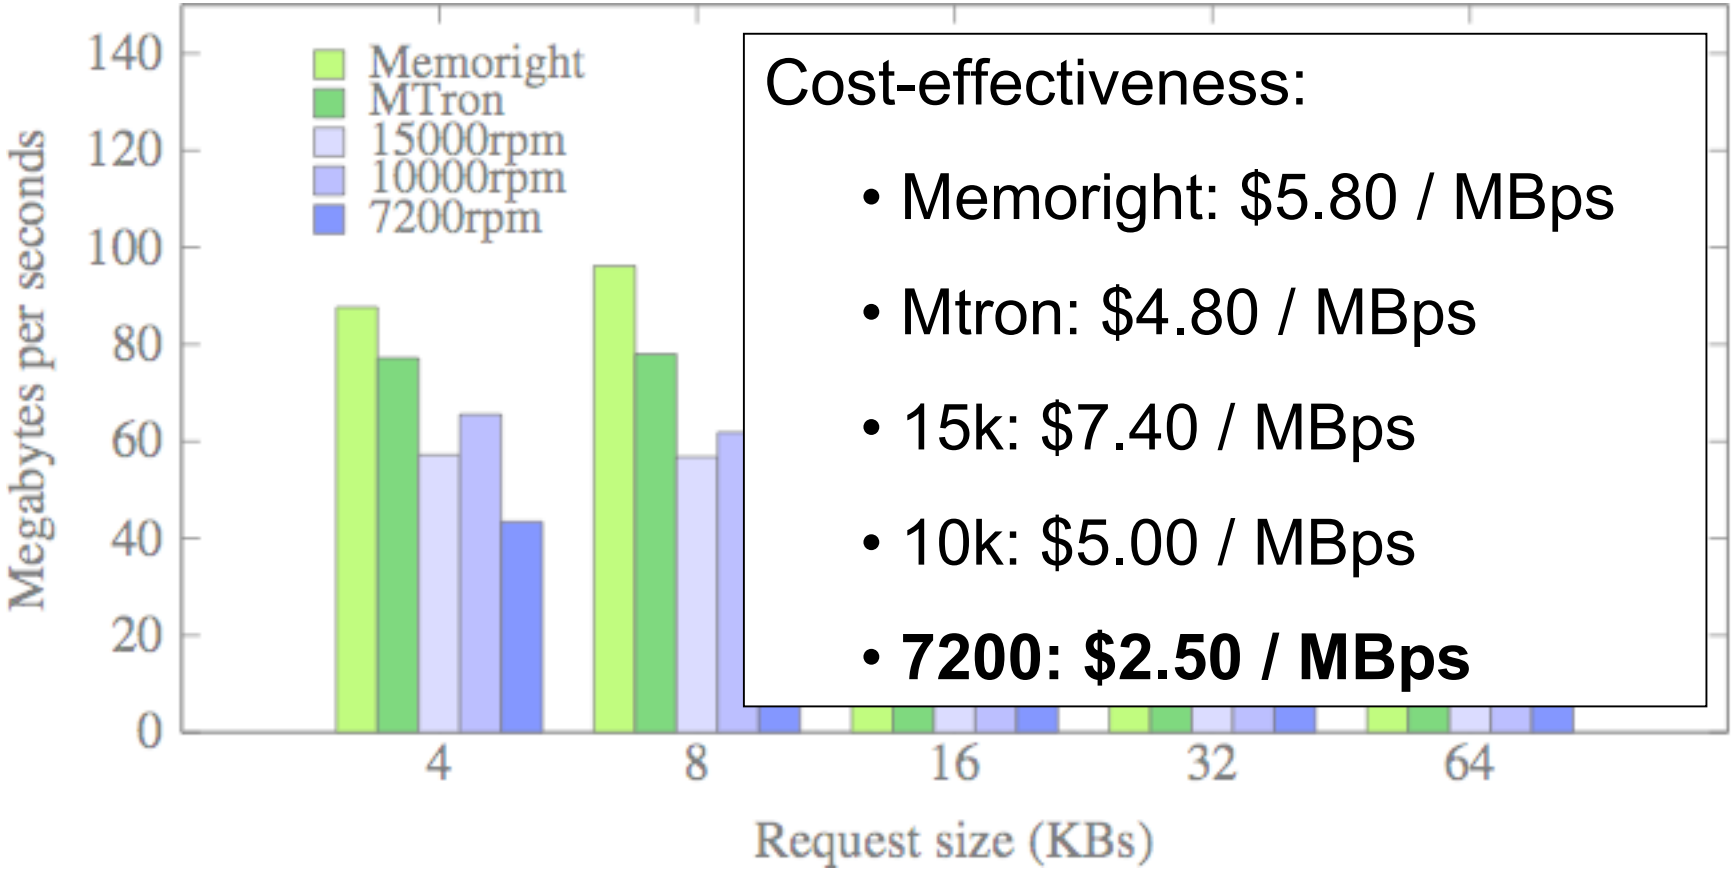

Cost-effectiveness:

- Memoright: $5.80 / MBps
- Mtron: $4.80 / MBps
- 15k: $7.40 / MBps
- 10k: $5.00 / MBps
- **7200: $2.50 / MBps**

Random Reads

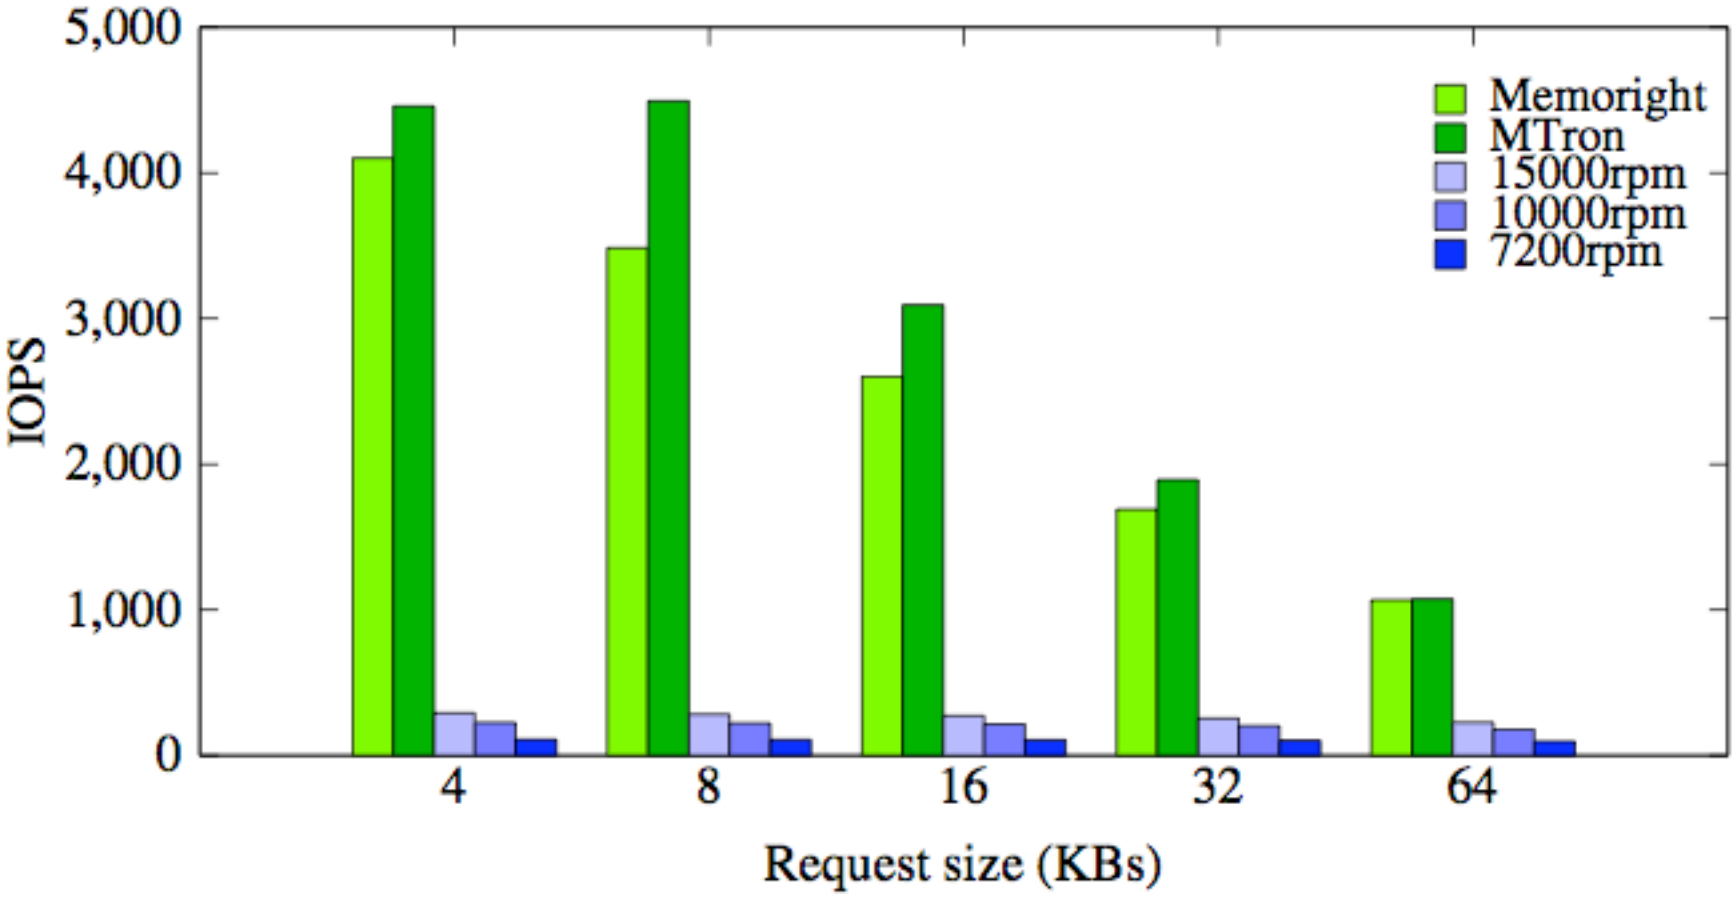

Random Reads

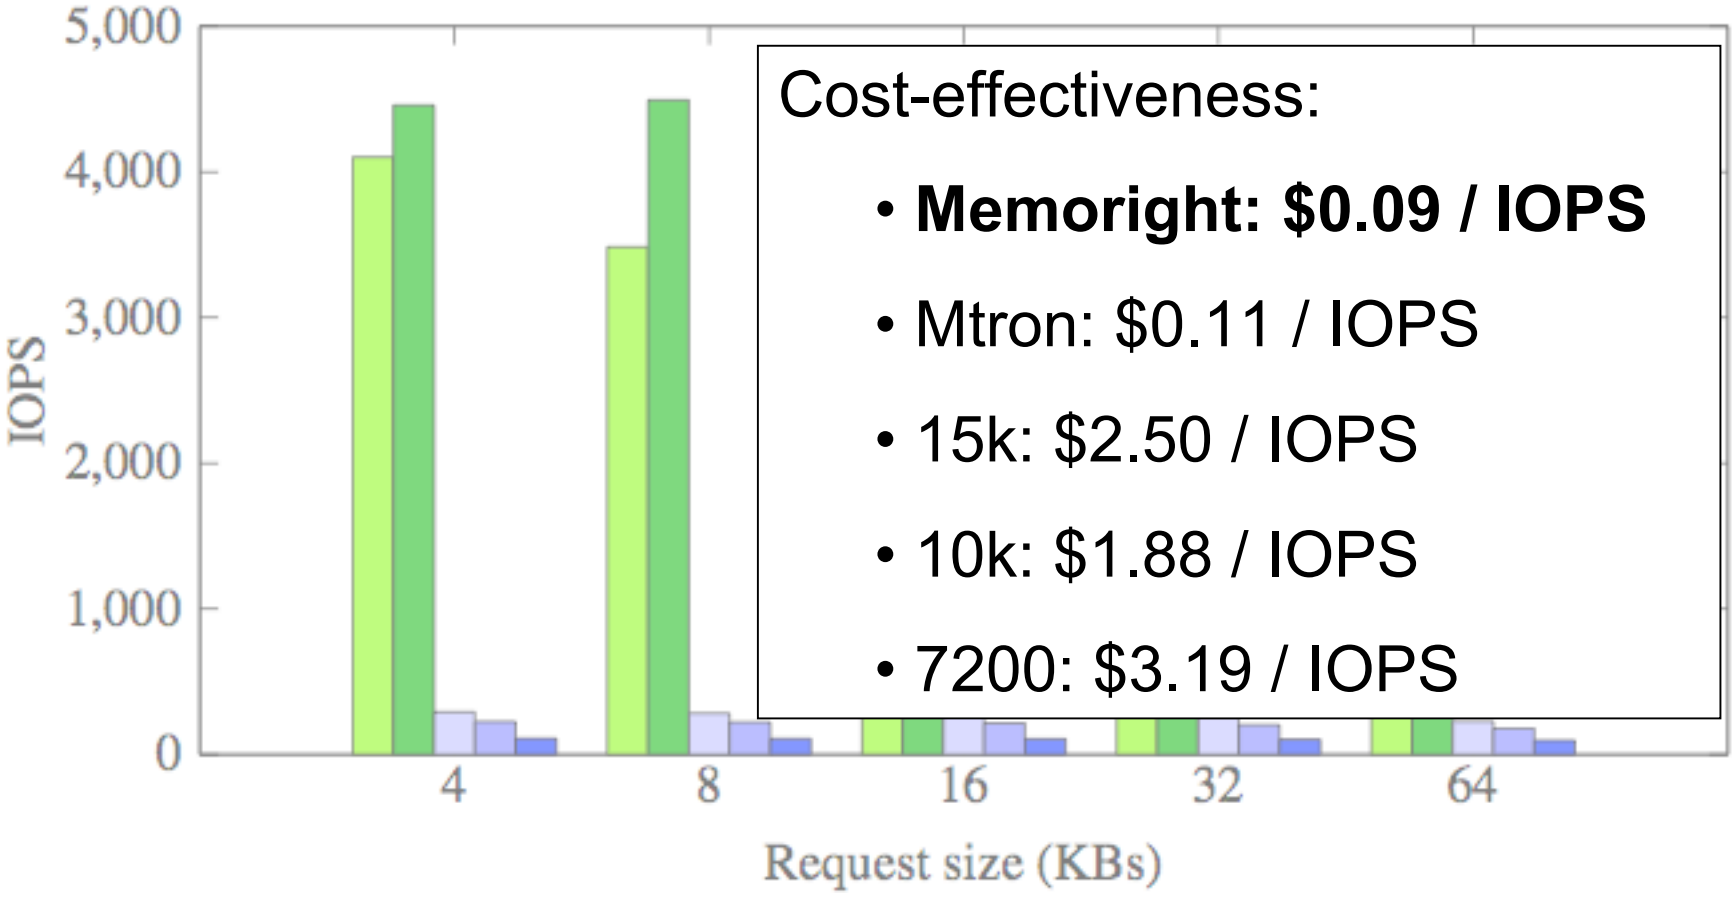

Cost-effectiveness:

- **Memoright: $0.09 / IOPS**
- Mtron: $0.11 / IOPS
- 15k: $2.50 / IOPS
- 10k: $1.88 / IOPS
- 7200: $3.19 / IOPS

Random Writes

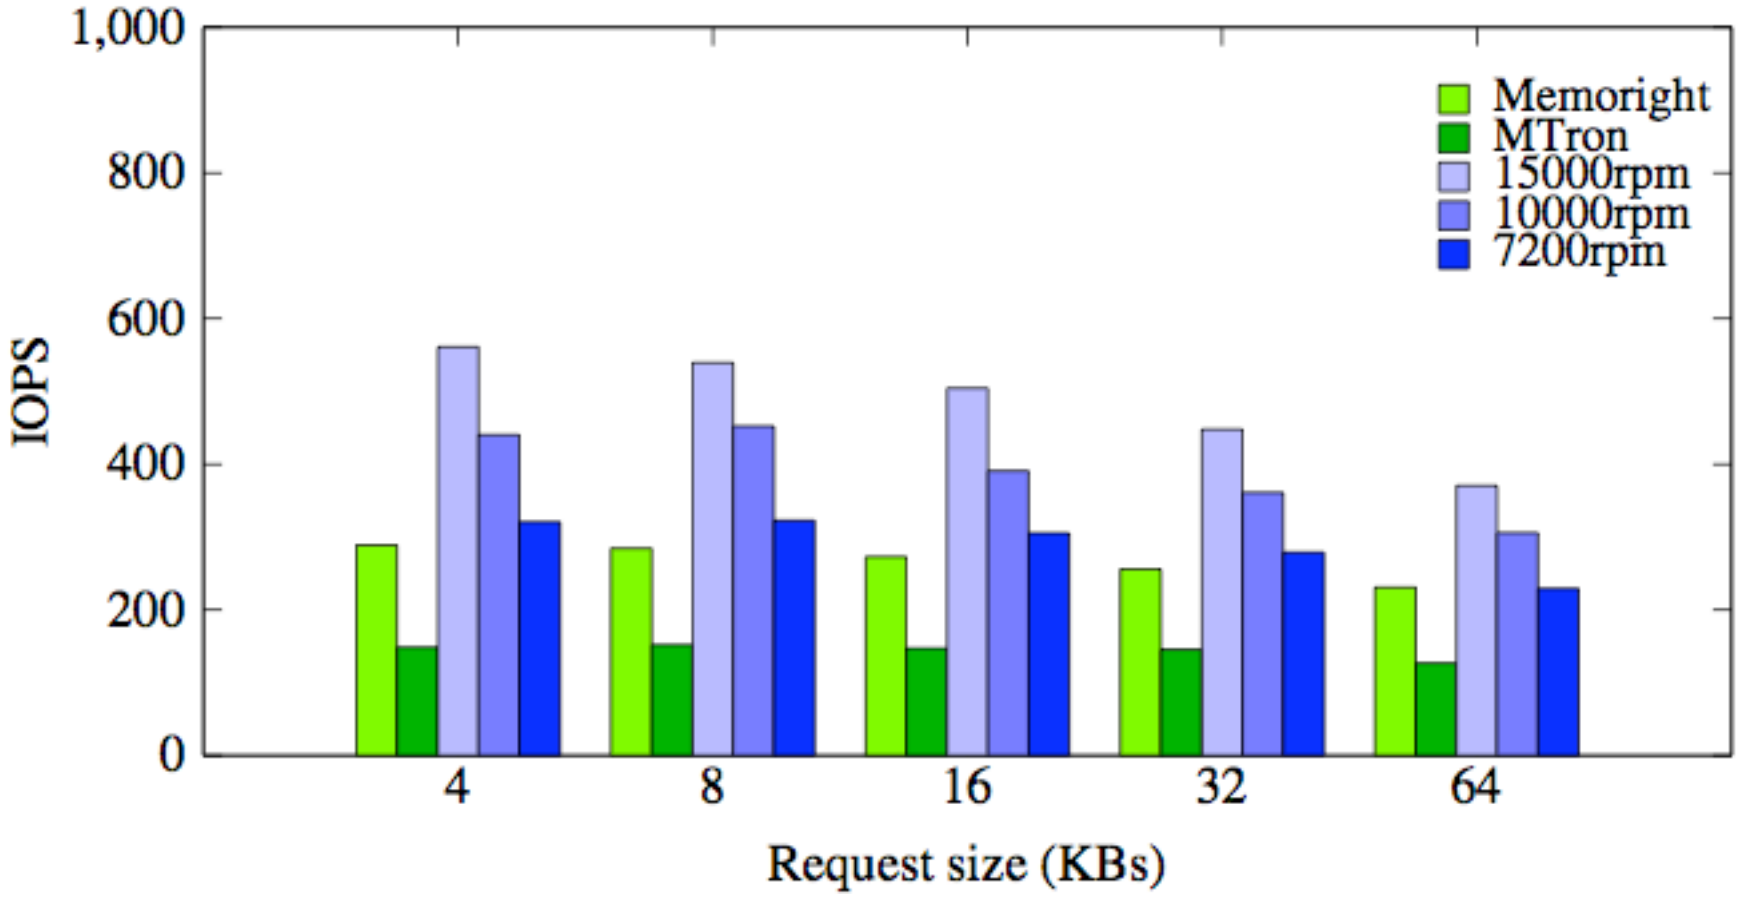

Random Writes

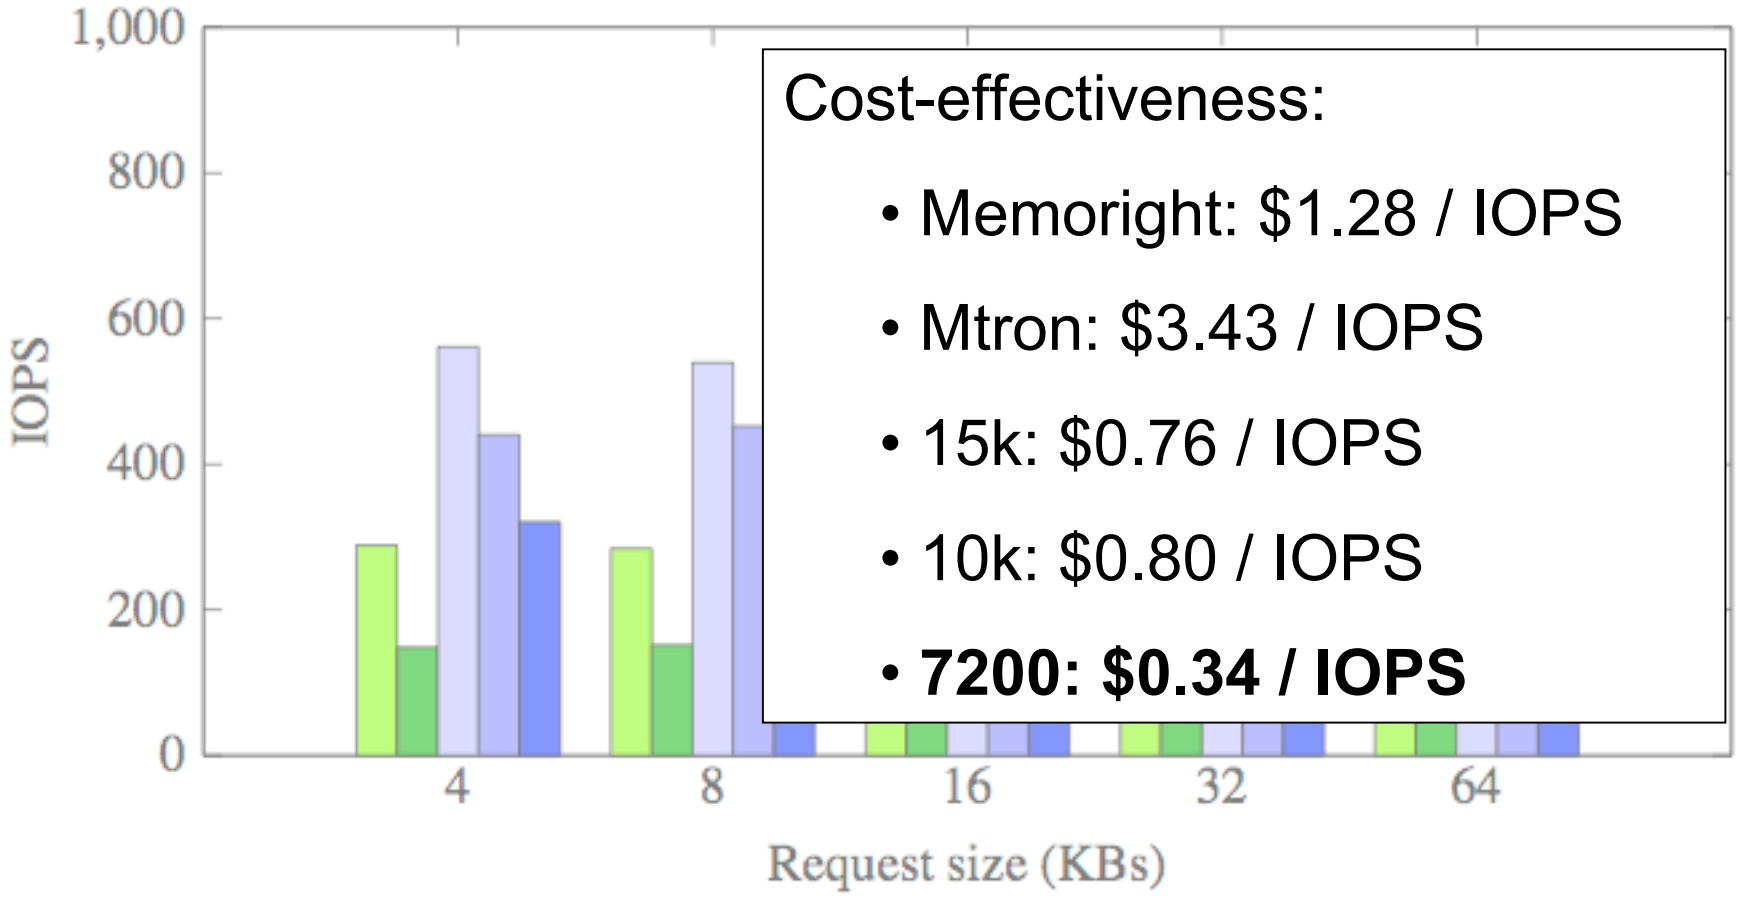

Cost-effectiveness:

- Memoright: $1.28 / IOPS
- Mtron: $3.43 / IOPS
- 15k: $0.76 / IOPS
- 10k: $0.80 / IOPS
- **7200: $0.34 / IOPS**

Interpretation

- For the measured devices:
  - Efficient log-structured I/O layer will improve relative SSD performance
  - Except for random reads, disk achieve
    - comparable price per performance
    - much better price per capacity
  - Target uses for consumer Flash:
    - High read to write ratio
    - Random access pattern
    - Small capacity

Measured Solid State Devices

‘08 high-end consumer and enterprise devices:

- Mtron Mobi 16GB @ $370
- Memoright GT 16GB @ $510
- Intel X25-M 80GB @ $730
- Intel X25-E 32GB @ $810
- FusionIO ioDrive 80GB @ $2400

Compared to 7200rpm, 10k rpm, and 15k rpm Seagate magnetic disks.

Improved Sequential Read Bandwidth

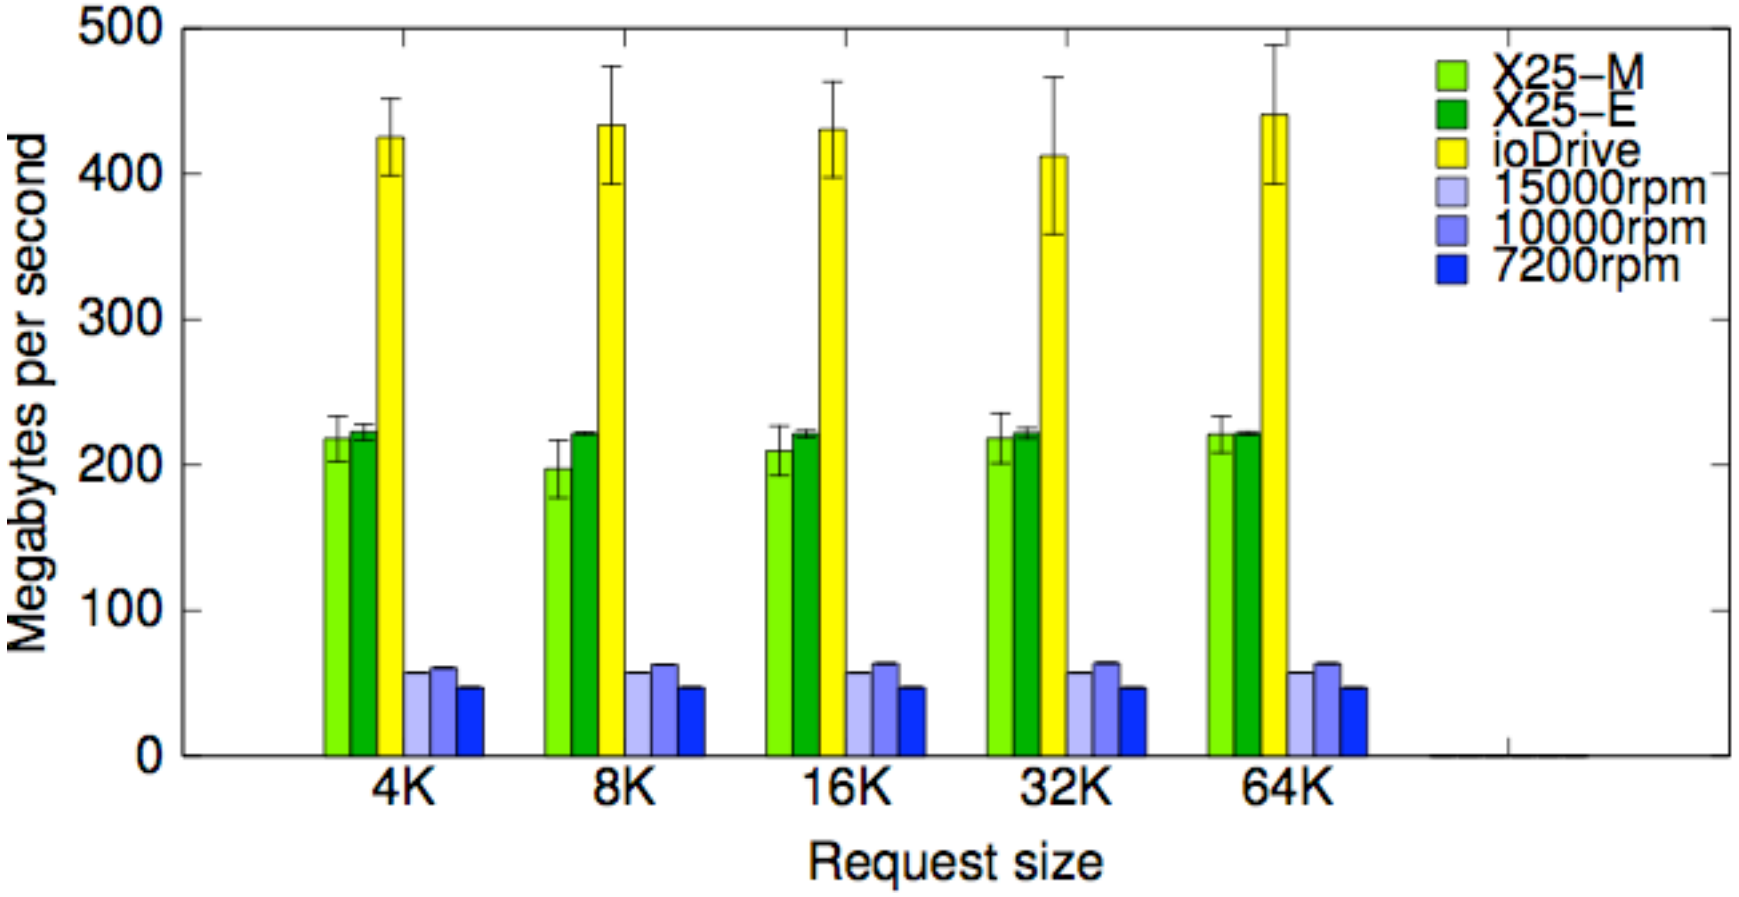

Improved Sequential Read Bandwidth

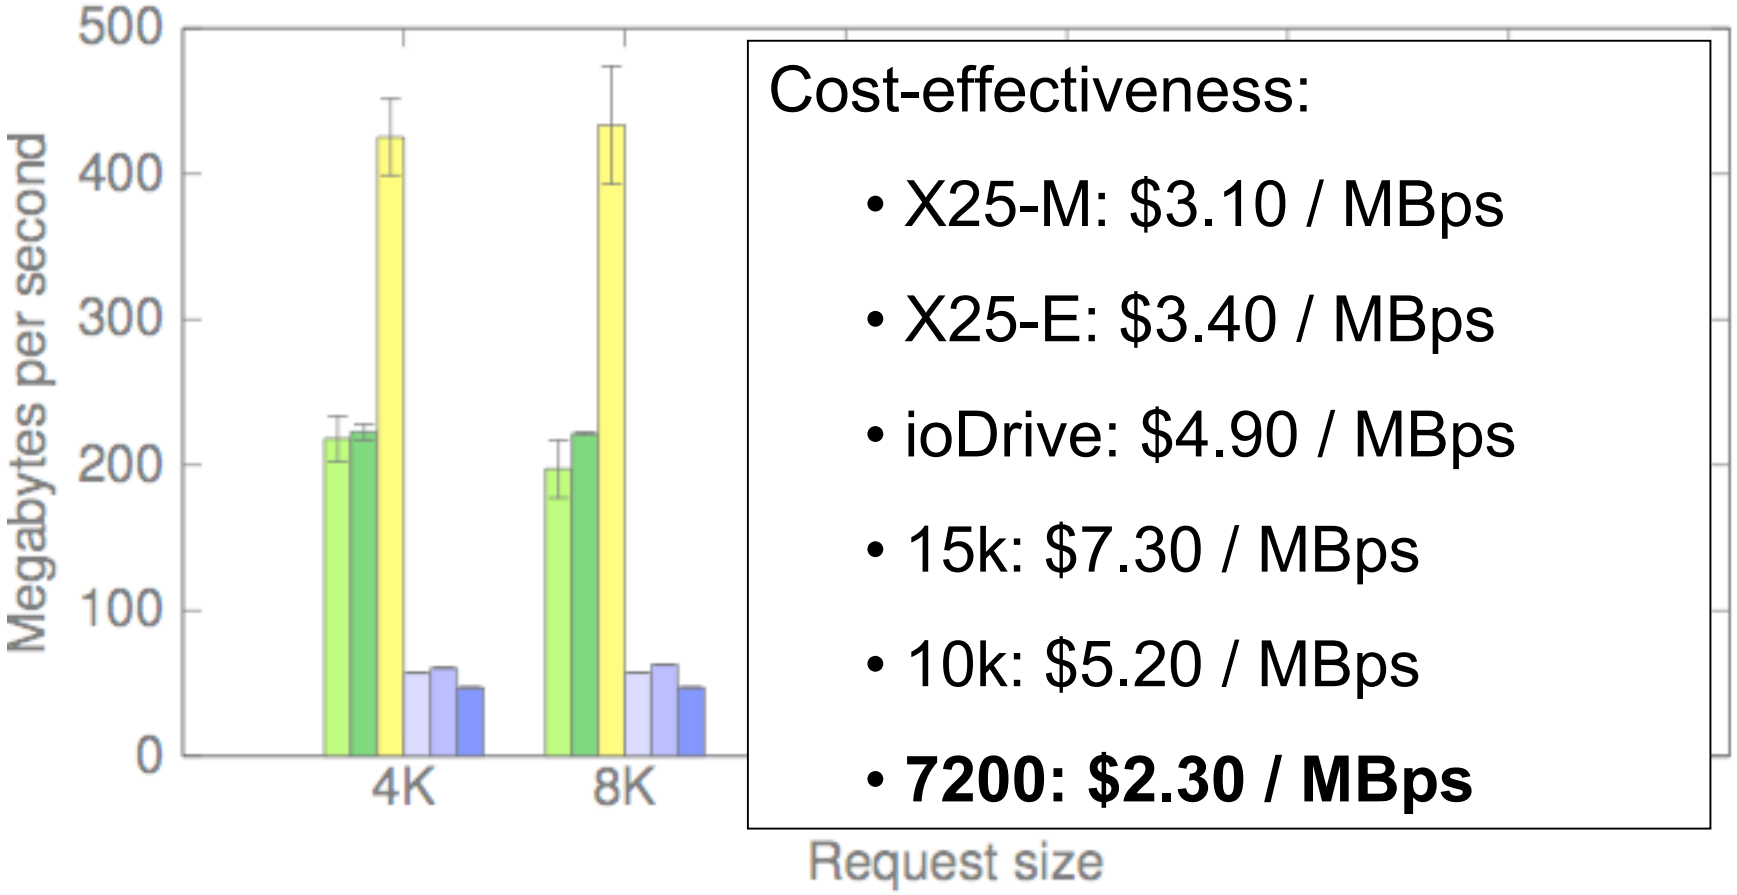

Cost-effectiveness:

- X25-M: $3.10 / MBps
- X25-E: $3.40 / MBps
- ioDrive: $4.90 / MBps
- 15k: $7.30 / MBps
- 10k: $5.20 / MBps
- **7200: $2.30 / MBps**

Improved Sequential Write Bandwidth

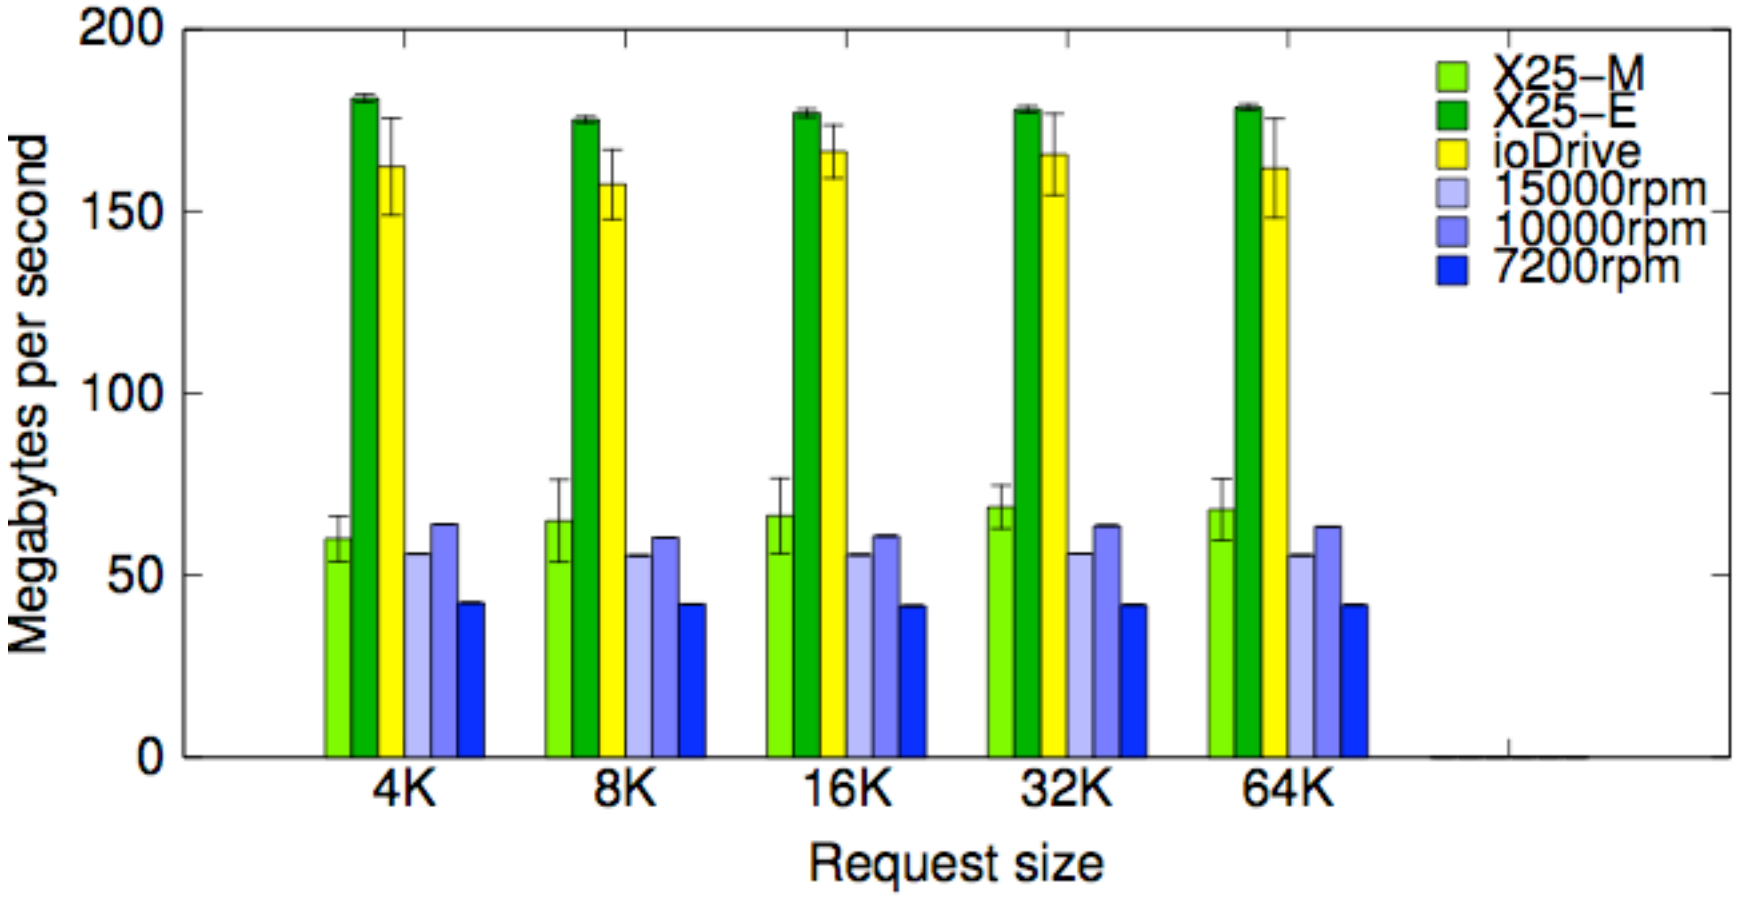

Improved Sequential Write Bandwidth

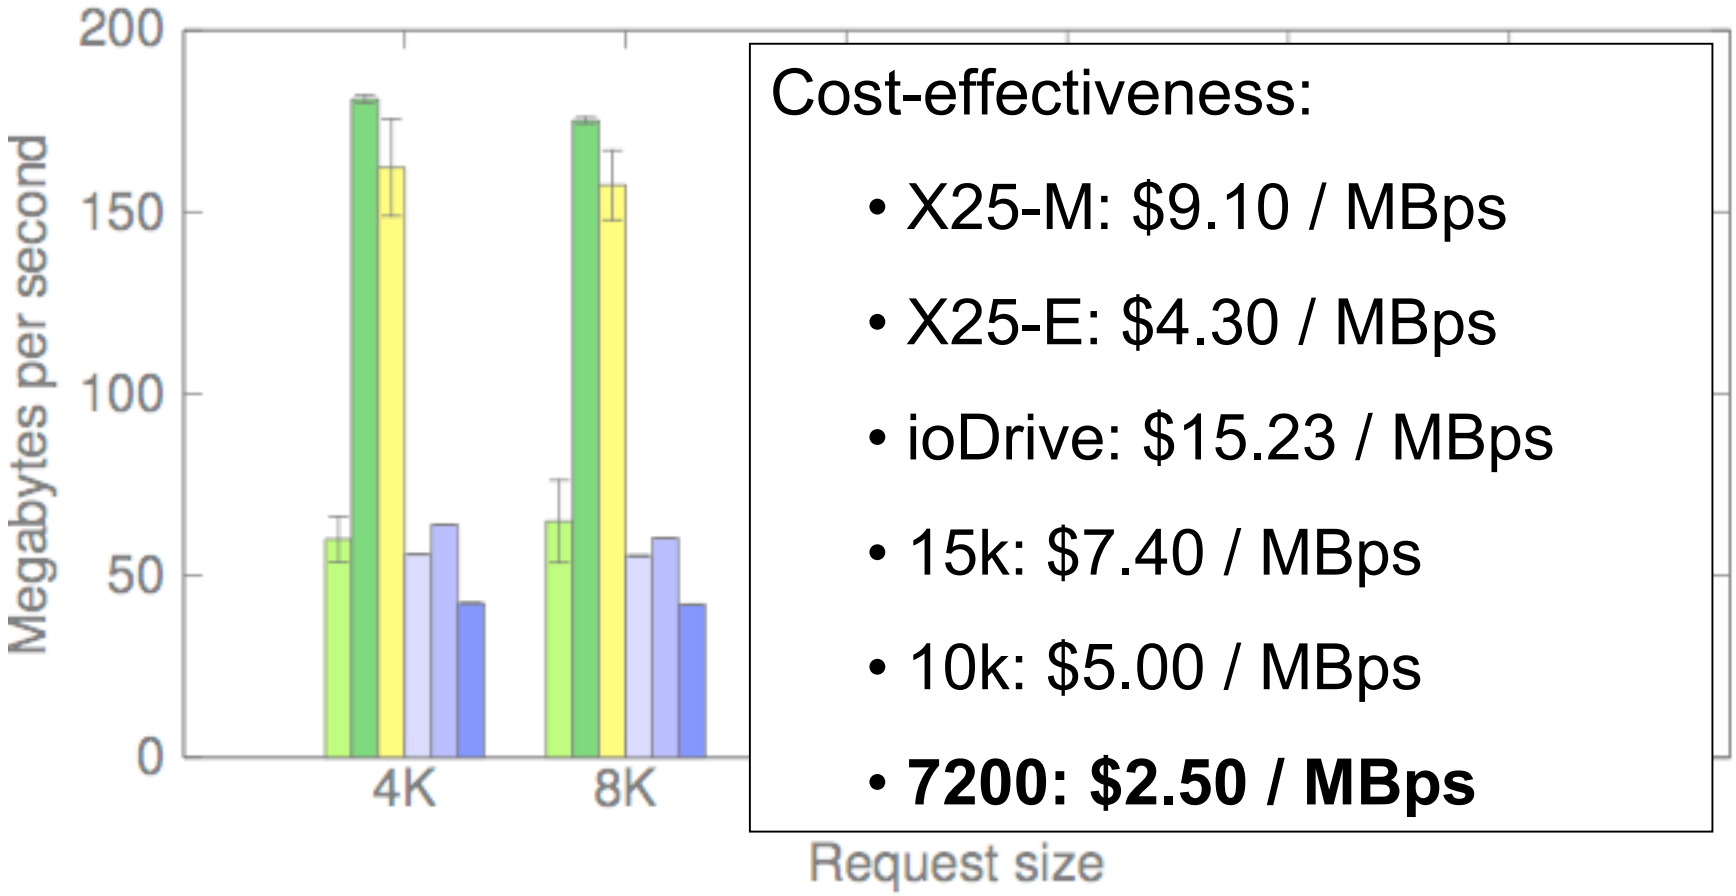

Cost-effectiveness:

- X25-M: $9.10 / MBps
- X25-E: $4.30 / MBps
- ioDrive: $15.23 / MBps
- 15k: $7.40 / MBps
- 10k: $5.00 / MBps
- **7200: $2.50 / MBps**

Great Random Read Performance

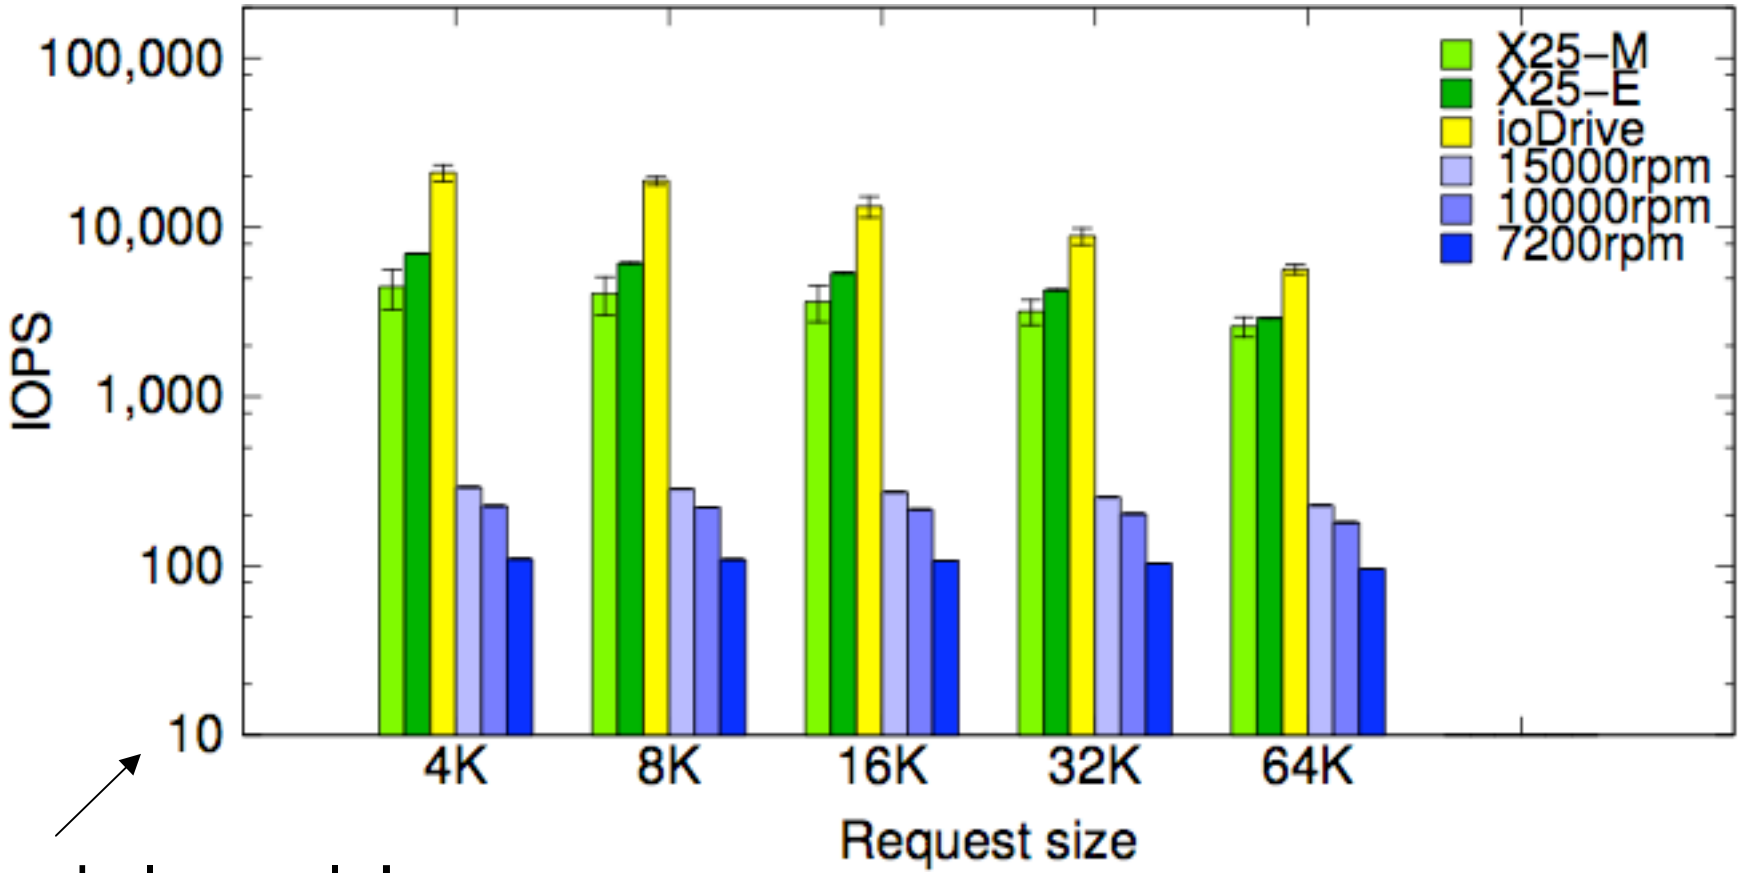

Look, logscale!

Great Random Read Performance

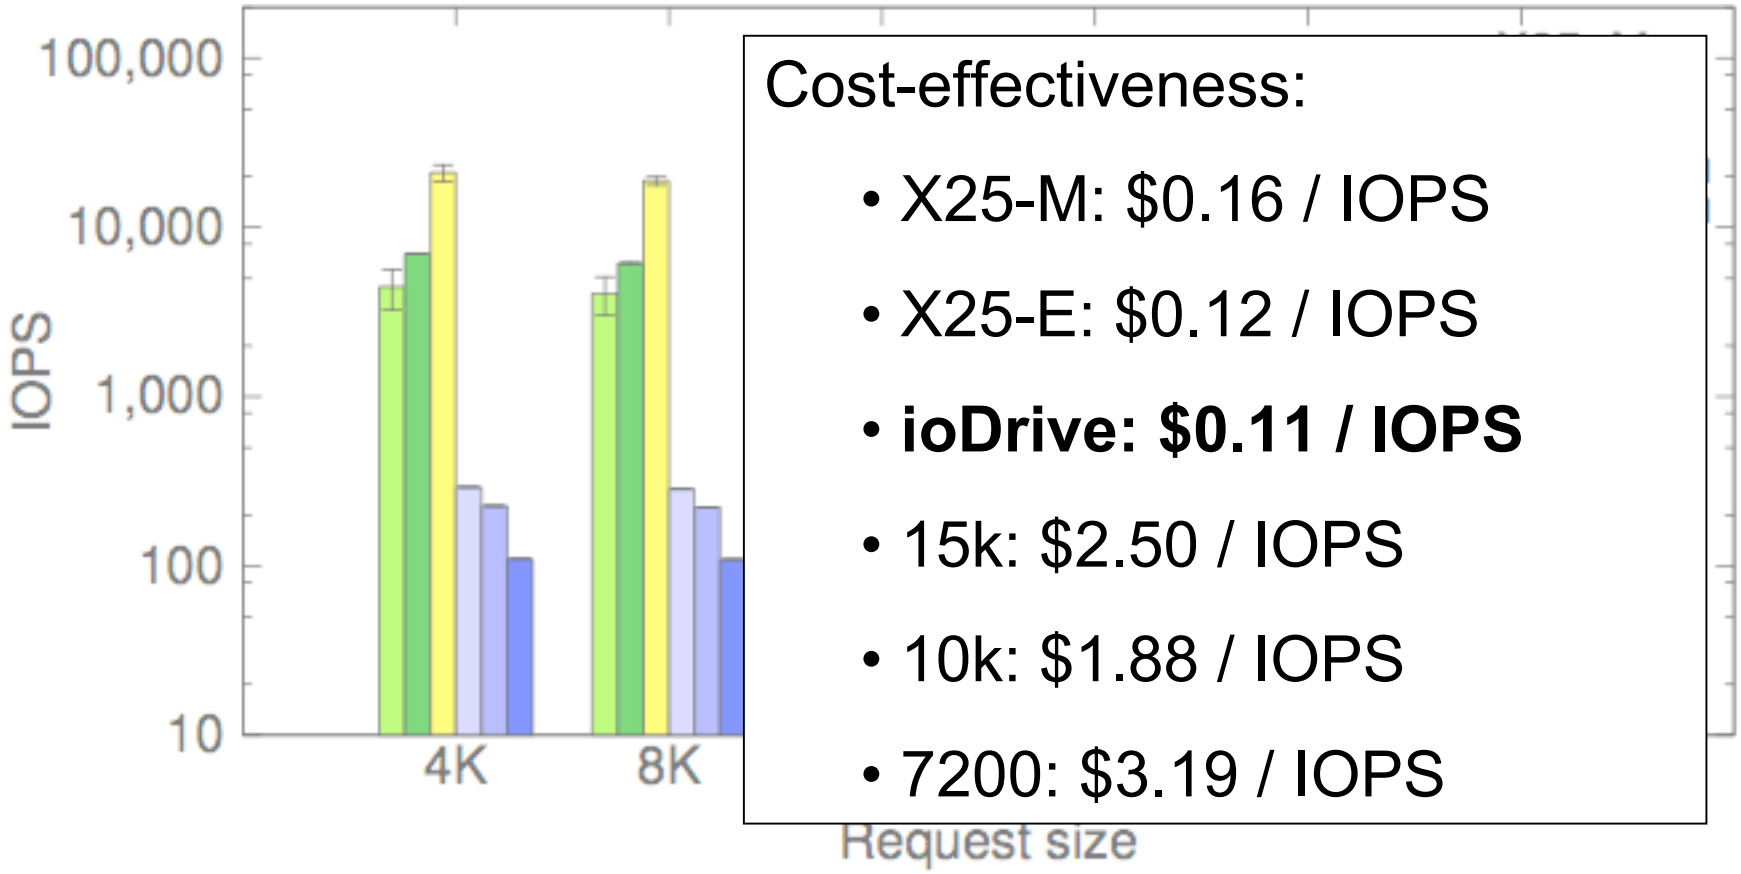

Cost-effectiveness:

- X25-M: $0.16 / IOPS
- X25-E: $0.12 / IOPS
- **ioDrive: $0.11 / IOPS**
- 15k: $2.50 / IOPS
- 10k: $1.88 / IOPS
- 7200: $3.19 / IOPS

Exceptional Random Write Performance

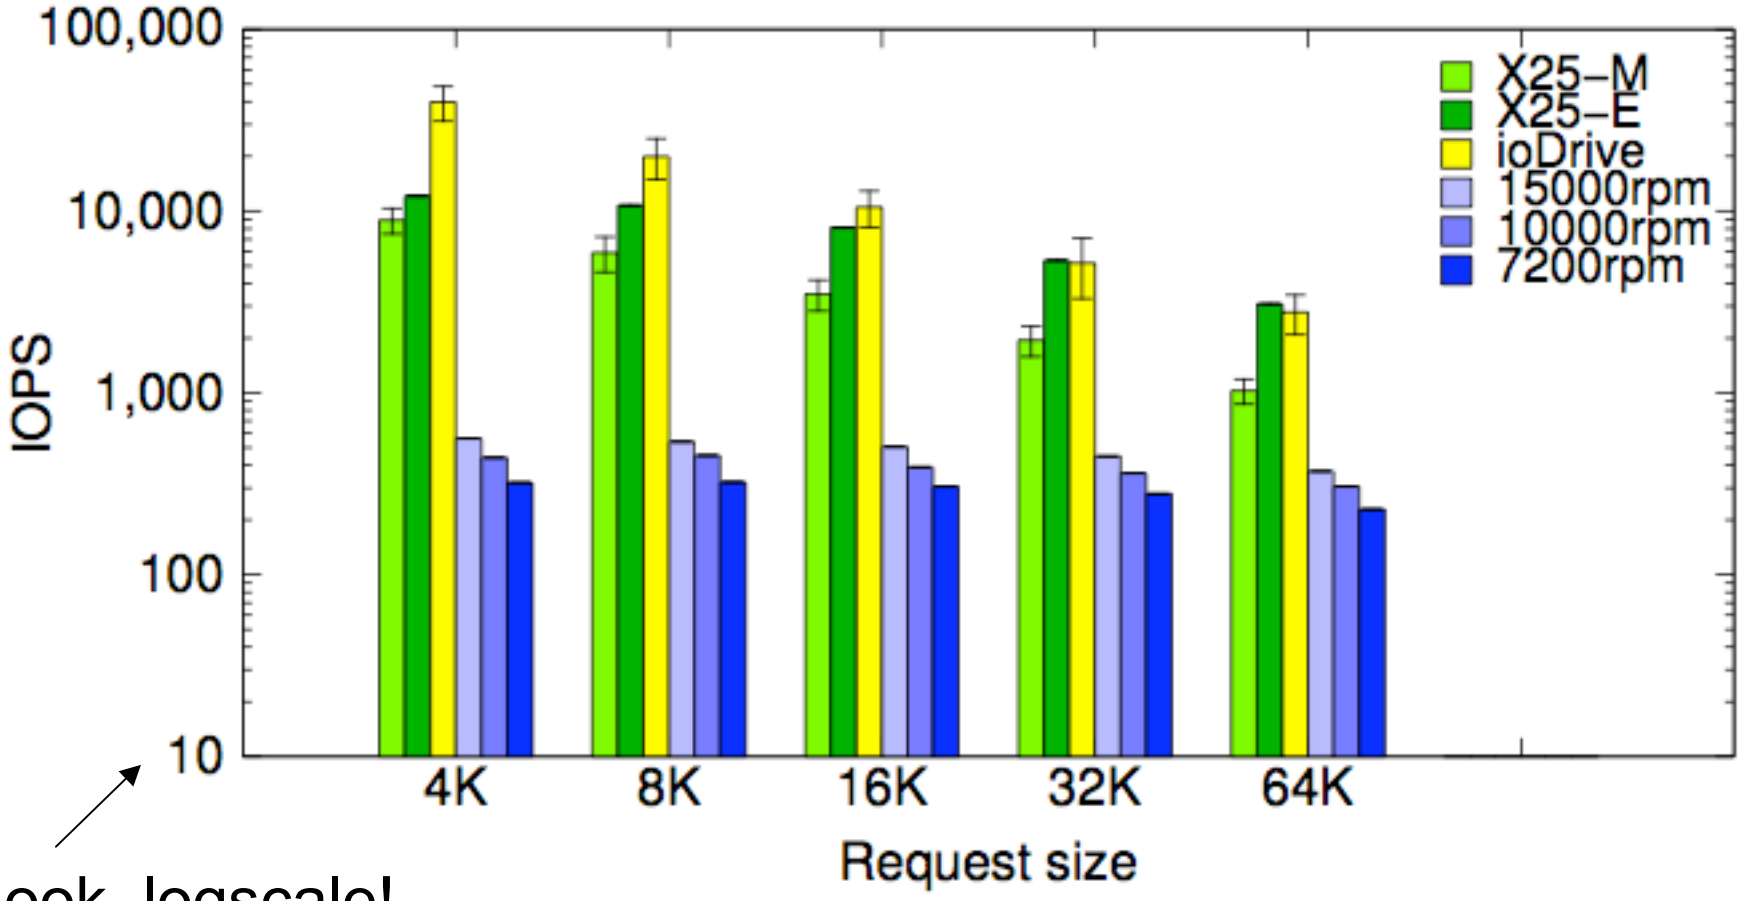

Look, logscale!

Exceptional Random Write Performance

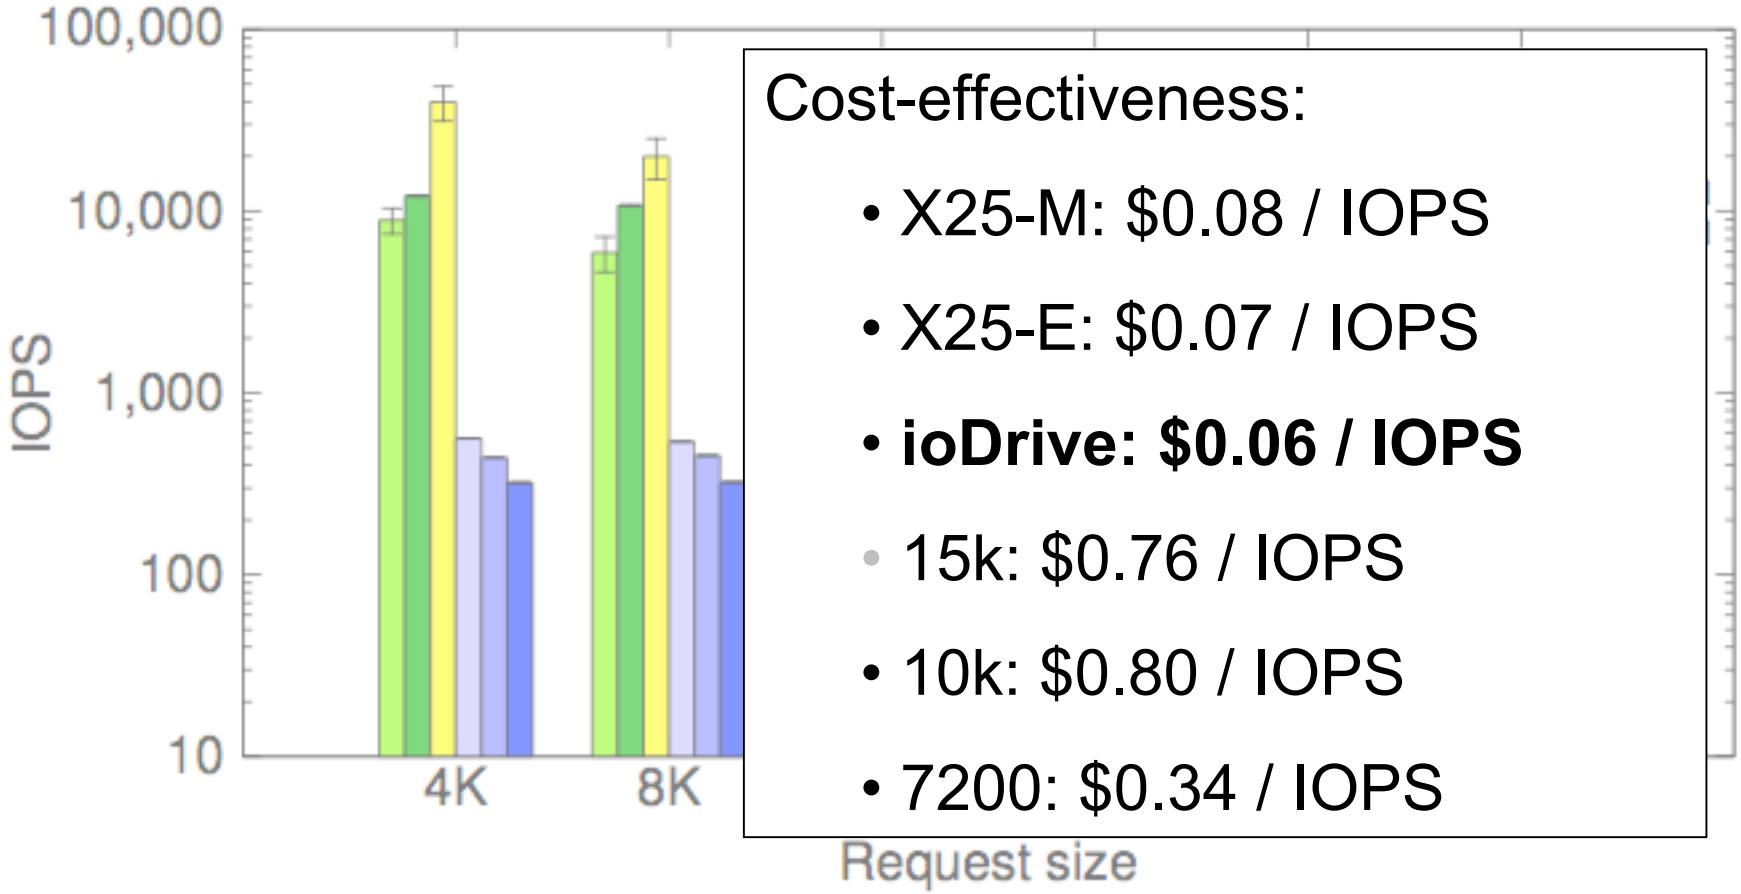

Cost-effectiveness:

- X25-M: $0.08 / IOPS
- X25-E: $0.07 / IOPS
- **ioDrive: $0.06 / IOPS**
- 15k: $0.76 / IOPS
- 10k: $0.80 / IOPS
- 7200: $0.34 / IOPS

Interpretation

- For the measured devices:
  - For any random I/O, SSDs achieve:
    - better $ per IOPS
    - similar $ per throughput
  - High performance enabled by parallelism and controller sophistication
  - Target workloads for Flash:
    - Random I/O
    - AMR check-pointing

Summary

Price per performance

	Sequential reads	Sequential writes	Random reads	Random writes
Memoright GT	$5.80	$5.96	**$0.12**	$1.77
Mtron Mobi	$4.81	$4.91	**$0.08**	$2.49
Intel X25-M	**$3.35**	$12.15	**$0.16**	**$0.08**
Intel X25-E	**$3.64**	$4.47	**$0.12**	**$0.07**
FusionIO ioDrive	$5.64	$14.79	**$0.11**	**$0.06**
15k rpm	$7.43	$7.61	$1.46	$0.76
10k rpm	$5.78	$5.47	$1.55	$0.80
7200 rpm	**$2.34**	**$2.59**	$1.00	$0.34

7200 rpm have cheapest sequential bandwidth

SSDs have much cheaper read IOPS

Enterprise SSDs have much cheaper write IOPS

Ongoing and Future Work

- Concurrent I/O and SSDs
  - measured numbers fall behind advertised ones
  - role of the file system
  - future measurements using IOMeter

- Infrastructure for hybrid storage
  - Virtual device aggregating disks and SSDs
  - Annotation of existing file systems
  - Caching algorithms
  - Data partitioning approaches

Questions?

Thank you for your attention.