CAMELab.org
Computer Architecture and Memory Systems Lab.

UT DALLAS

Power, Energy and Thermal Considerations in SSD-Based I/O Acceleration

Jie Zhang, Mustafa Shihab and Myoungsoo Jung

Computer Architecture and Memory Systems Lab

Department of Electrical Engineering

The University of Texas at Dallas

Summary

Challenges: State-of-art SSD integrates many more internal resources compared to traditional ones. Until now, it has not been considered how this growing number of internal resources impact **power, energy and thermal factor** on SSD.

Summary

Challenges: State-of-art SSD integrates many more internal resources compared to traditional ones. Until now, it has not been considered how this growing number of internal resources impact **power, energy and thermal factor** on SSD.

Goals: Introduce the power, energy and temperature impact in high-end SSDs and address the critical issues related with growing number of internal resources.

Summary

Challenges: State-of-art SSD integrates many more internal resources compared to traditional ones. Until now, it has not been considered how this growing number of internal resources impact **power, energy and thermal factor** on SSD.

Goals: Introduce the power, energy and temperature impact in high-end SSDs and address the critical issues related with growing number of internal resources.

Contributions: We quantitatively analyze challenges for SSDs in integrating more resources:

- Operating temperature
- Dynamic power behaviors
- Energy consumption
- Overheating and power throttling issues

Outline

- Motivation
- Evaluation Setup
- Findings
  - Operating temperature
  - Dynamic power behavior
  - Energy consumption
  - Overheating problem and power throttling issue
- Conclusion

Increased Resources in SSDs

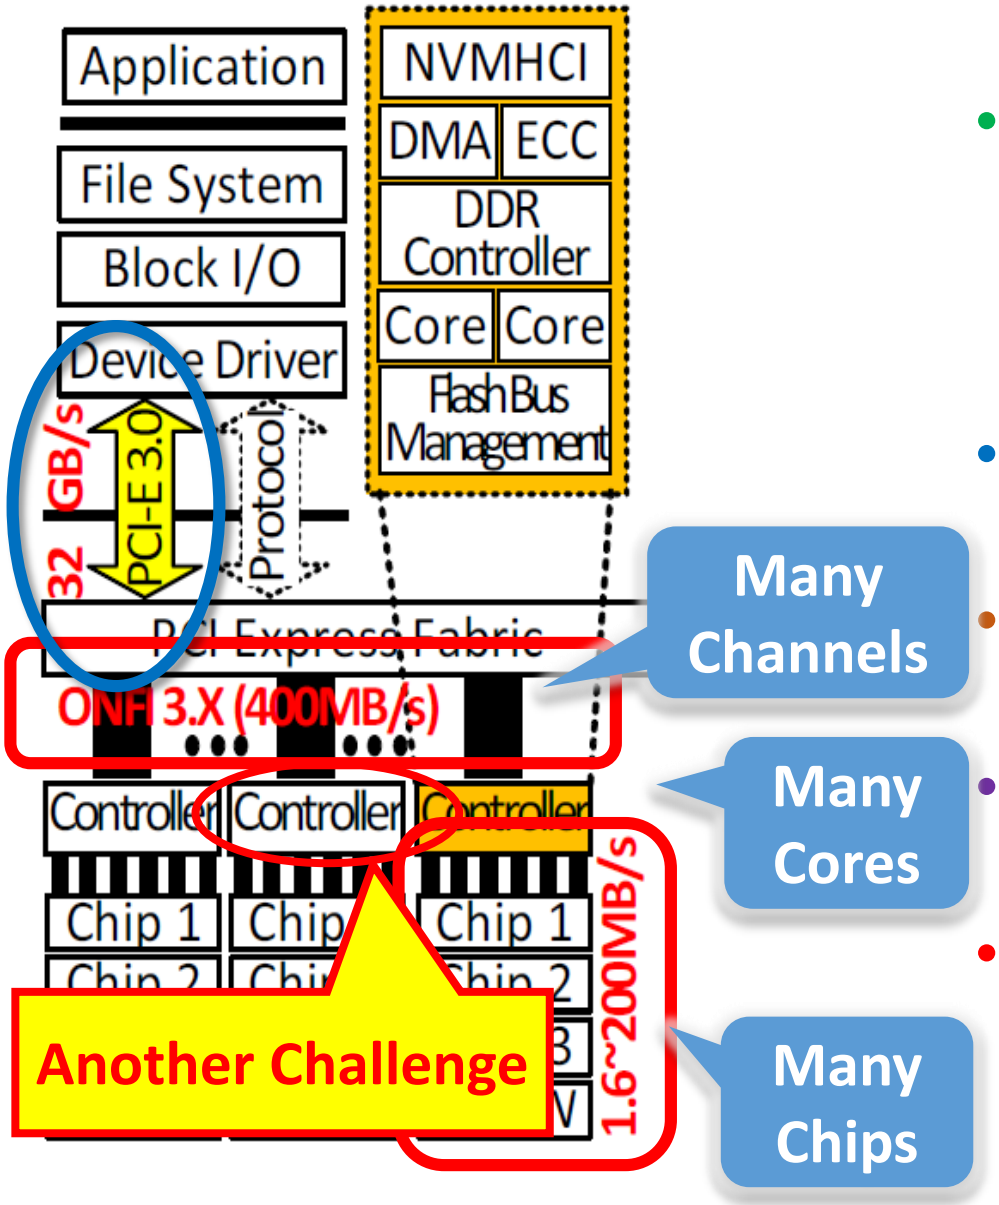

- Huge bandwidth gap between single flash chips/channels and the host interface
- Many chips to utilize chip-level parallelism
- Many channels to utilize channel-level parallelism
- Many controllers/cores to handle multiple tasks
- Overall, more and more resources are integrated in state-of-art SSDs

Resource Integration Trend

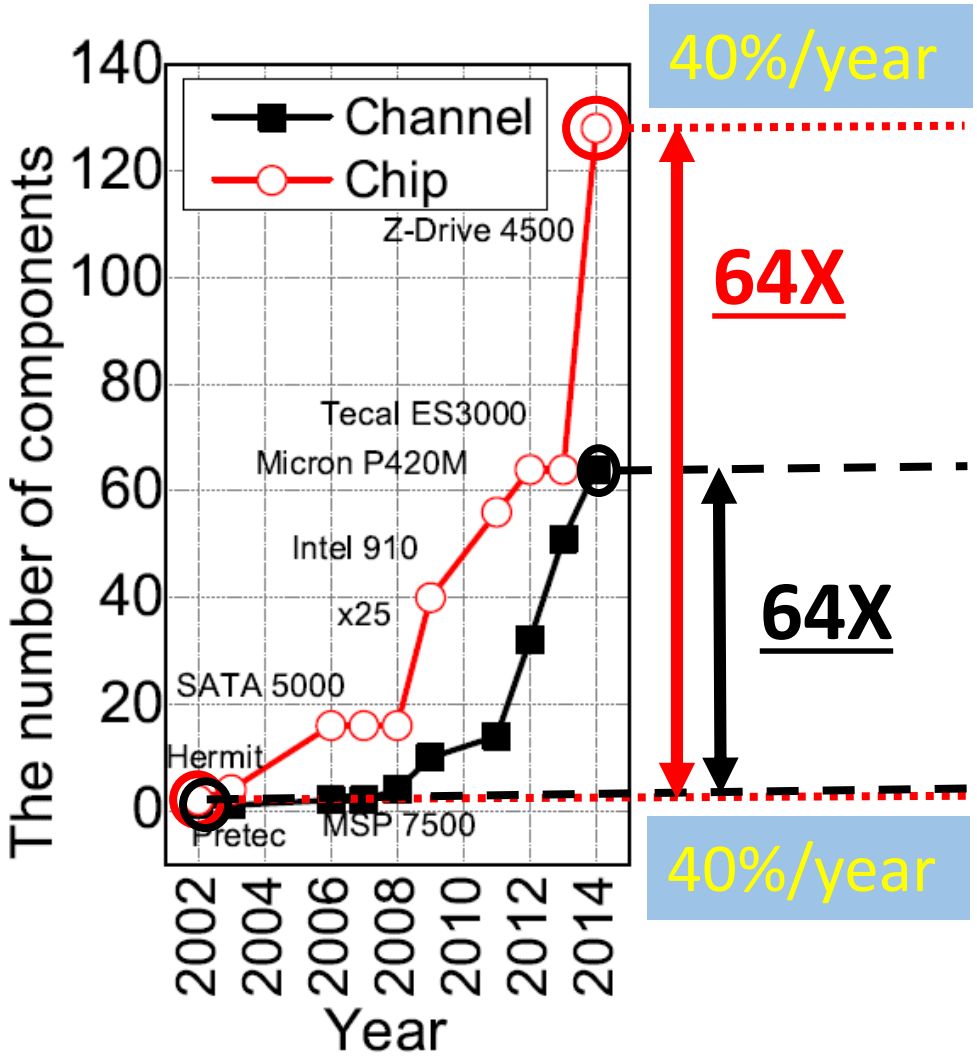

- Over the past 12 years, the number of chips and channels has increased by 64 times
- Following this trend, there will be even more internal resources in the future
- These integrated power contributors cause potential power concern
- Essential to study the power, energy, and thermal factors of modern SSDs

Outline

- Motivation
- Evaluation Setup
- Findings
  - Operating temperature
  - Dynamic power behavior
  - Energy consumption
  - Overheating problem and power throttling issue
- Conclusion

Evaluation Setup: SSD Testbeds

Baseline SSD

Multi-core SSD

Multi-resource SSD

	SP-SSD	MC-SSD	MR-SSD
Feature	Multi-Channel	Multi-Core	Multi-Resource
Interface	SATA 6Gbps	PCIe 2.0	PCIe 2.0
Core	3	16	16
# of Channels	8	8	32
# of Chips	64	64	128
DRAM Sizes	256MB	2GB	2.25GB
Storage Cap.	512GB	400GB	512GB

- MC-SSD vs. SP-SSD: more cores, larger DRAM
- MR-SSD vs. SP-SSD: more channels, more chips

Evaluation Setup: Workload & Dev. Status

Data Type	SEQ			RND		
Data Size	4KB	8KB	16KB	32KB	64KB	128KB
	256KB	512KB	1MB	2MB	4MB	
Device Status	Pristine			Aged		

- Data type: **Sequential data** and **Random data**
- Data size: **4KB** to **4MB**
- Device status:
  - Pristine device: SSD with factory default status
  - Aged device: mimic actual aged SSD
- **Based on these, there are 198 different configurations**

Outline

- Motivation
- Evaluation Setup
- Findings
  - Operating temperature
  - Dynamic power behavior
  - Energy consumption
  - Overheating problem and power throttling issue
- Conclusion

Operating Temperature

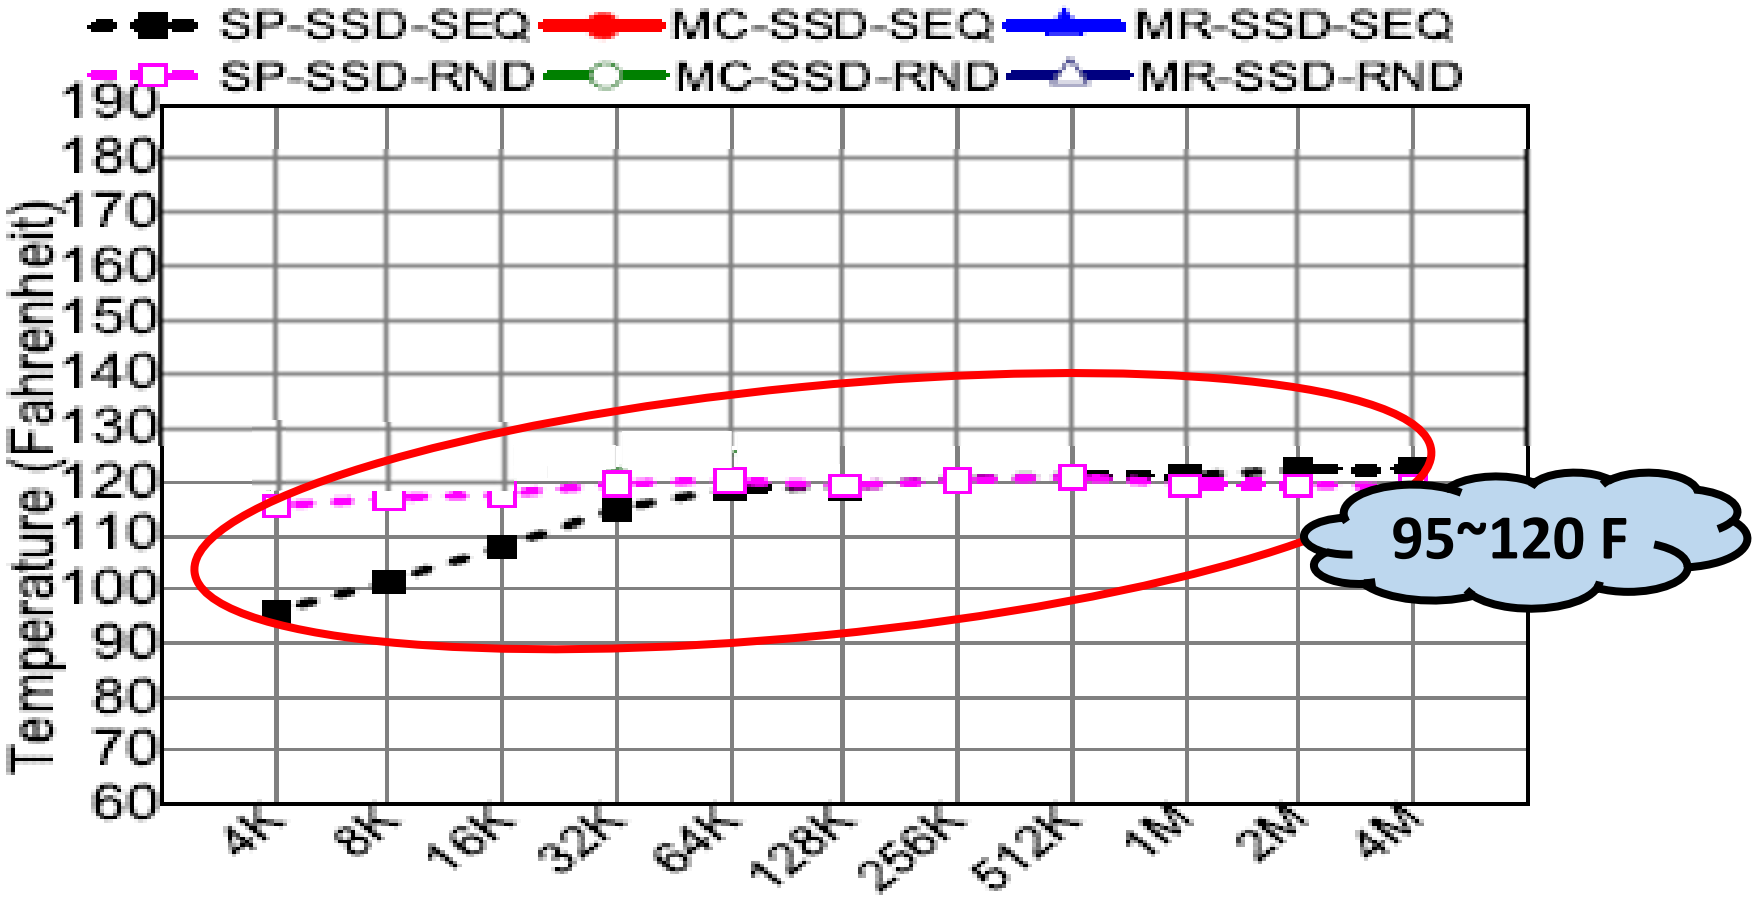

Write Access on Pristine Device

- Different request sizes have little impact on SP-SSD's temperature

Operating Temperature

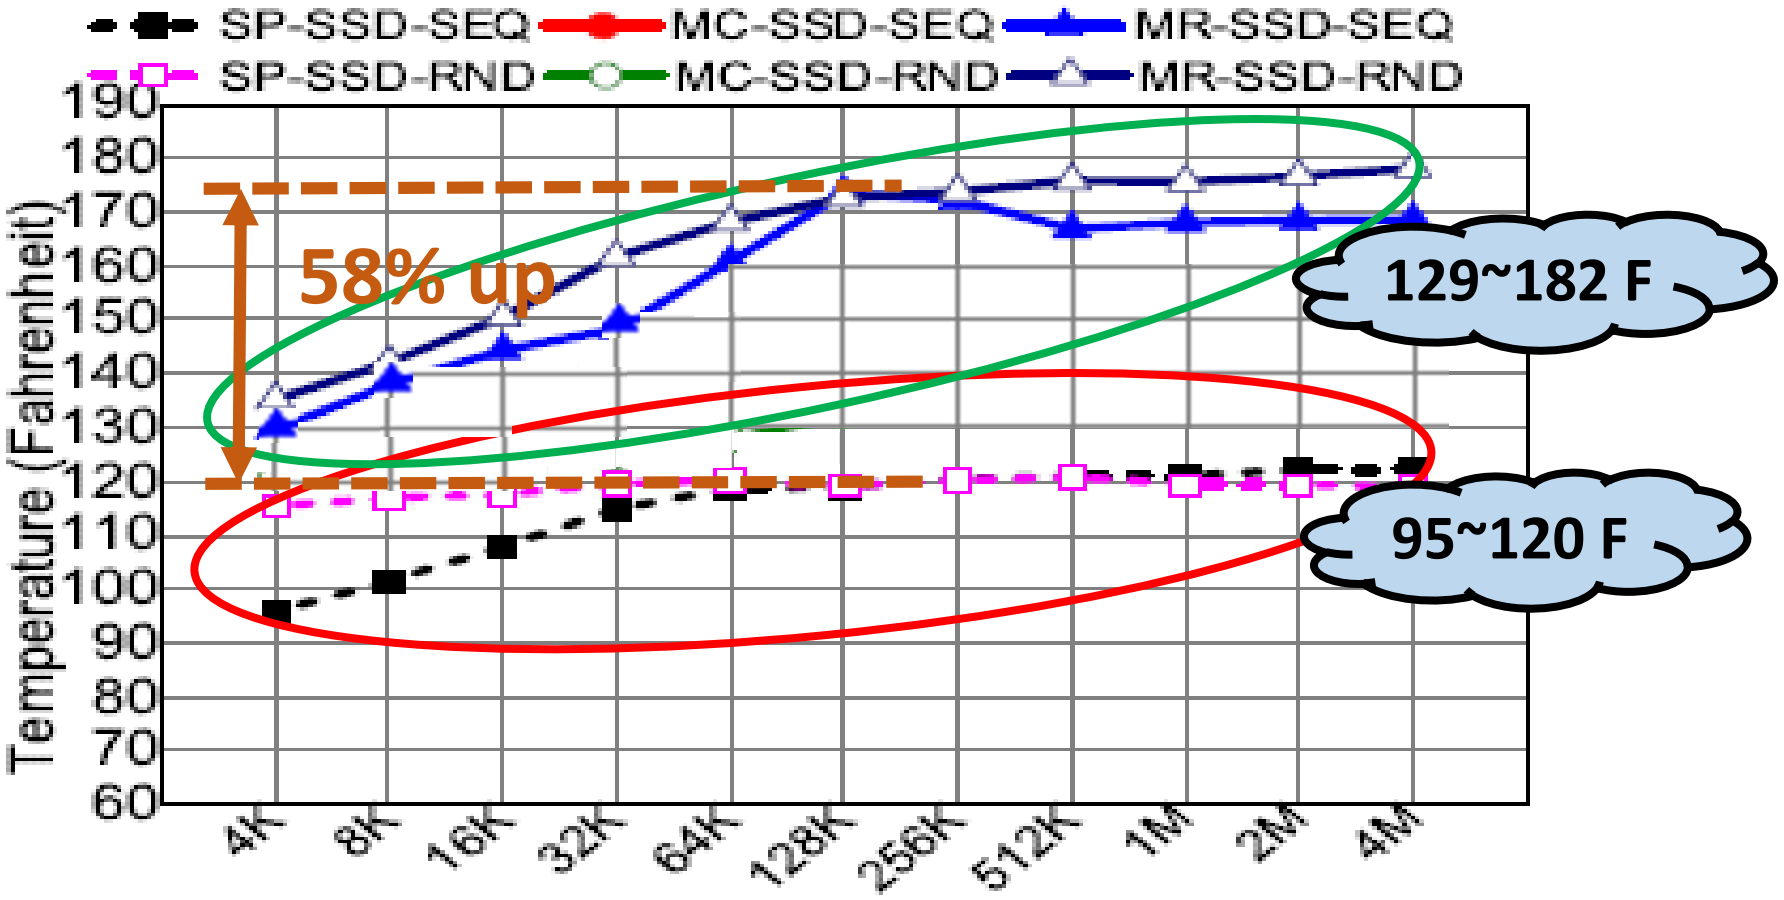

- Different request sizes have little impact on SP-SSD's temperature
- MR-SSDs exhibit higher temperature than SP-SSDs

Operating Temperature

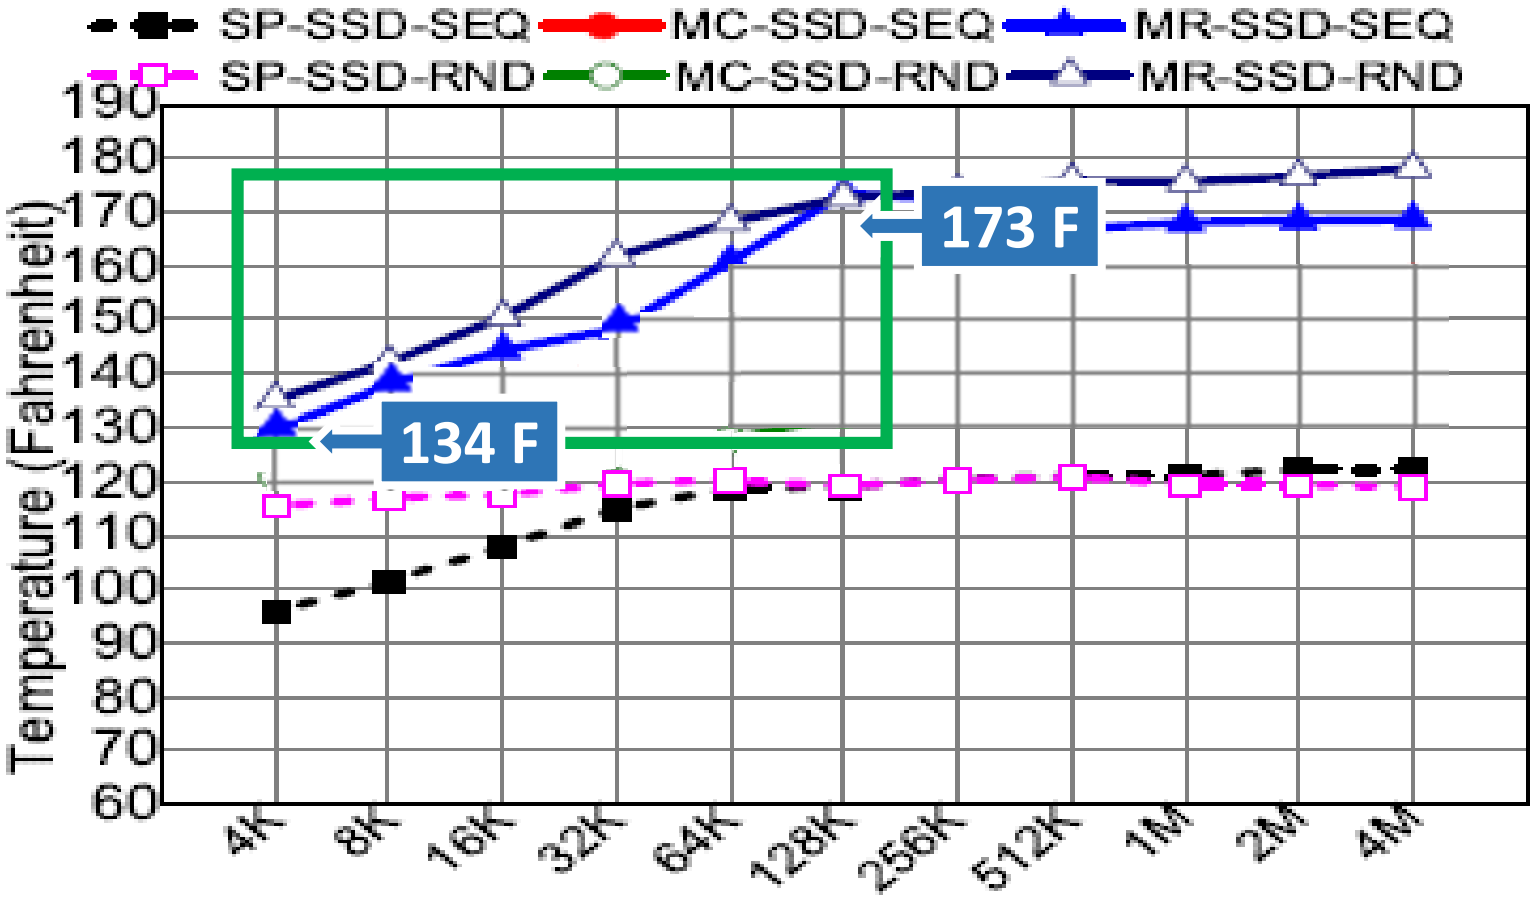

Write Access on Pristine Device

- Different request sizes have little impact on SP-SSD's temperature
- MR-SSDs exhibit higher temperature than SP-SSDs
- As for MR-SSDs, temperature increases as the data size increases

Operating Temperature

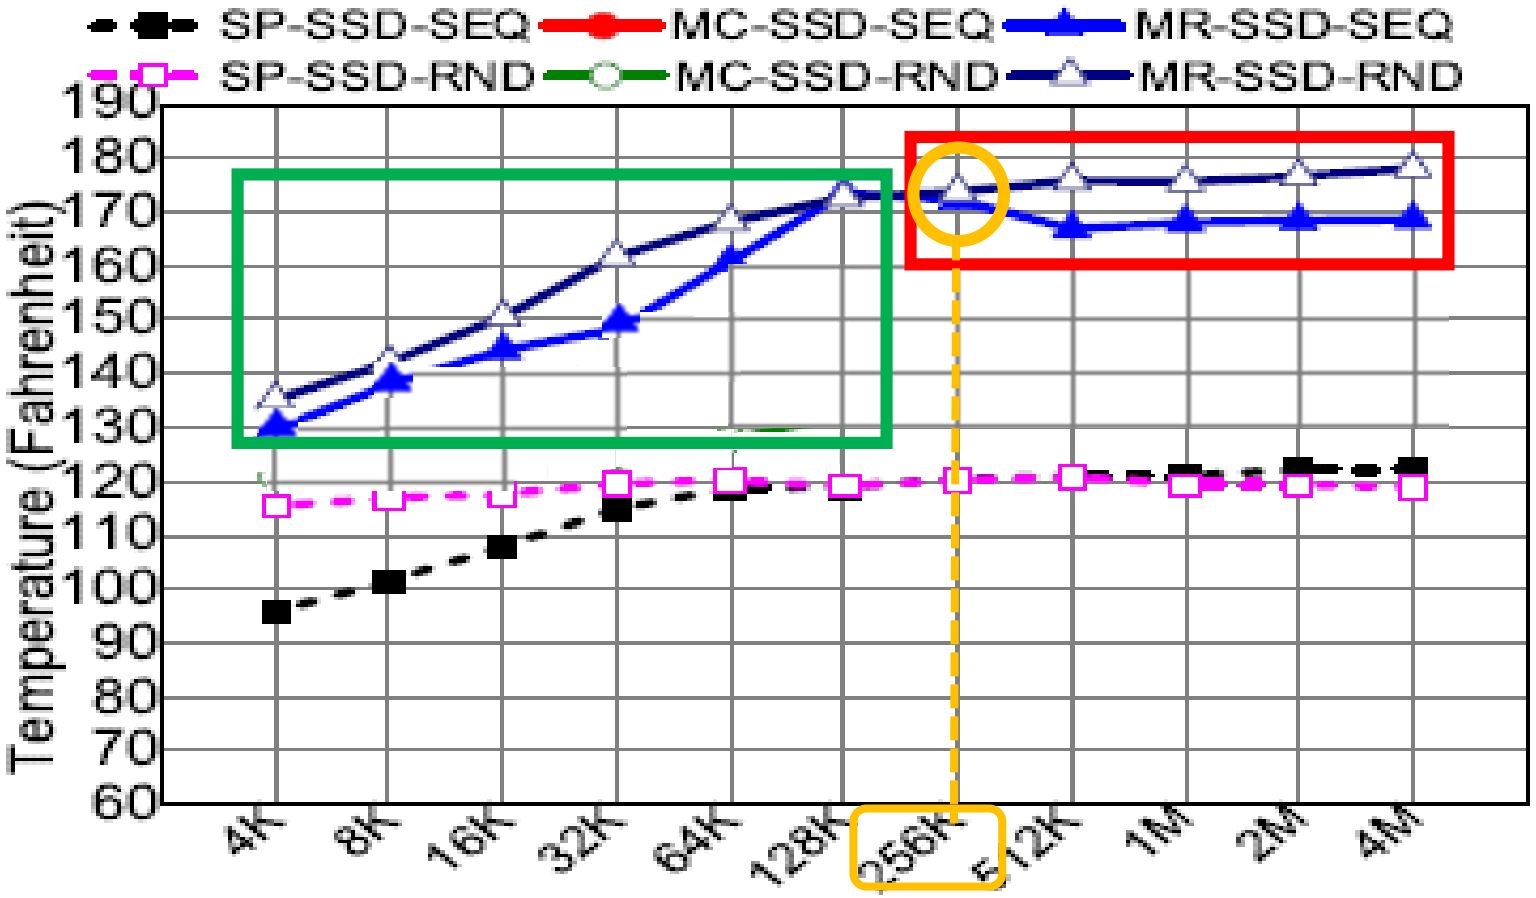

Write Access on Pristine Device

- Different request sizes have little impact on SP-SSD's temperature
- MR-SSDs exhibit higher temperature than SP-SSDs
- As for MR-SSDs, temperature increases as the data size increases
- Temperature saturates when data size increases to a certain level

Outline

- Motivation
- Evaluation Setup
- Findings
  - Operating temperature
  - Dynamic power behavior
  - Energy consumption
  - Overheating problem and power throttling issue
- Conclusion

Dynamic Power Behavior

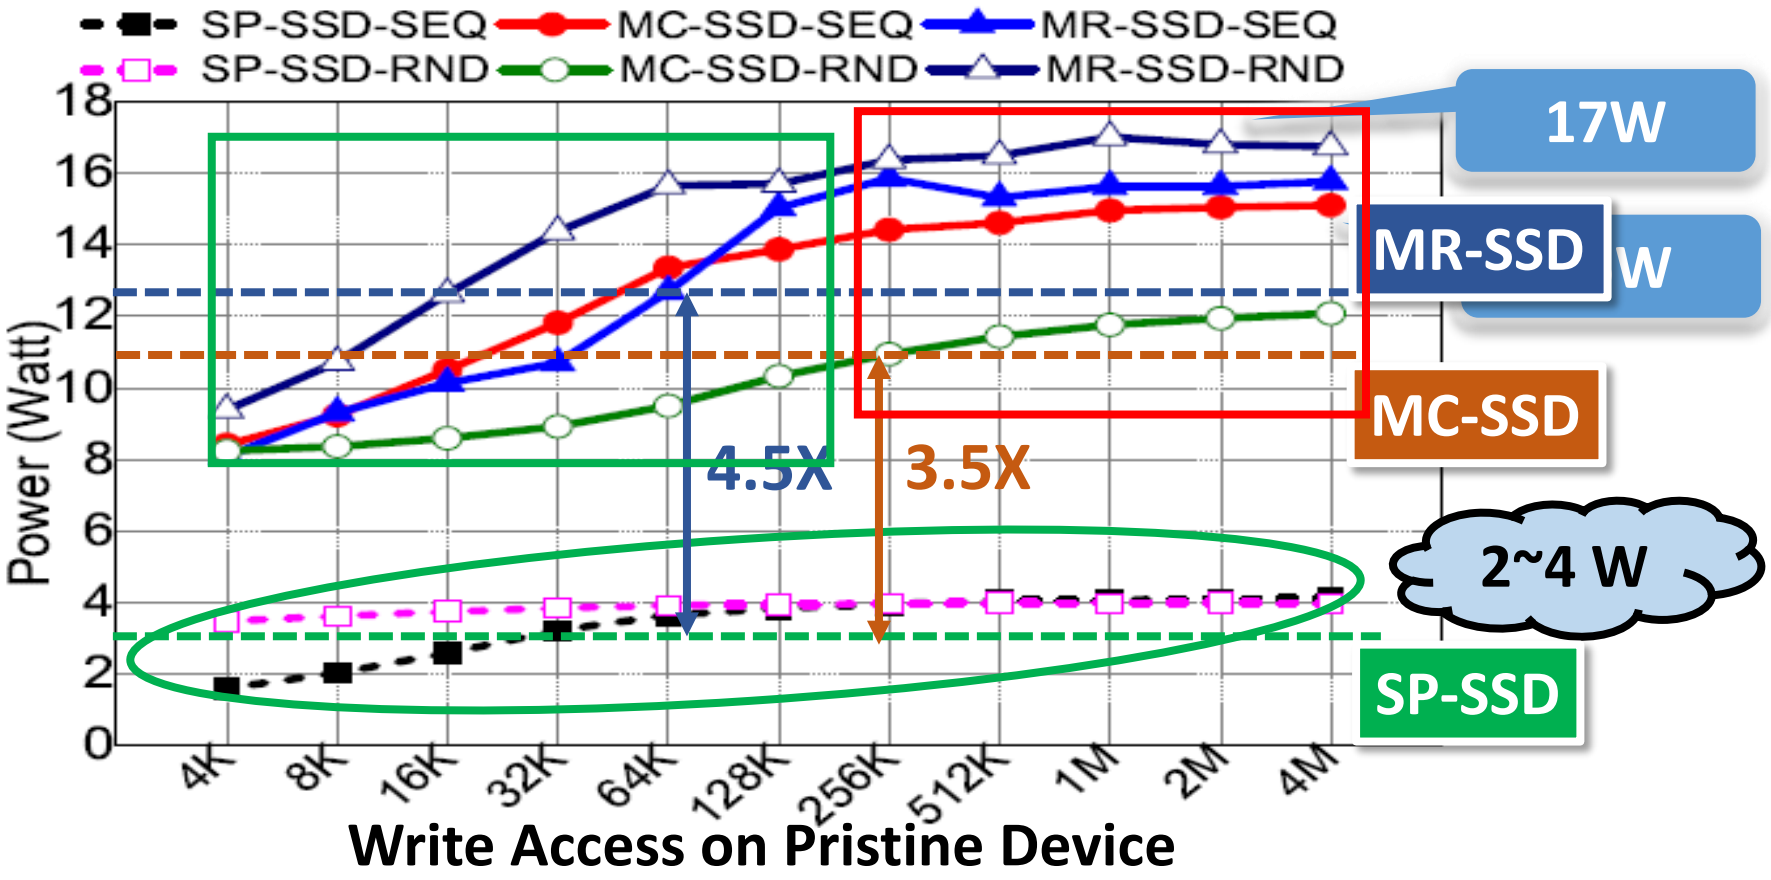

- Different request sizes have little impact on power of SP-SSD
- MC/MR-SSDs consume 3.5X and 4.5X more power than SP-SSDs
- As for MC/MR-SSDs, power increases as the data size increases
- Power saturates when data size increases to a certain level

Outline

- Motivation
- Evaluation Setup
- Findings
  - Operating temperature
  - Dynamic power behavior
  - Energy consumption
  - Overheating problem and power throttling issue
- Conclusion

Energy Consumption

- Common presumption: MR/MC-SSDs save energy because of perf.

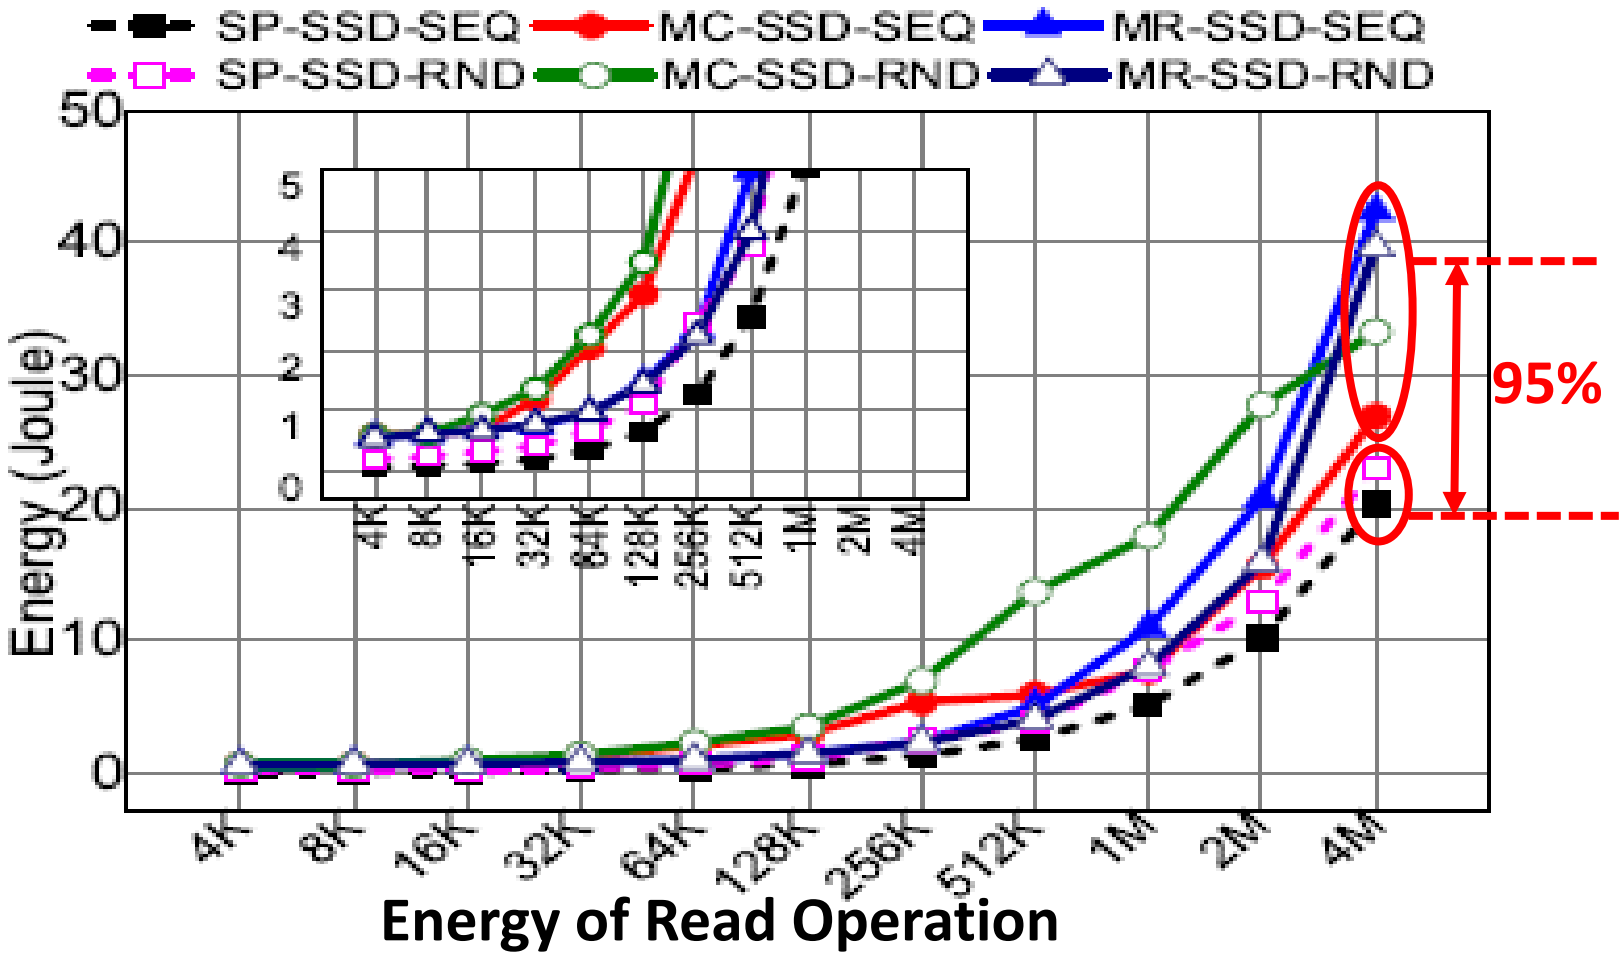

Energy Consumption

- ~~Common presumption: MR/MC-SSDs save energy because of perf.~~

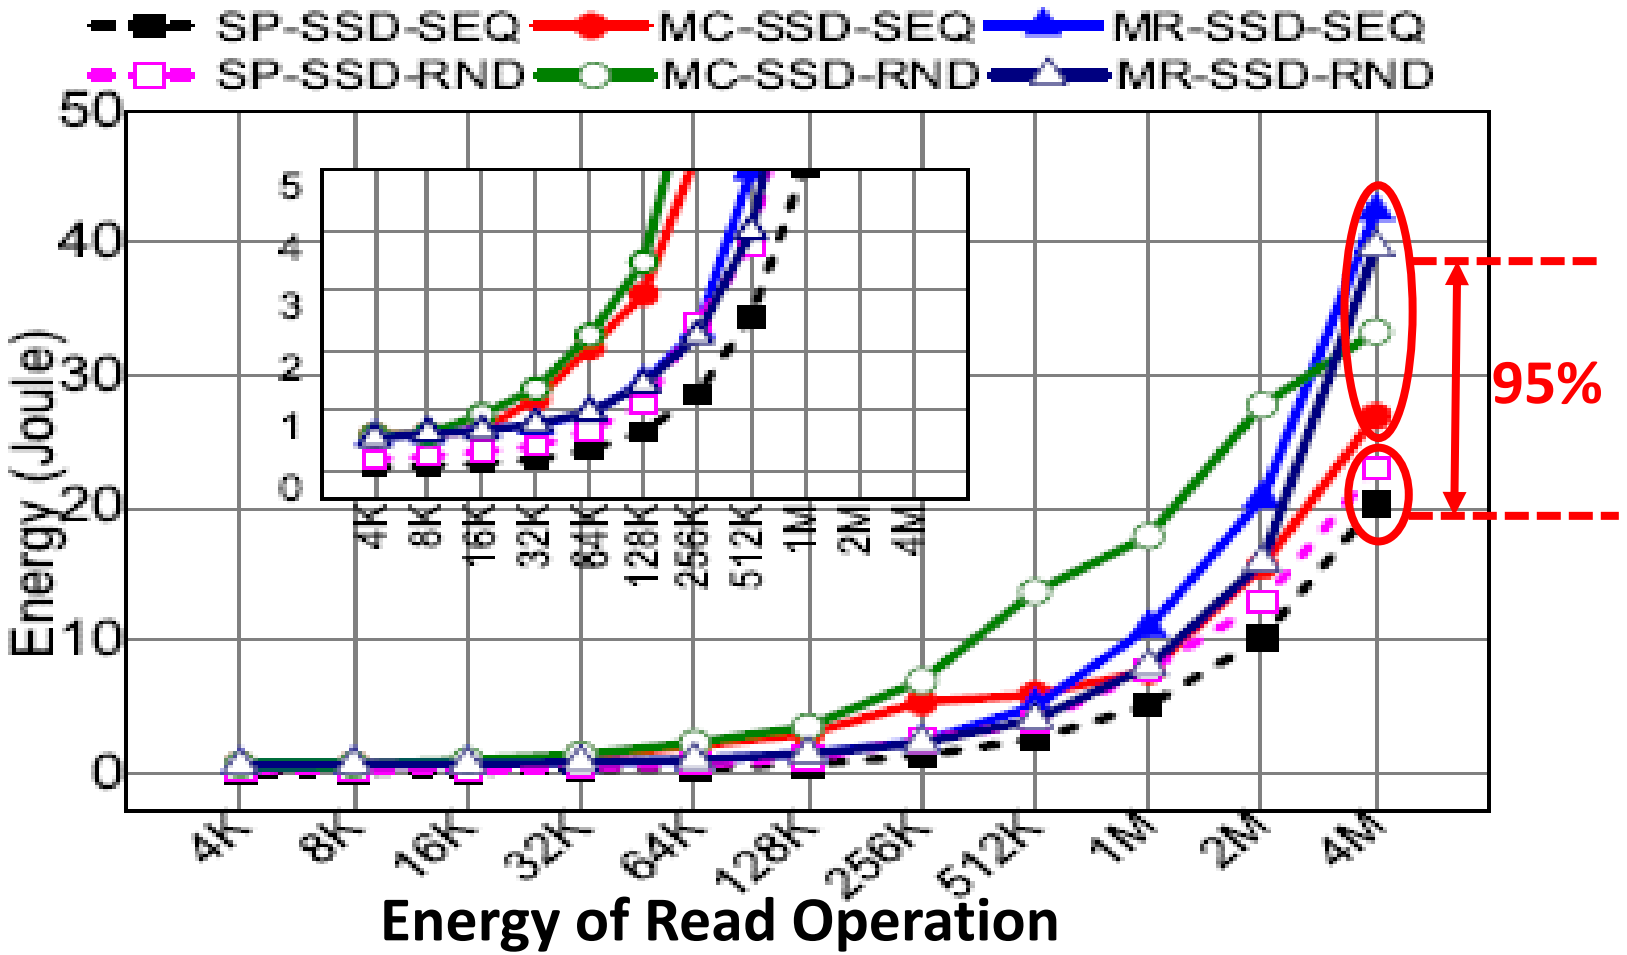

- MC/MR-SSDs cost more energy than SP-SSDs
- MC/MR-SSDs require more power to feed its internal resources

Outline

- Motivation
- Evaluation Setup
- Findings
  - Operating temperature
  - Dynamic power behavior
  - Energy consumption
  - Overheating problem and power throttling issue
- Conclusion

Overheating Problem

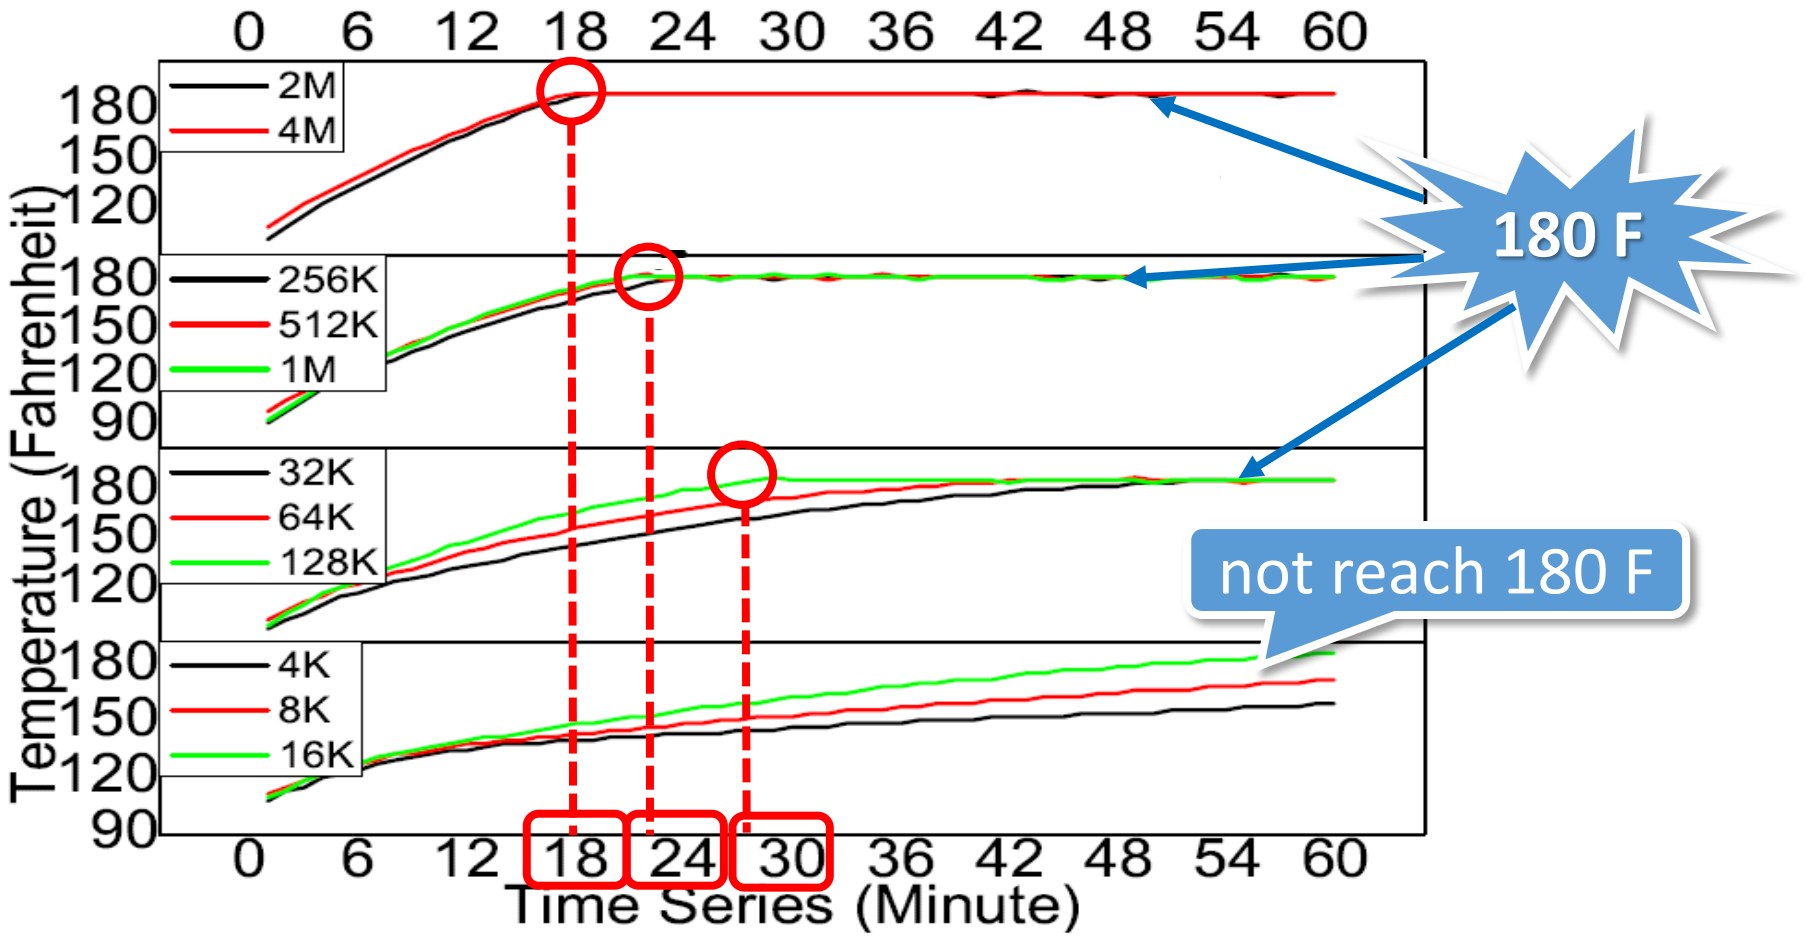

- Overheating problem: heat output exceeds certain level which may lead to malfunction of SSDs
- Device-level protection mechanism dynamically reduces heat output

Overheating Problem

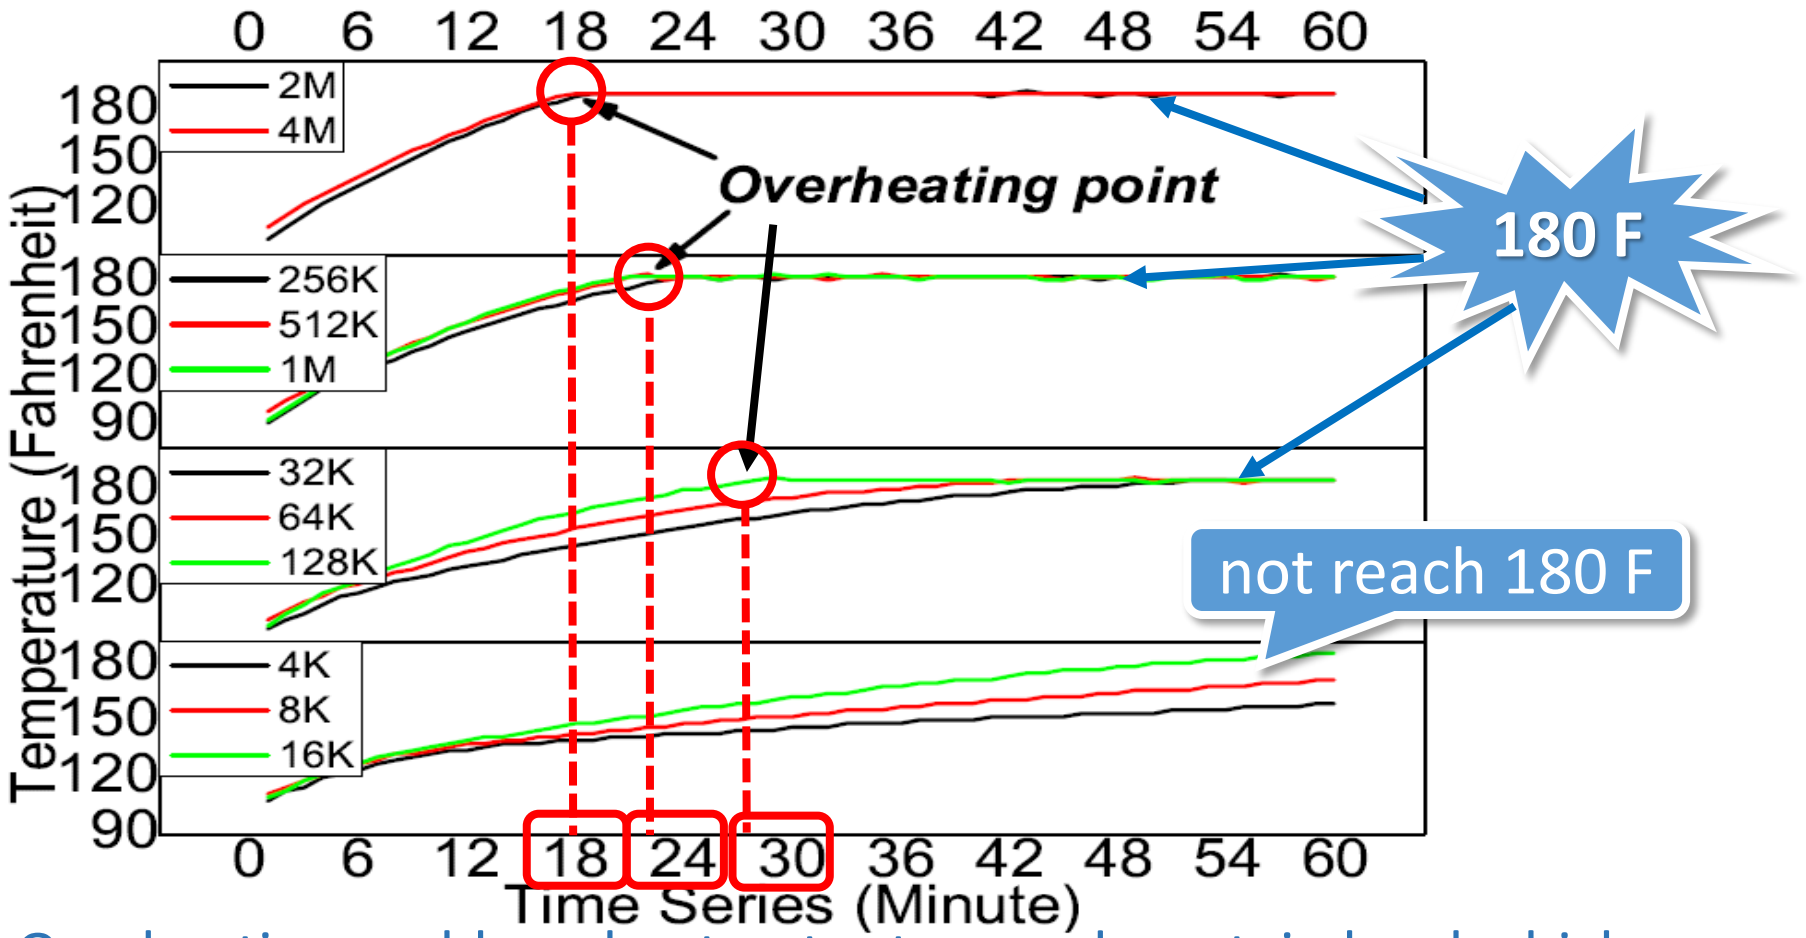

- Overheating problem: heat output exceeds certain level which may lead to malfunction of SSDs
- Device-level protection mechanism dynamically reduces heat output
- Devices with larger data sizes reach overheating points earlier
- How is the temperature controlled?

Power Throttling

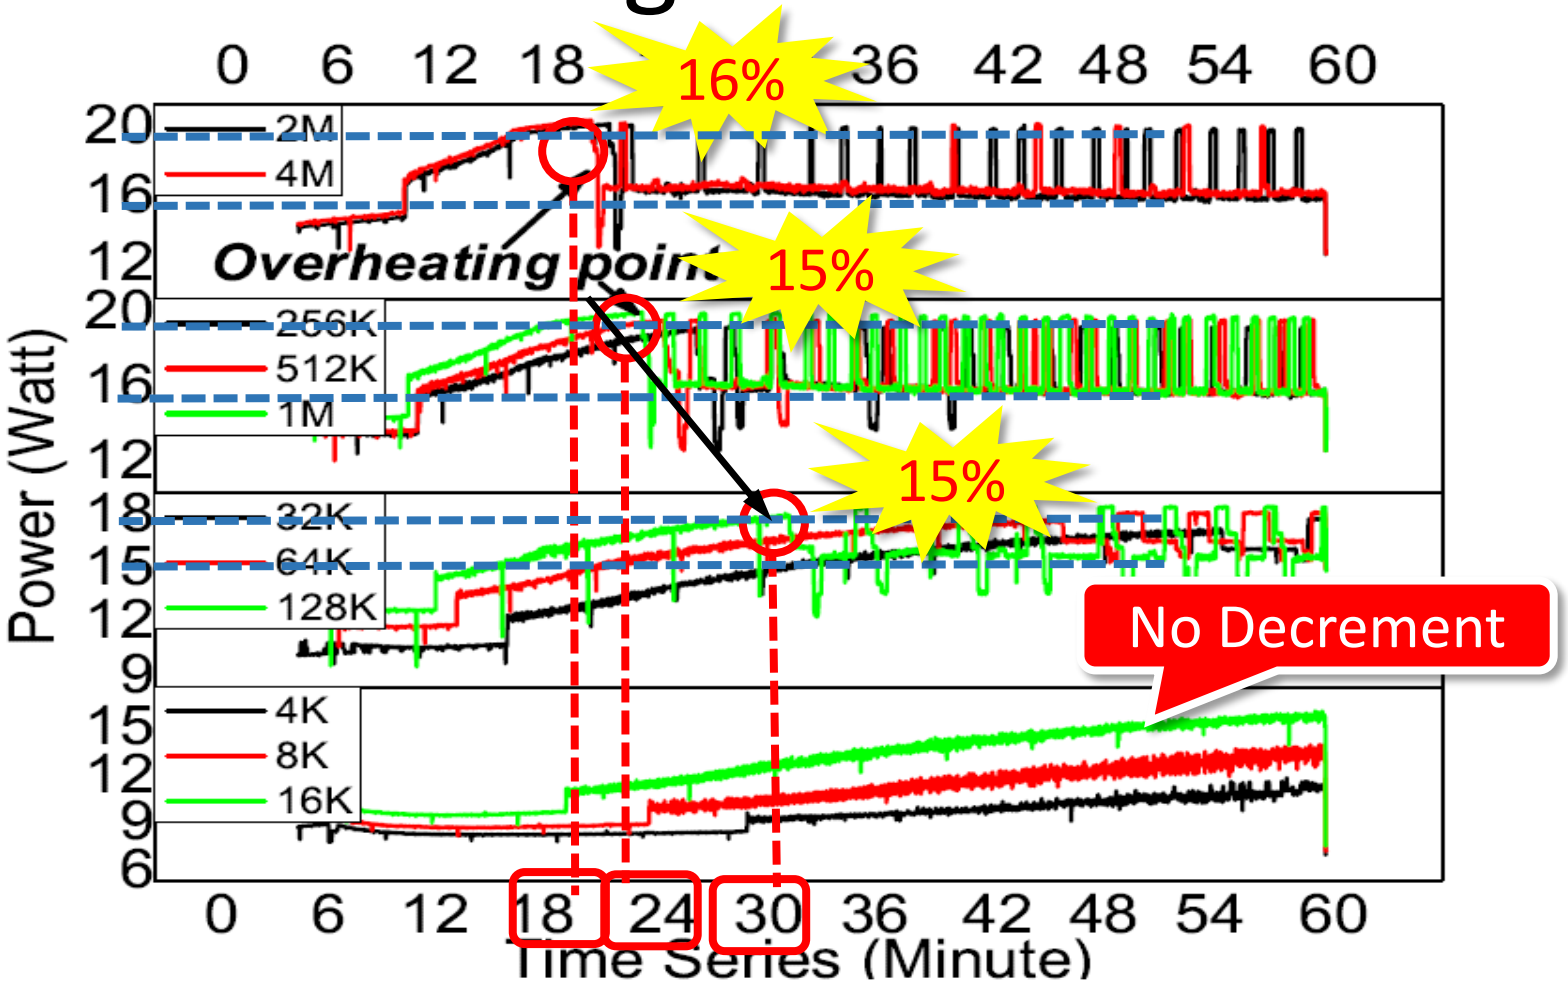

- Power throttling : power is automatically decreased at the overheating points to control heat output
- This decrement is carried out by device-level mechanism

Is simply decreasing power good enough?

Power Throttling's Drawback

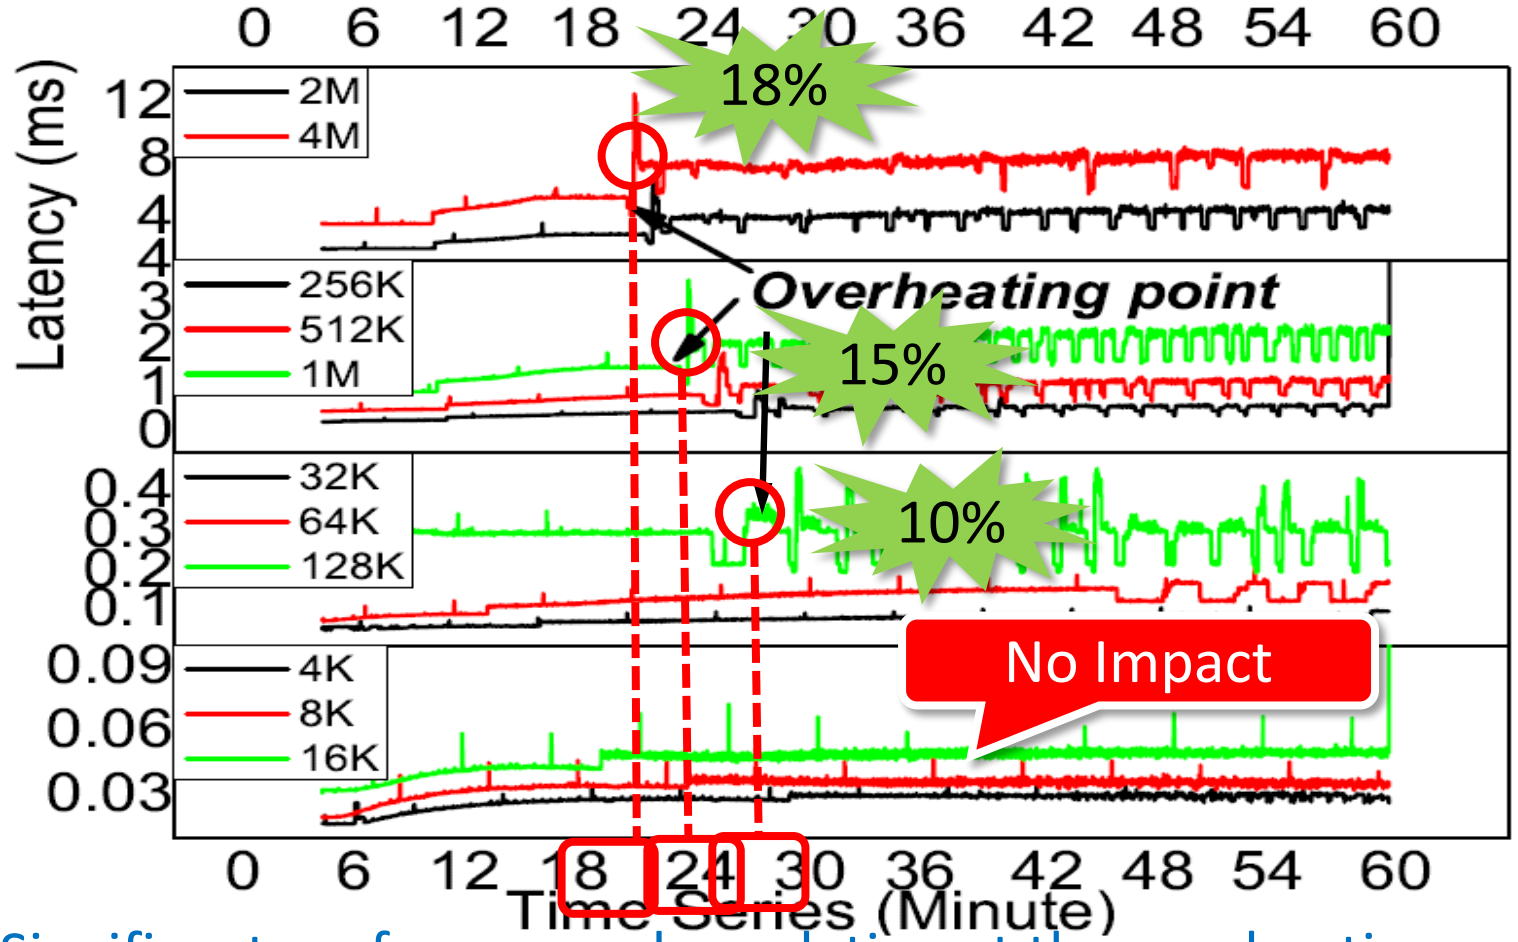

- Significant performance degradation at the overheating points
- Protection mechanism throttles power by reducing active resources
- Reduced power lowers temperature, but also limits active resources
- Overheating problem and power throttling issue hinder SSDs from integrating more resources

Outline

- Motivation
- Evaluation Setup
- Findings
  - Operating temperature
  - Dynamic power behavior
  - Overheating problem and power throttling issue
  - Energy consumption
- Conclusion

Conclusion

Critical Findings:

- MC/MR-SSDs generate higher temperature, consume more power and energy than conventional SSDs, which may not be acceptable for many work conditions
- Overheating problems and power throttling issues are holding back state-of-the-art SSDs from achieving potential performance gains

Future Work:

- HW/SW optimization studies are required to improve power and energy efficiency of SSDs

Thank you!