


Heat-Based Dynamic Data Caching: A Load Balancing Strategy for Energy-Efficient Parallel Storage Systems with Buffer Disks

Ziliang Zong, Xiao Qin, Xiaojun Ruan and Mais Nijim
Department of Mathematics and Computer Science
South Dakota School of Mines and Technology



Buffer Disk Architecture

Energy-Efficient Disk Systems



Active



Sleep



Idle

The states of disks






Buffer Disk Architecture


Energy-Efficient Disk Systems

Parameters	IBM 36Z15 Ultrastar (high perf.)	IBM 73LZX Ultrastar (low perf.)
Number of platters	4	2
Rotations per minute	15000	10000
Average seek time	3.4 ms	4.9 ms
Average rotation time	2 ms	3 ms
Transfer rate	55 MB/sec	53 MB/sec
Power (active)	13.5 W	9.5 W
Power (idle)	10.2 W	6.0 W
Power (sleep)	2.5W	1.4W
Energy (spin down)	13.0 J	10.0 J
Time (spin down)	1.5 s	1.7 s
Energy (spin up)	135.0 J	97.9 J
Time (spin up)	10.9 s	10.1 s

Buffer Disk Architecture



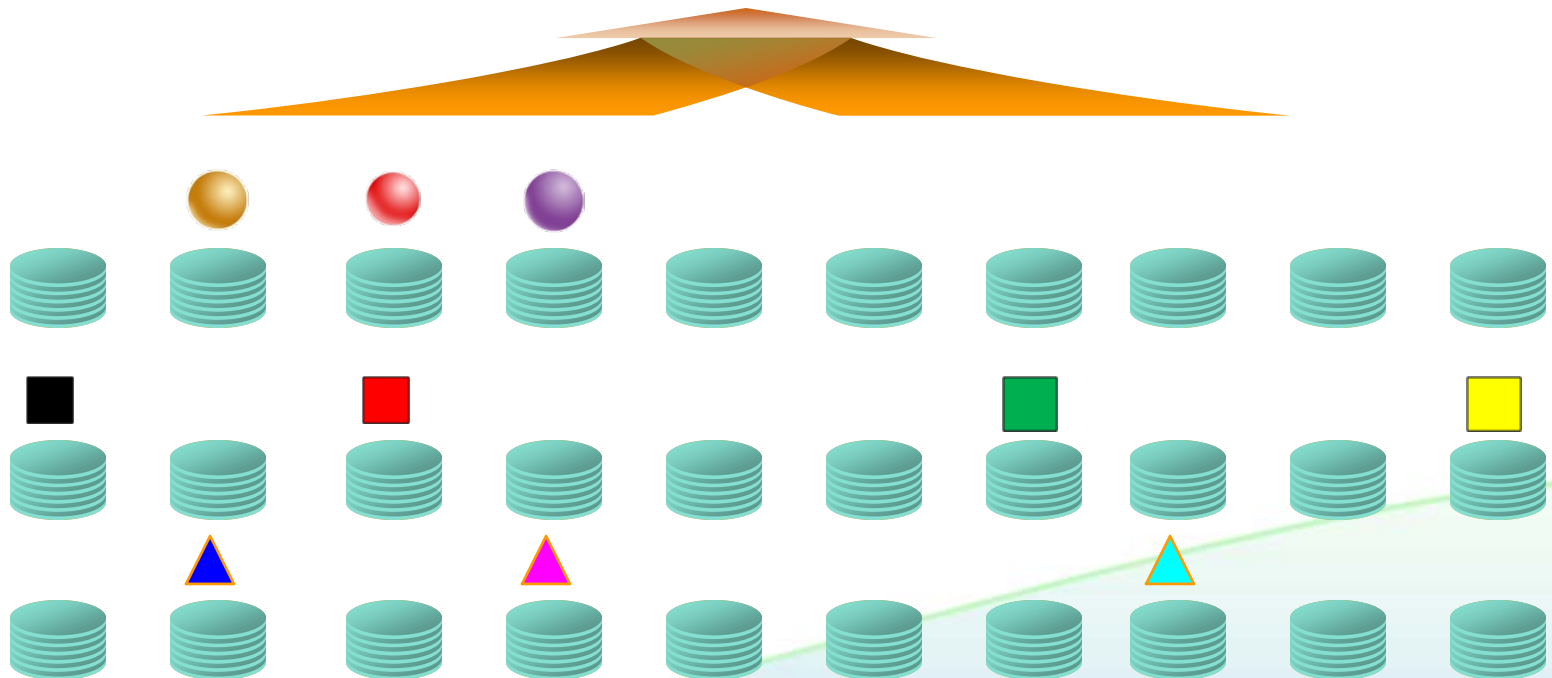



Requests Queue



Buffer Disk Architecture

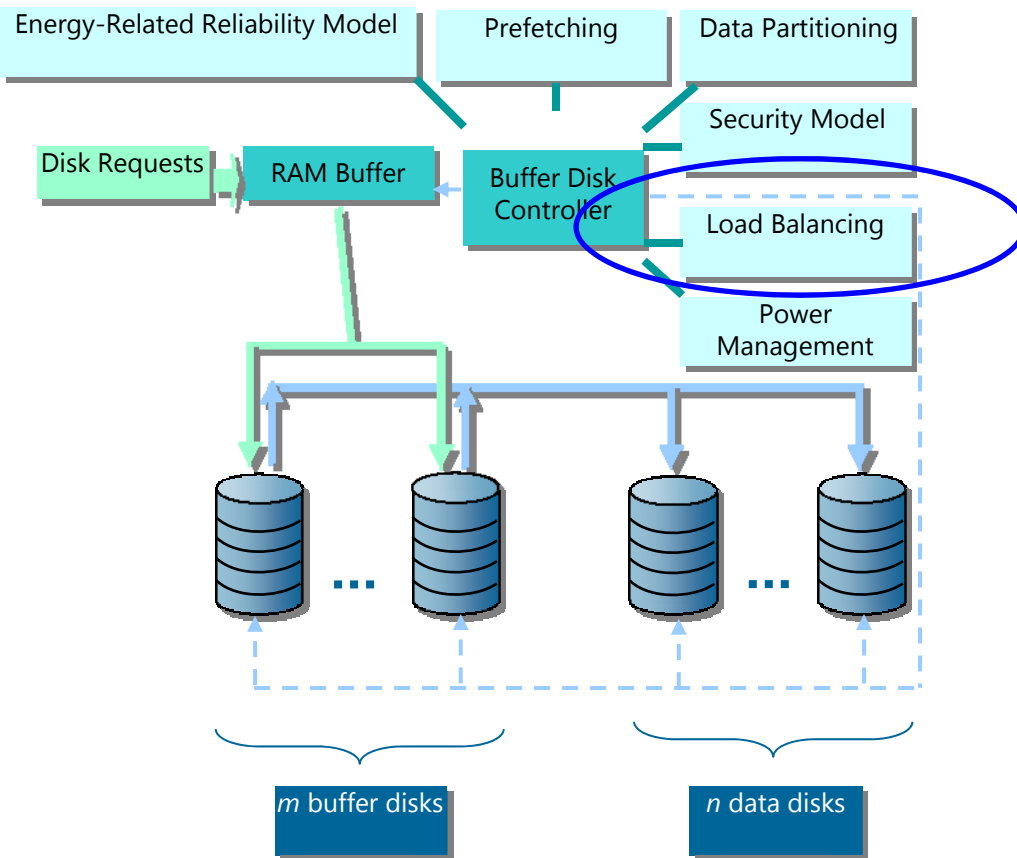
Energy-Efficient Disk Systems

Requests Queue



Buffer Disk Architecture

Energy-Efficient Disk Systems



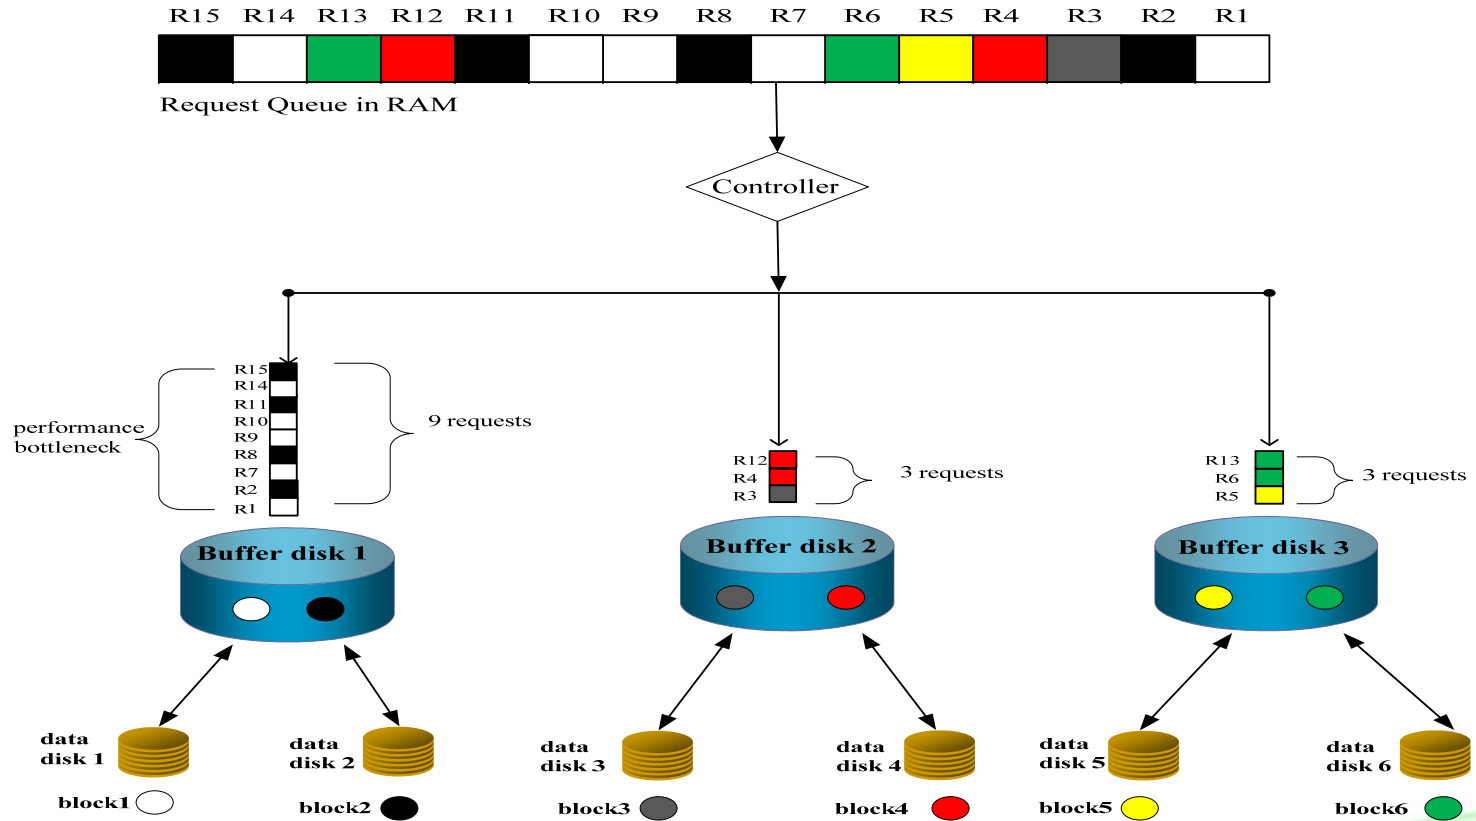
A general framework

Large number of data disks

Small number of buffer disks

Buffer Disk Architecture

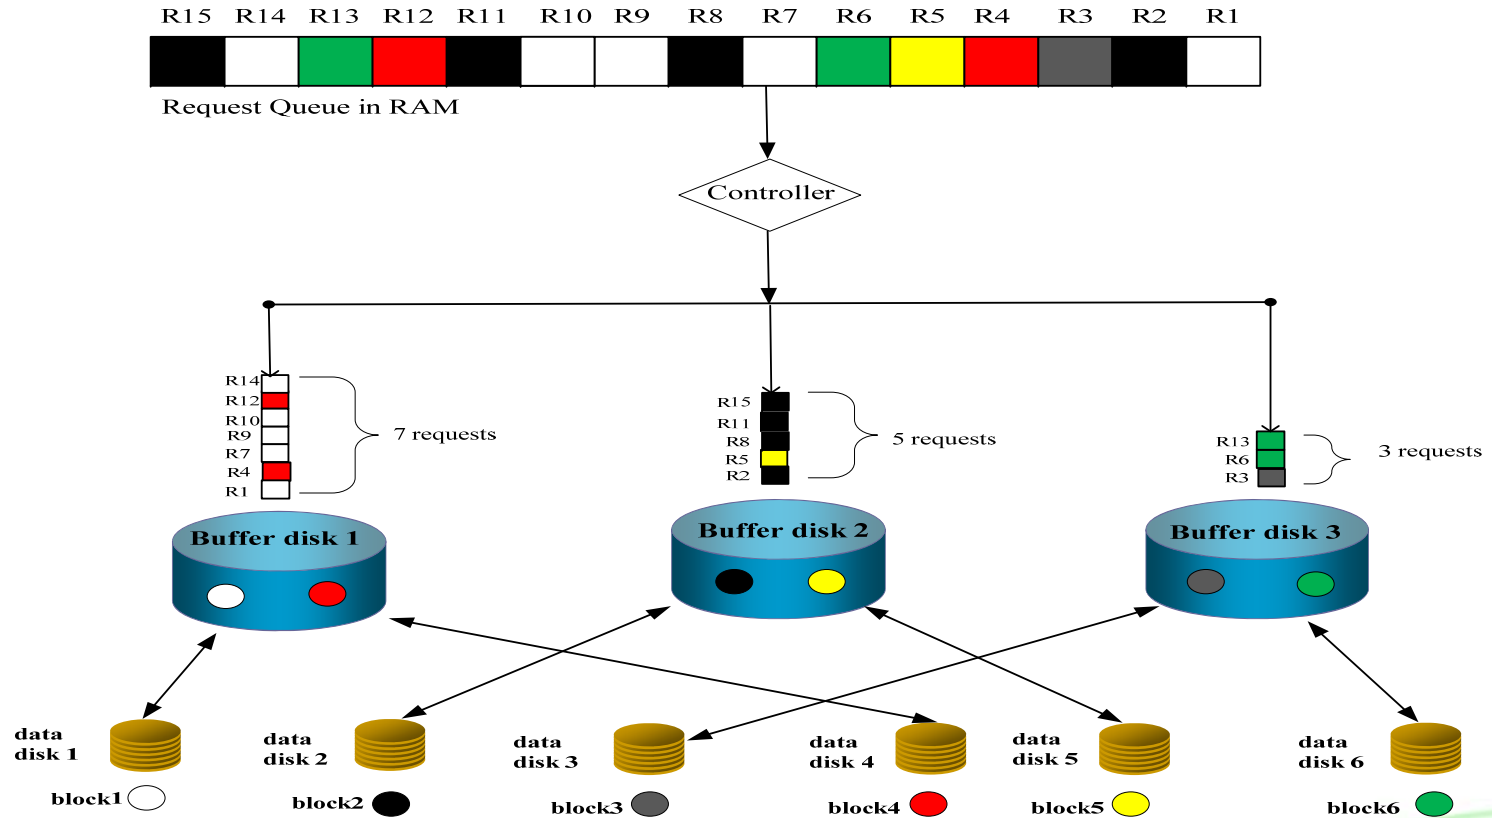
Energy-Efficient Disk Systems



Sequential Mapping Results

Buffer Disk Architecture

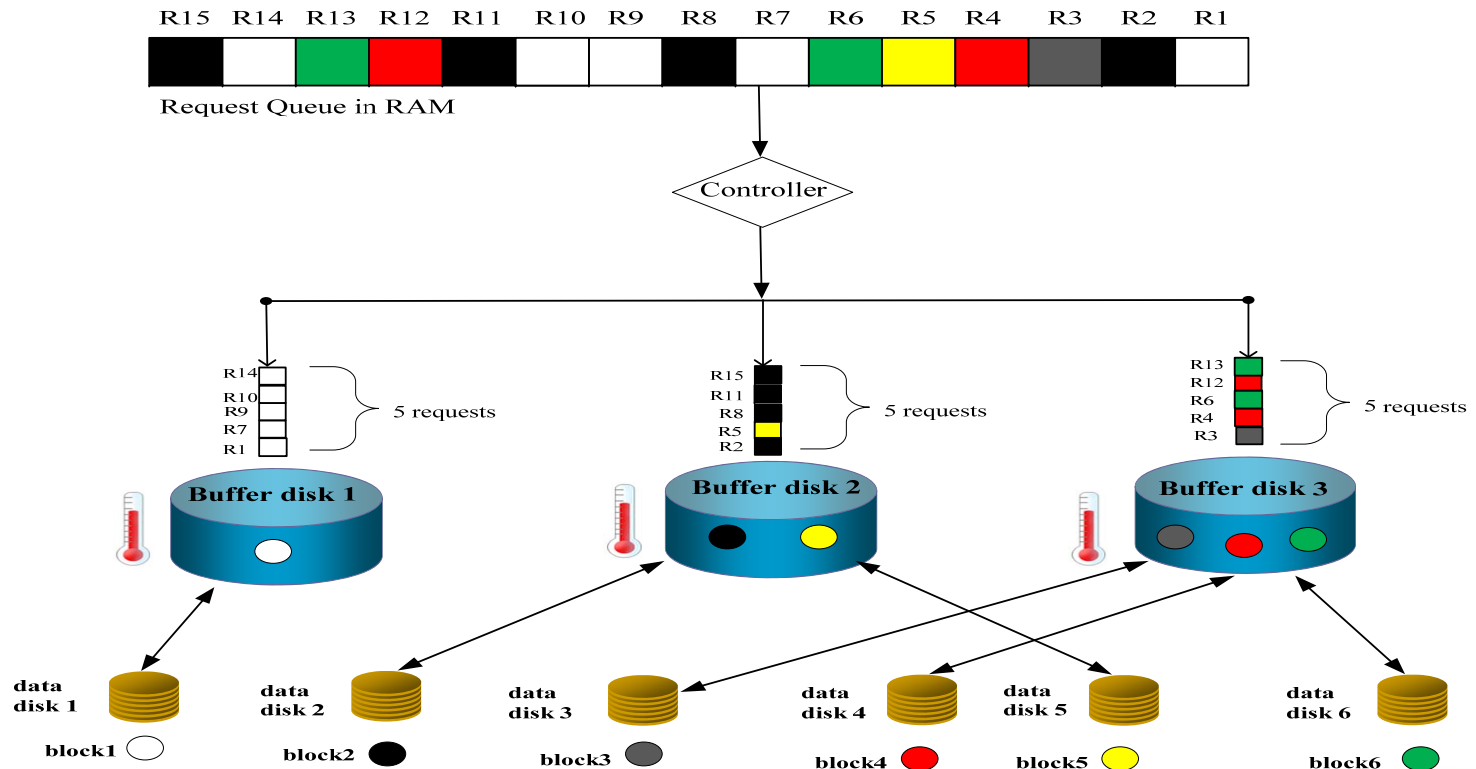
Energy-Efficient Disk Systems



Round-Robin Mapping Results

Buffer Disk Architecture

Energy-Efficient Disk Systems



Heat-Based Load Balancing Results

Buffer Disk Architecture

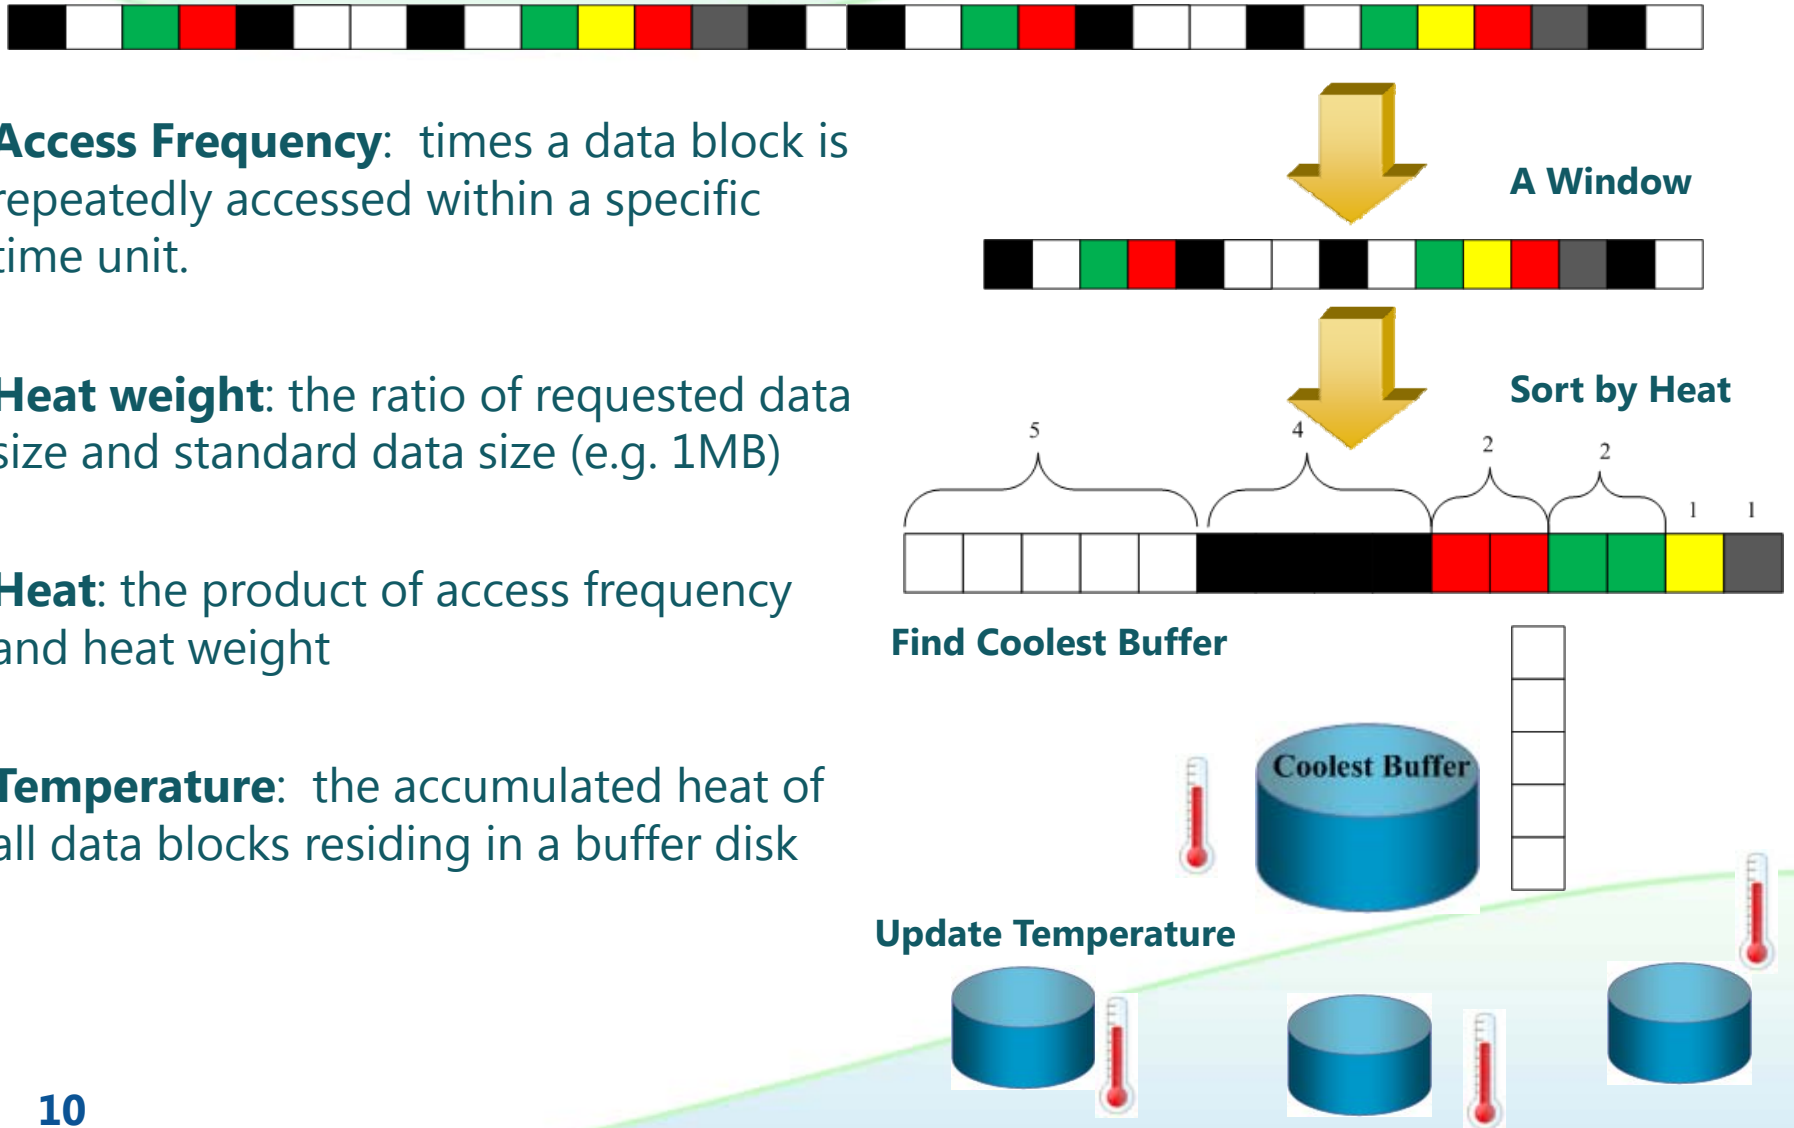
Energy-Efficient Disk Systems

Access Frequency: times a data block is repeatedly accessed within a specific time unit.

Heat weight: the ratio of requested data size and standard data size (e.g. 1MB)

Heat: the product of access frequency and heat weight


Temperature: the accumulated heat of all data blocks residing in a buffer disk



Buffer Disk Architecture


Energy-Efficient Disk Systems

Large Reads: average 84.4% improvement (64MB)



Energy Consumption for Large Reads.

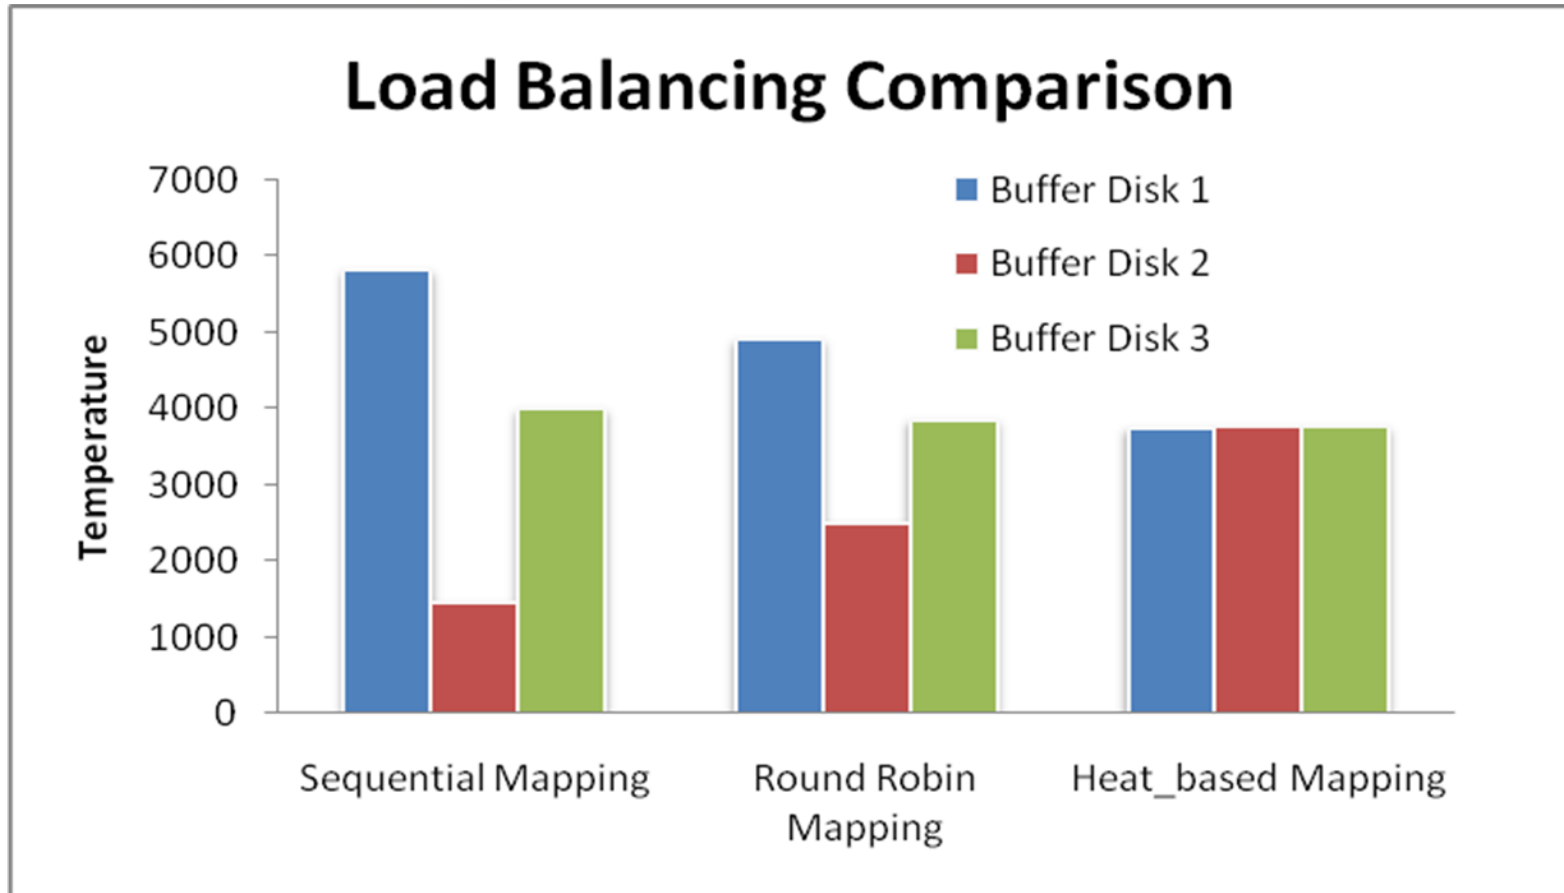
Small Reads: average 78.77% improvement (64KB)



Energy Consumption for Small Reads.

Buffer Disk Architecture


Energy-Efficient Disk Systems



Load Balancing Comparison of Three Mapping Strategies.

Buffer Disk Architecture

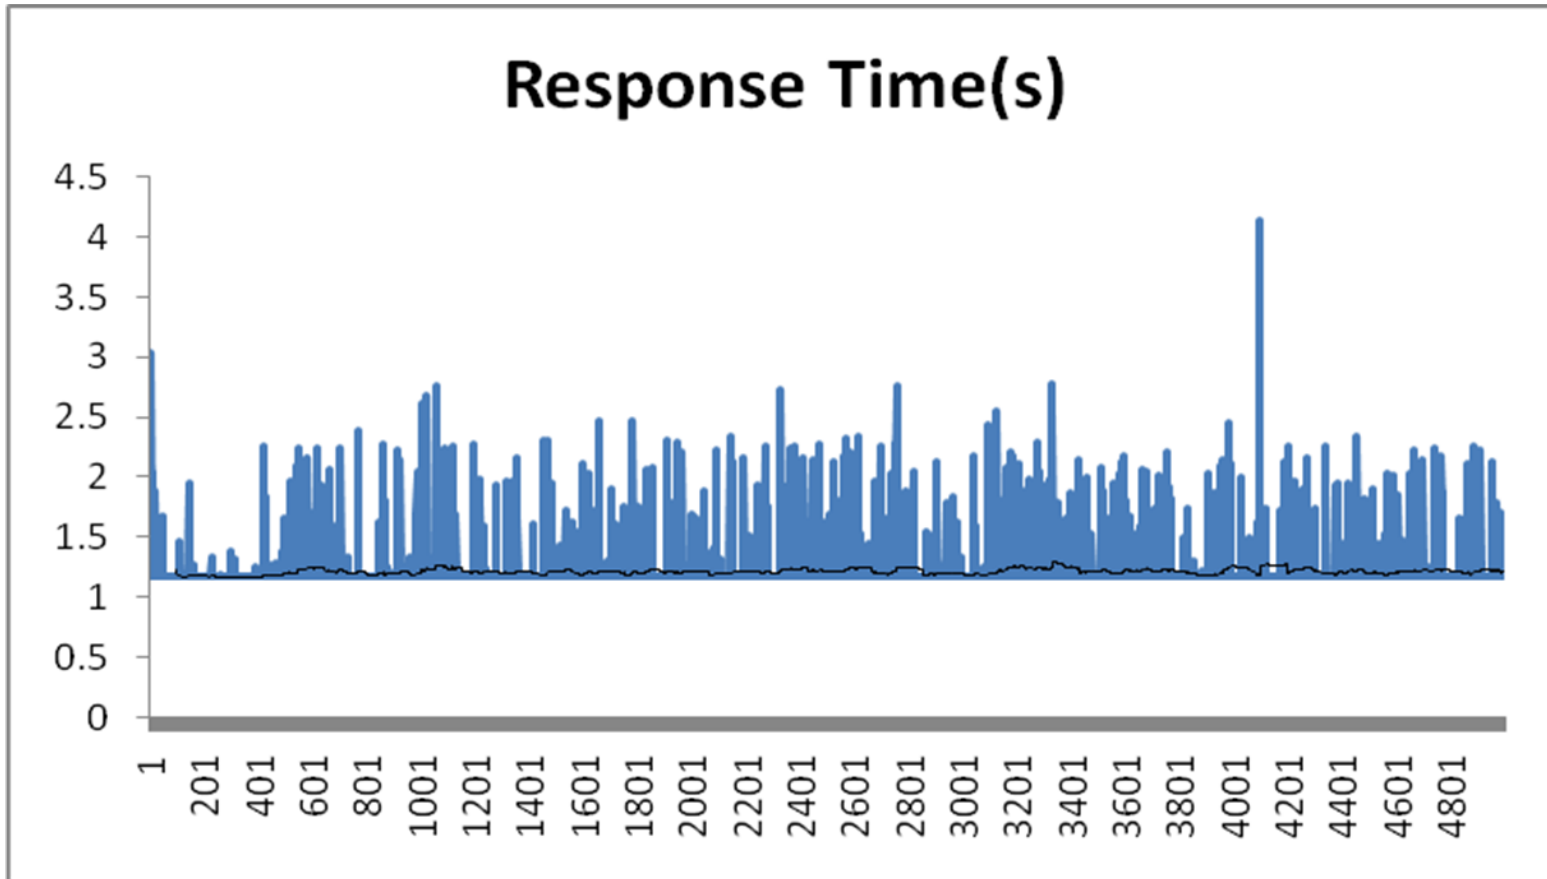
Energy-Efficient Disk Systems



Response Time Trace before Caching (Large Reads).

Buffer Disk Architecture


Energy-Efficient Disk Systems



Response Time Trace after Caching (Large Reads).

Buffer Disk Architecture

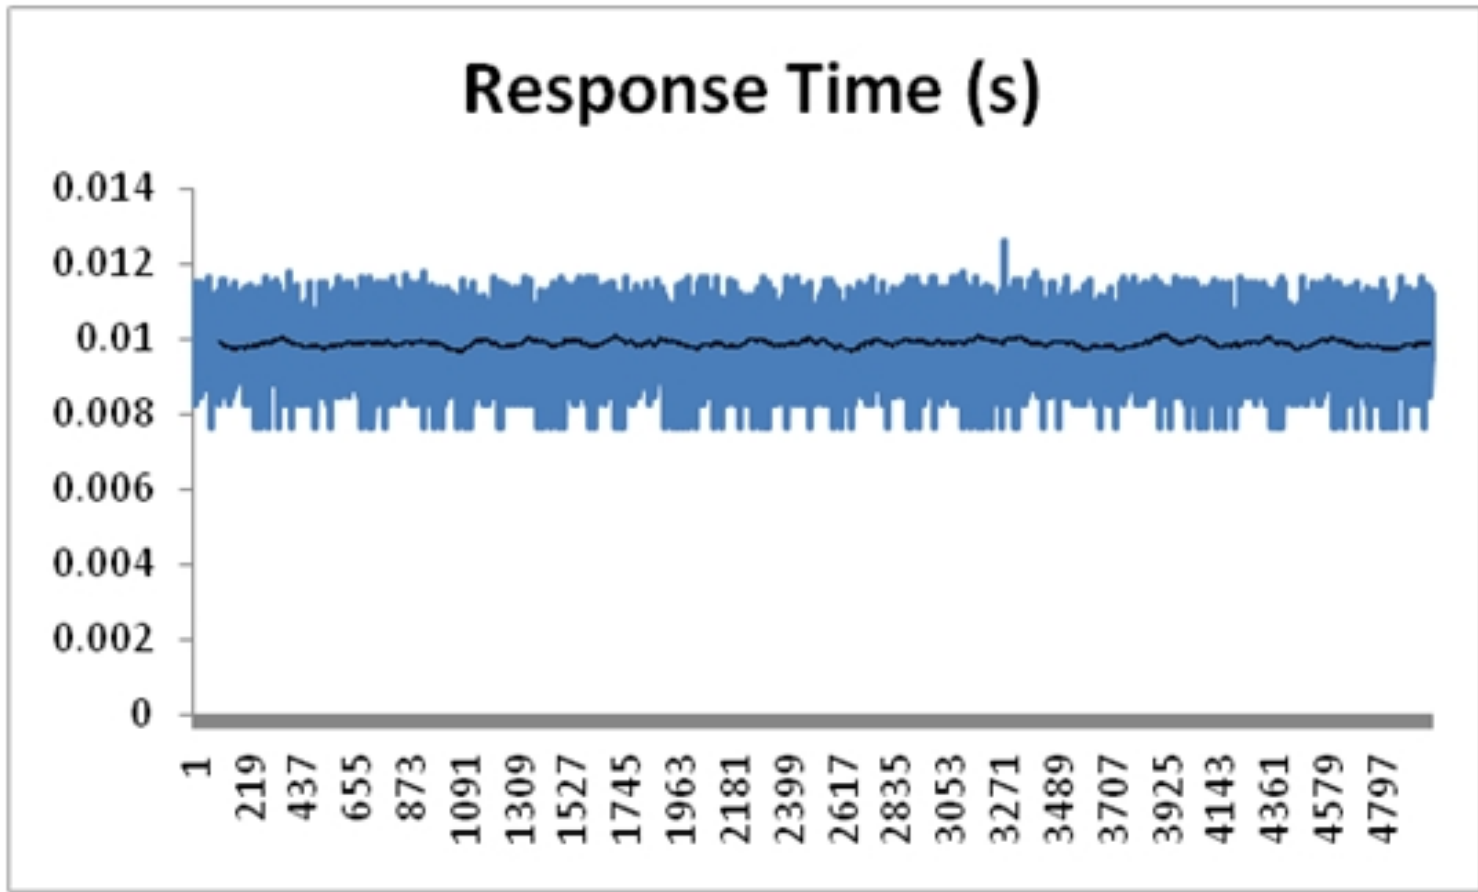
Energy-Efficient Disk Systems



Response Time Trace before Caching (Small Reads).

Buffer Disk Architecture

Energy-Efficient Disk Systems



Response Time Trace after Caching (Small Reads).

Buffer Disk Architecture

Energy-Efficient Disk Systems

		Average Response Time
Large requests	Before caching (64MB):	5.614s
	After caching (64MB):	1.219s
	No-Buffer(64MB)	1.216s
Small requests	Before caching (64KB):	0.767s
	After caching (64KB):	0.01s
	No-Buffer(64KB)	0.01s



Thank you !

Questions?

