Large-Scale Study of In-the-Field Flash Failures

Onur Mutlu
omutlu@ethz.ch

(joint work with Justin Meza, Qiang Wu, Sanjeev Kumar)

August 10, 2016

Flash Memory Summit 2016, Santa Clara, CA

Original Paper (I)

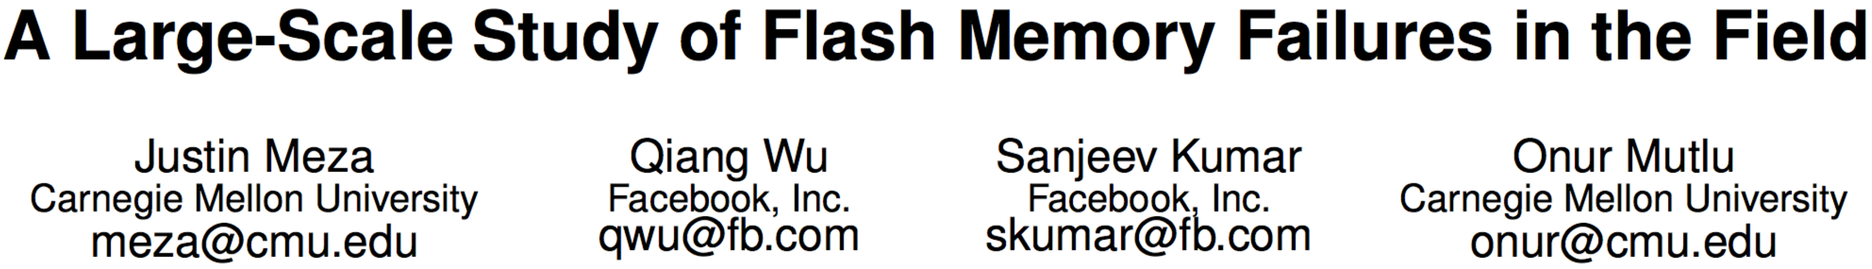

A Large-Scale Study of Flash Memory Failures in the Field

Justin Meza
Carnegie Mellon University
meza@cmu.edu

Qiang Wu
Facebook, Inc.
qwu@fb.com

Sanjeev Kumar
Facebook, Inc.
skumar@fb.com

Onur Mutlu
Carnegie Mellon University
onur@cmu.edu

Original Paper (II)

- Presented at the ACM SIGMETRICS Conference in June 2015.
- Full paper for details:
  - Justin Meza, Qiang Wu, Sanjeev Kumar, and Onur Mutlu, **"A Large-Scale Study of Flash Memory Errors in the Field"** *Proceedings of the ACM International Conference on Measurement and Modeling of Computer Systems (**SIGMETRICS**)*, Portland, OR, June 2015. [Slides (pptx) (pdf)]
  - [Coverage at ZDNet] [Coverage on The Register] [Coverage on TechSpot] [Coverage on The Tech Report]
  - https://users.ece.cmu.edu/~omutlu/pub/flash-memory-failures-in-the-field-at-facebook_sigmetrics15.pdf

A Large-Scale Study of Flash Memory Errors in the Field

Justin Meza
Qiang Wu
Sanjeev Kumar
Onur Mutlu

facebook
Carnegie Mellon University

Overview

First study of flash reliability:

- at a large scale
- in the field

Overview

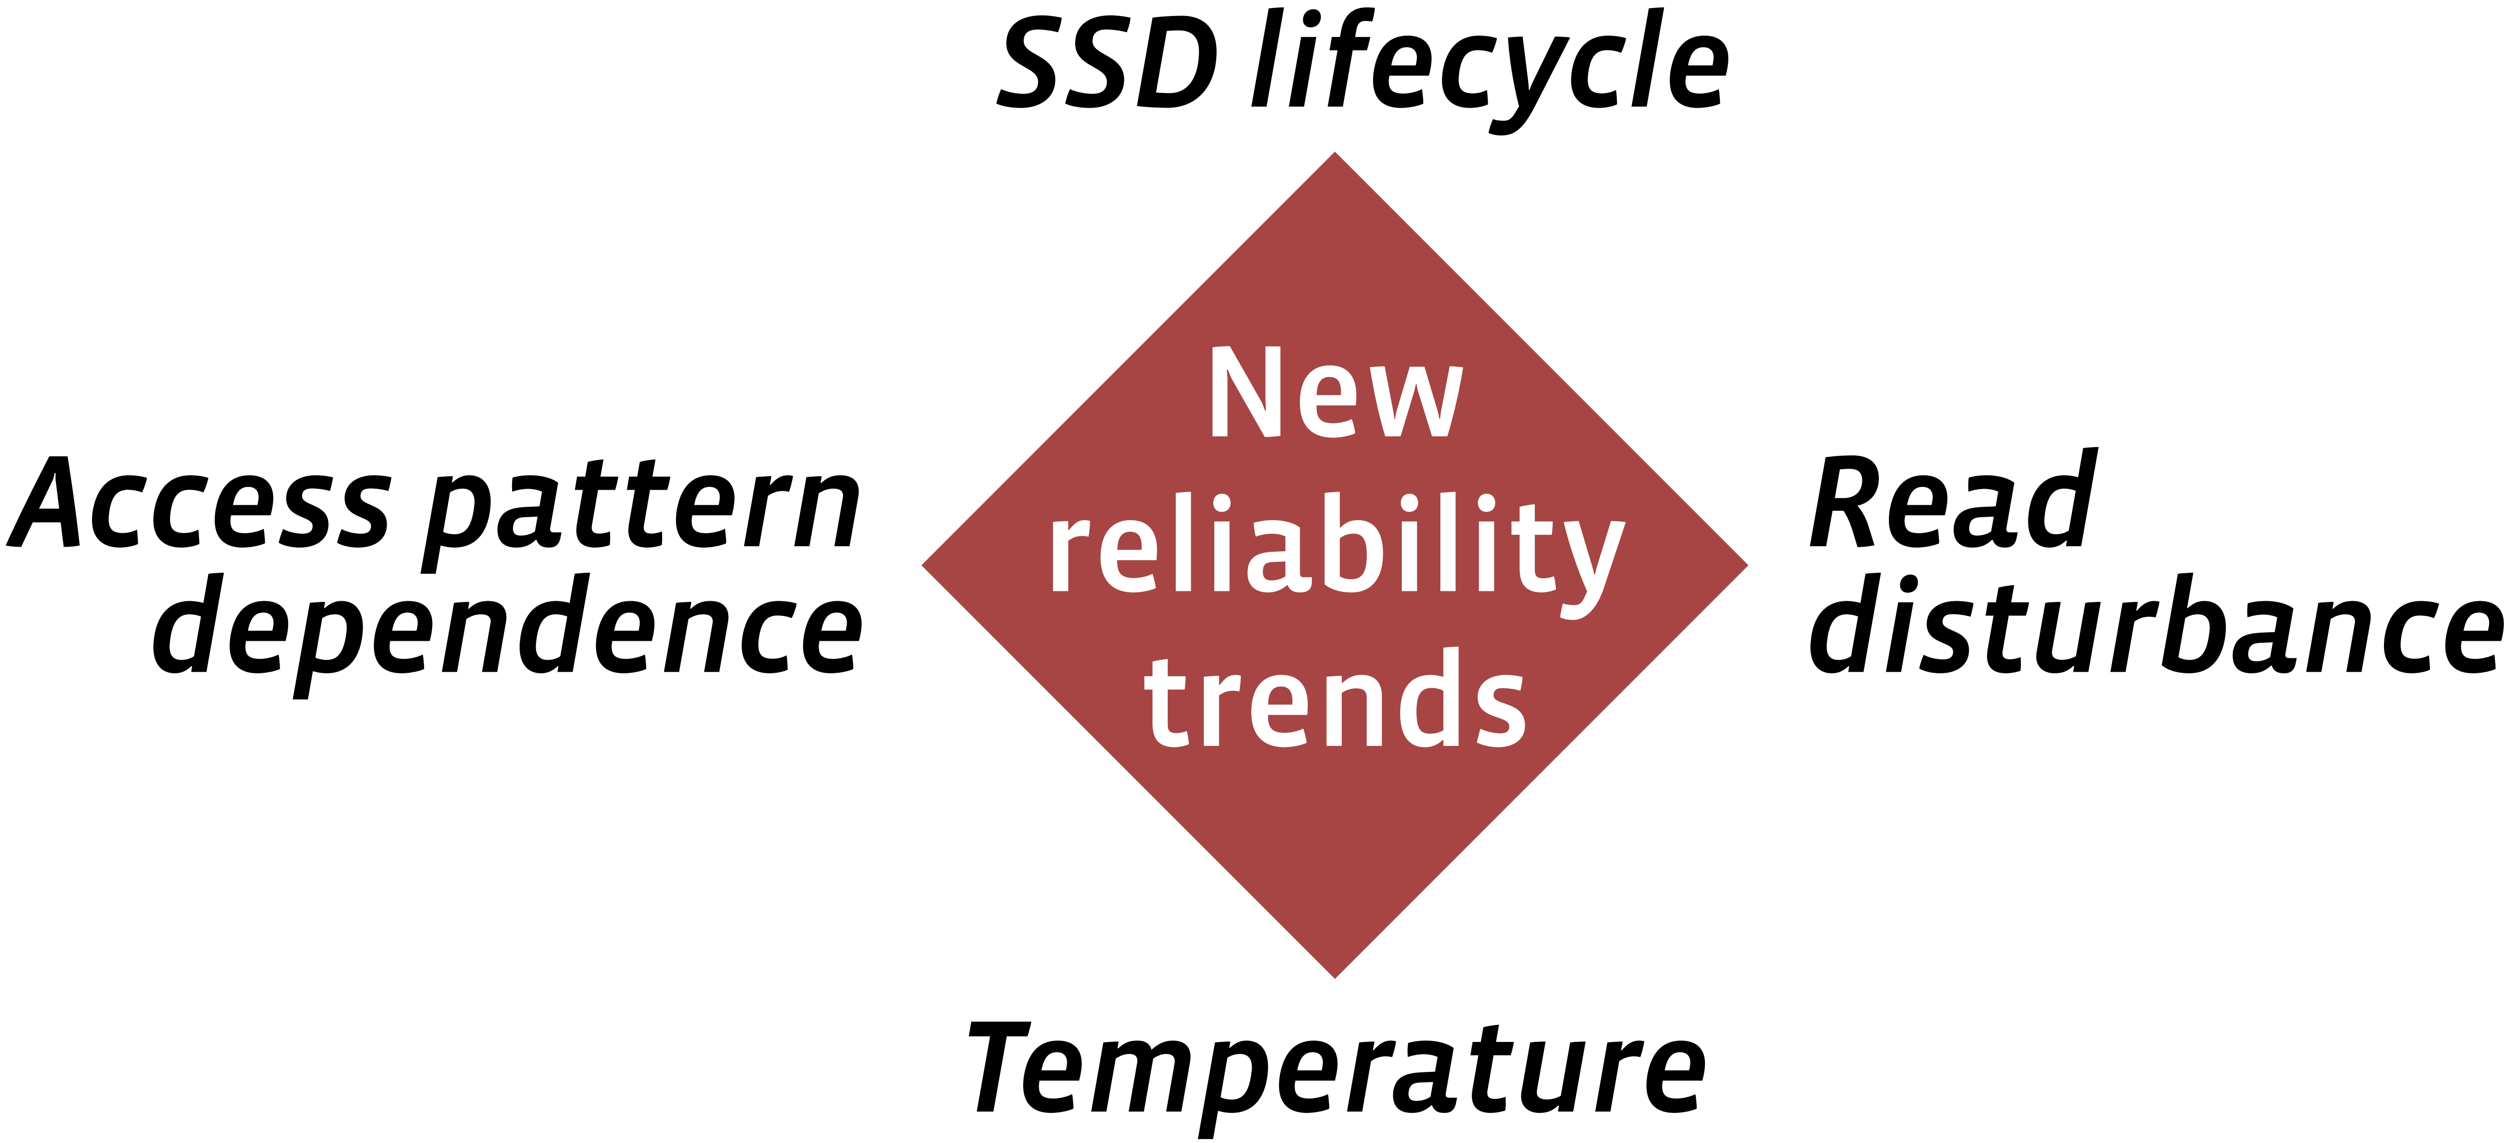

SSD lifecycle

Access pattern dependence

New reliability trends

Read disturbance

Temperature

Overview

SSD lifecycle

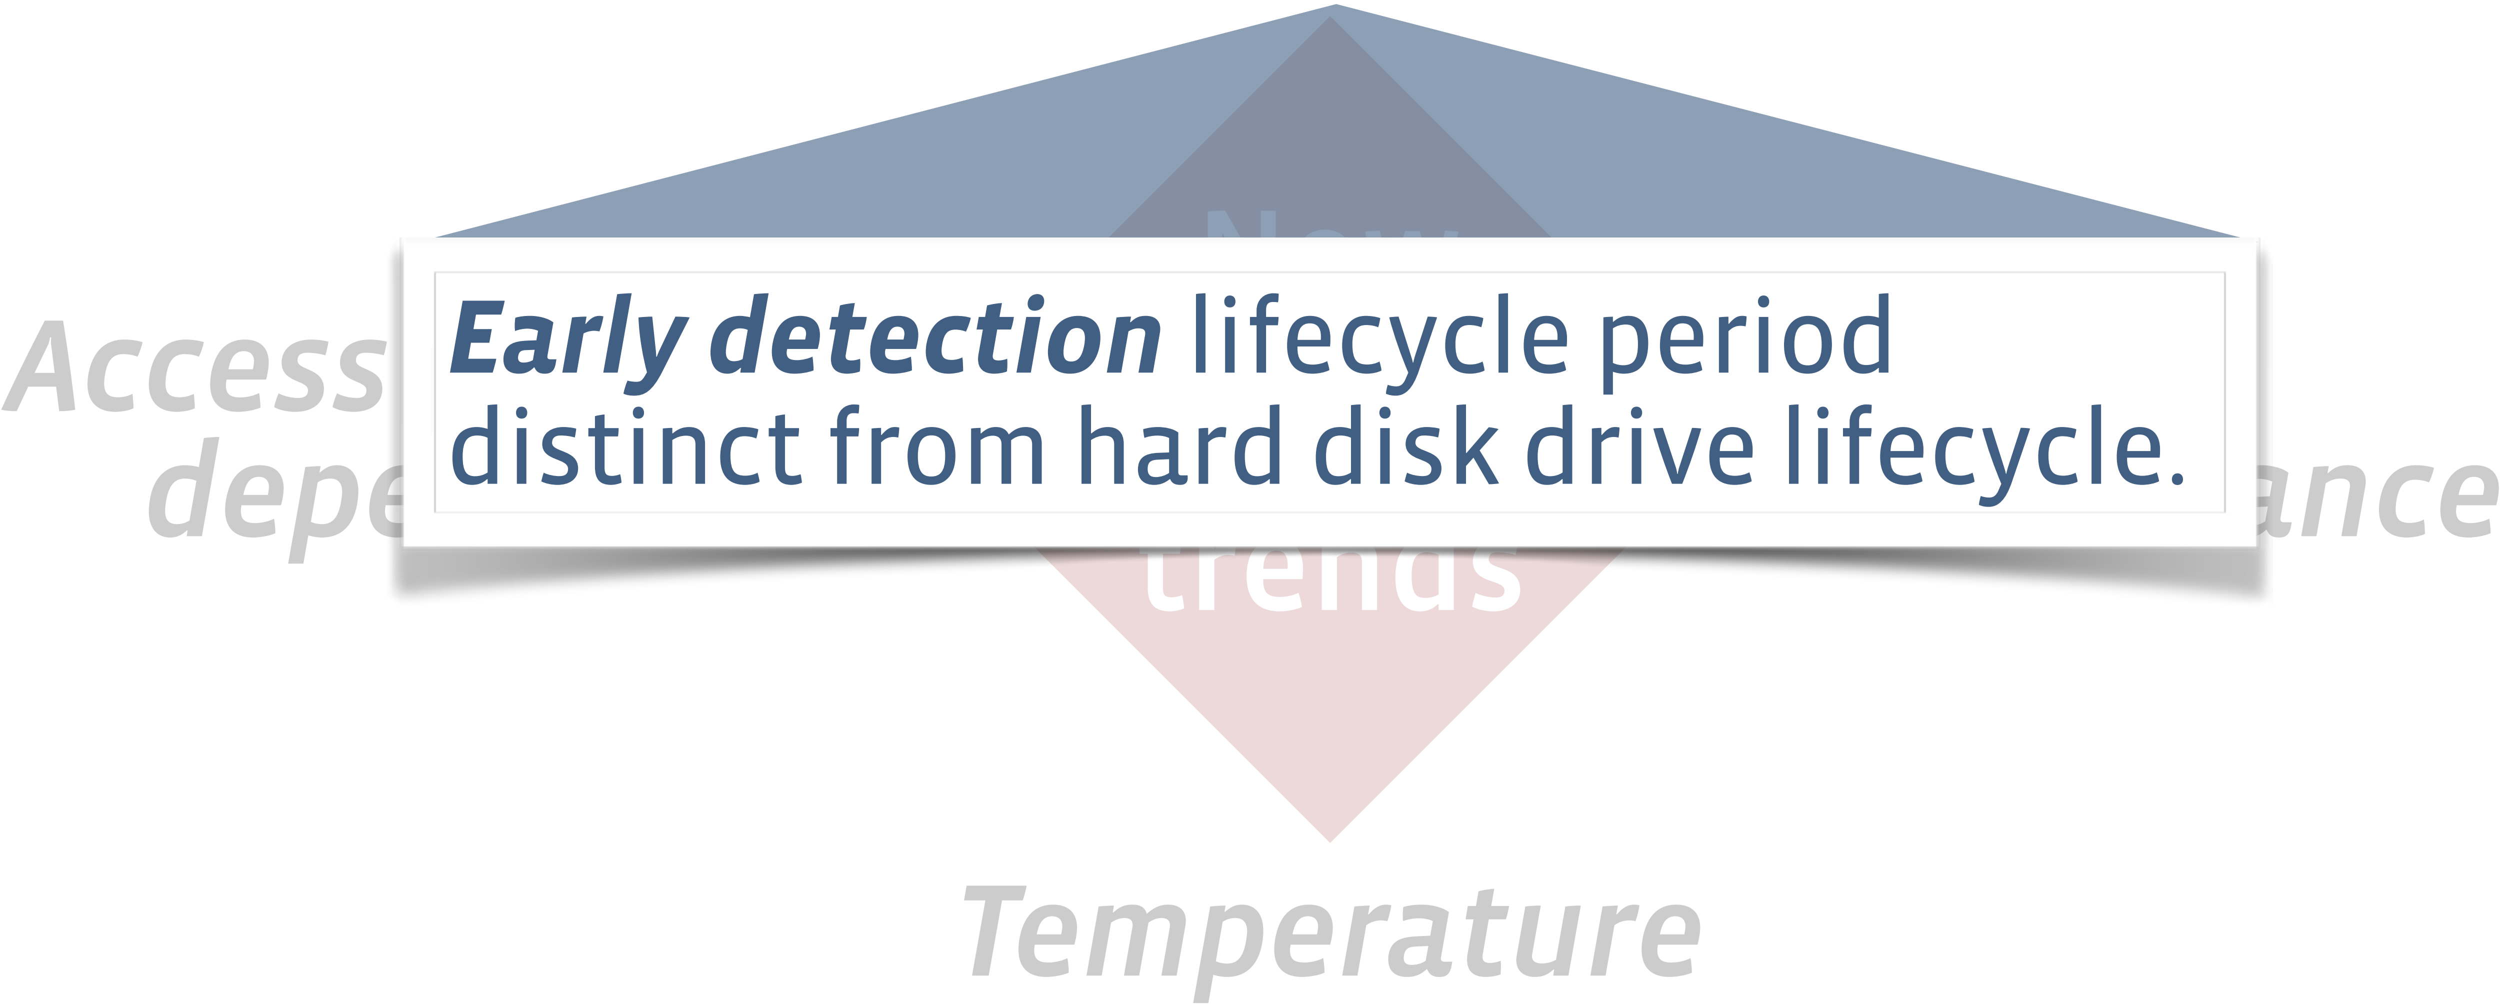

Early detection lifecycle period distinct from hard disk drive lifecycle.

Overview

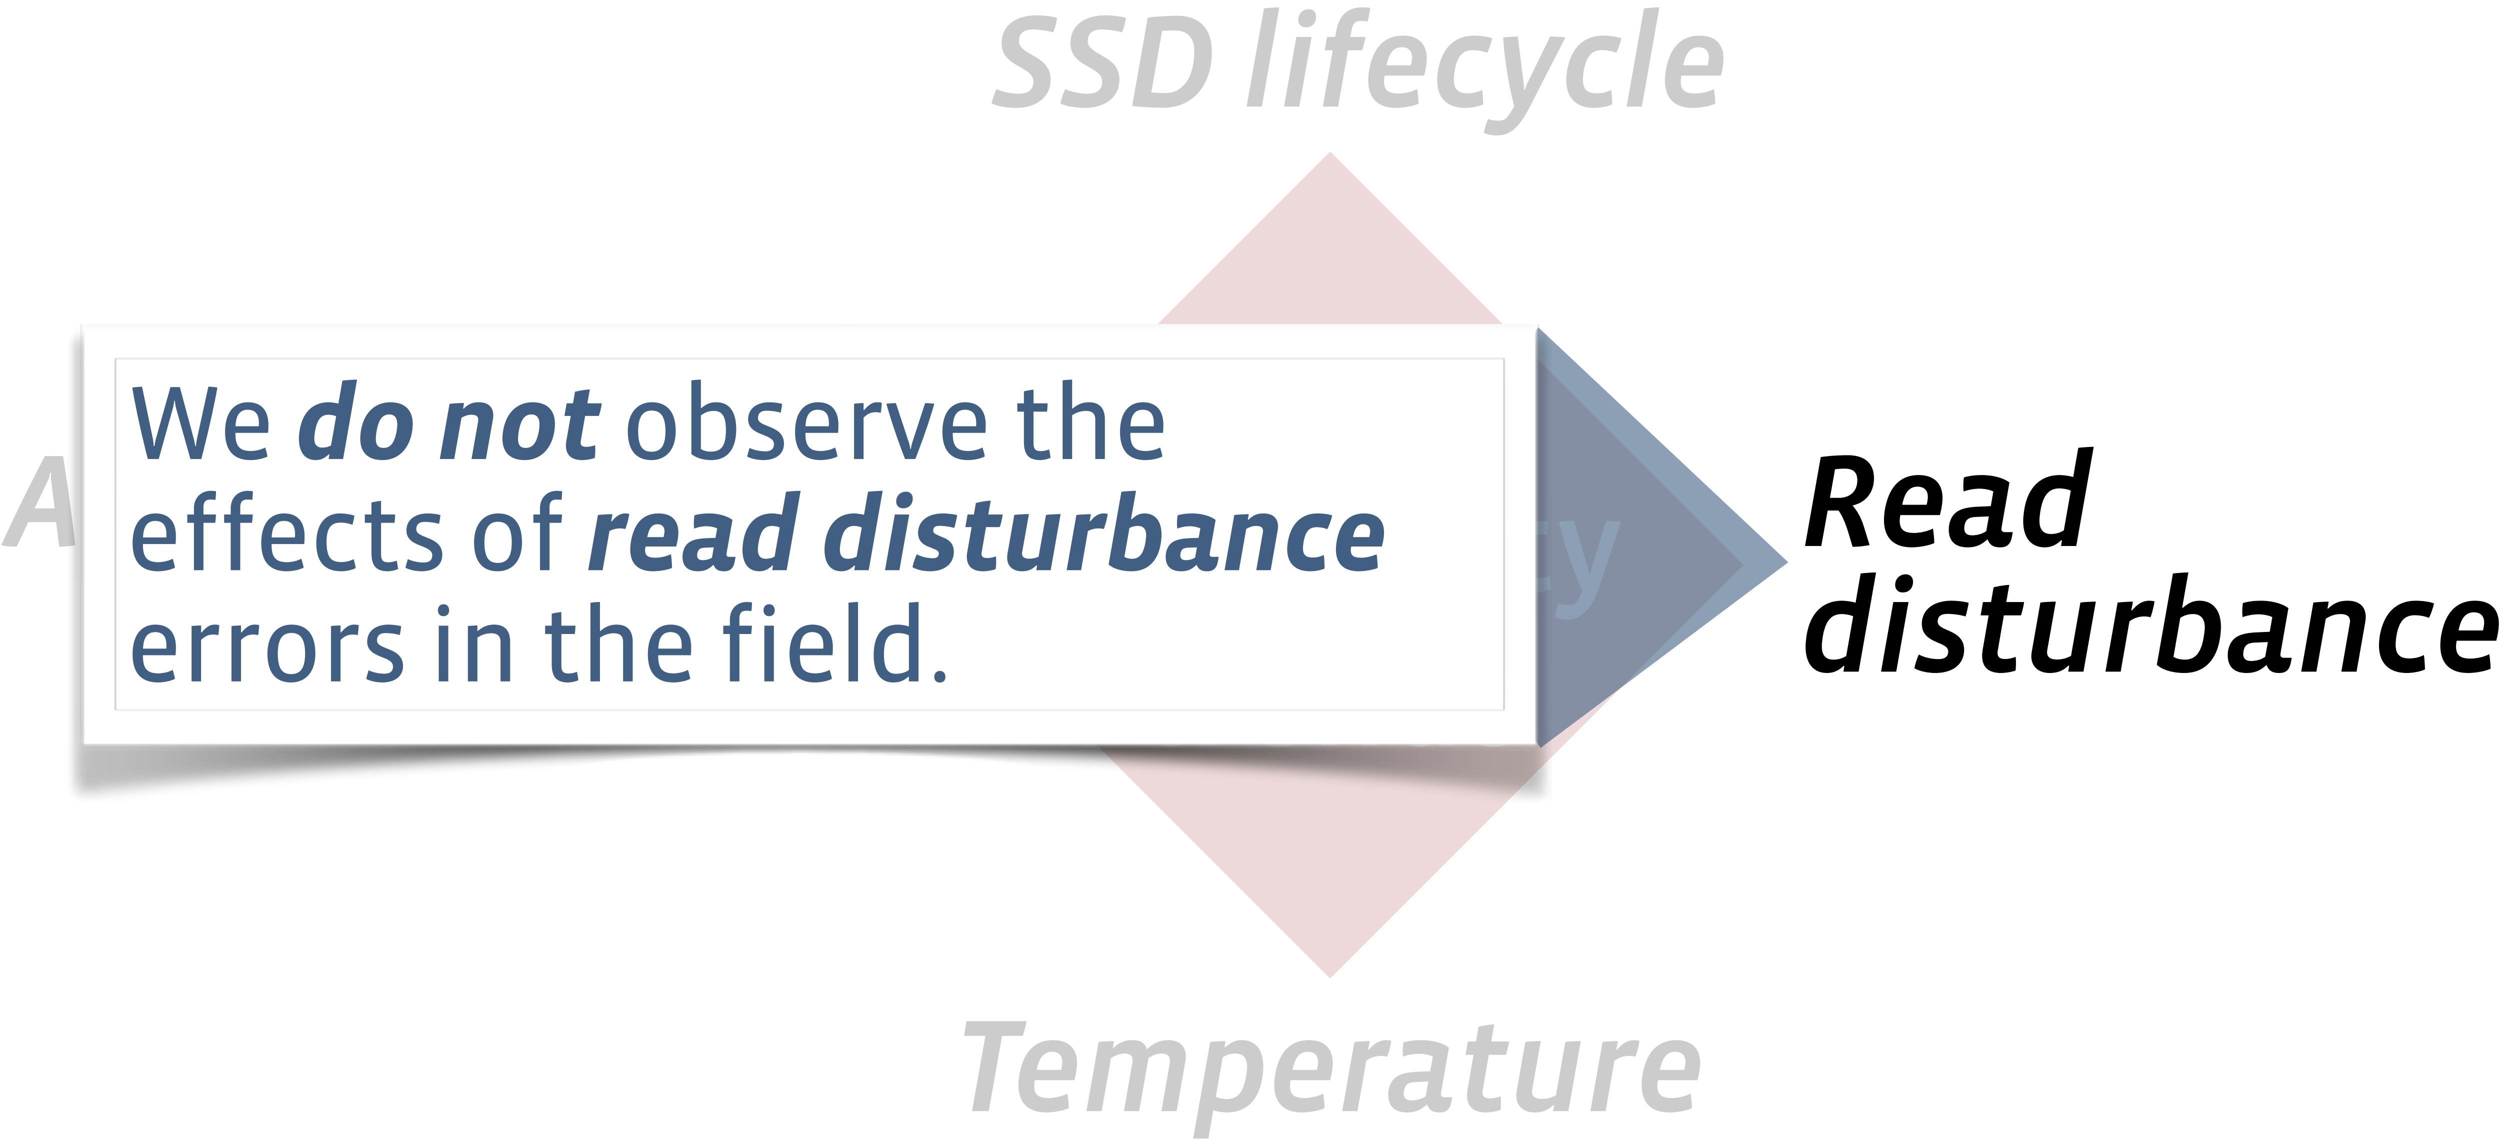

We ***do not*** observe the effects of ***read disturbance*** errors in the field.

Overview

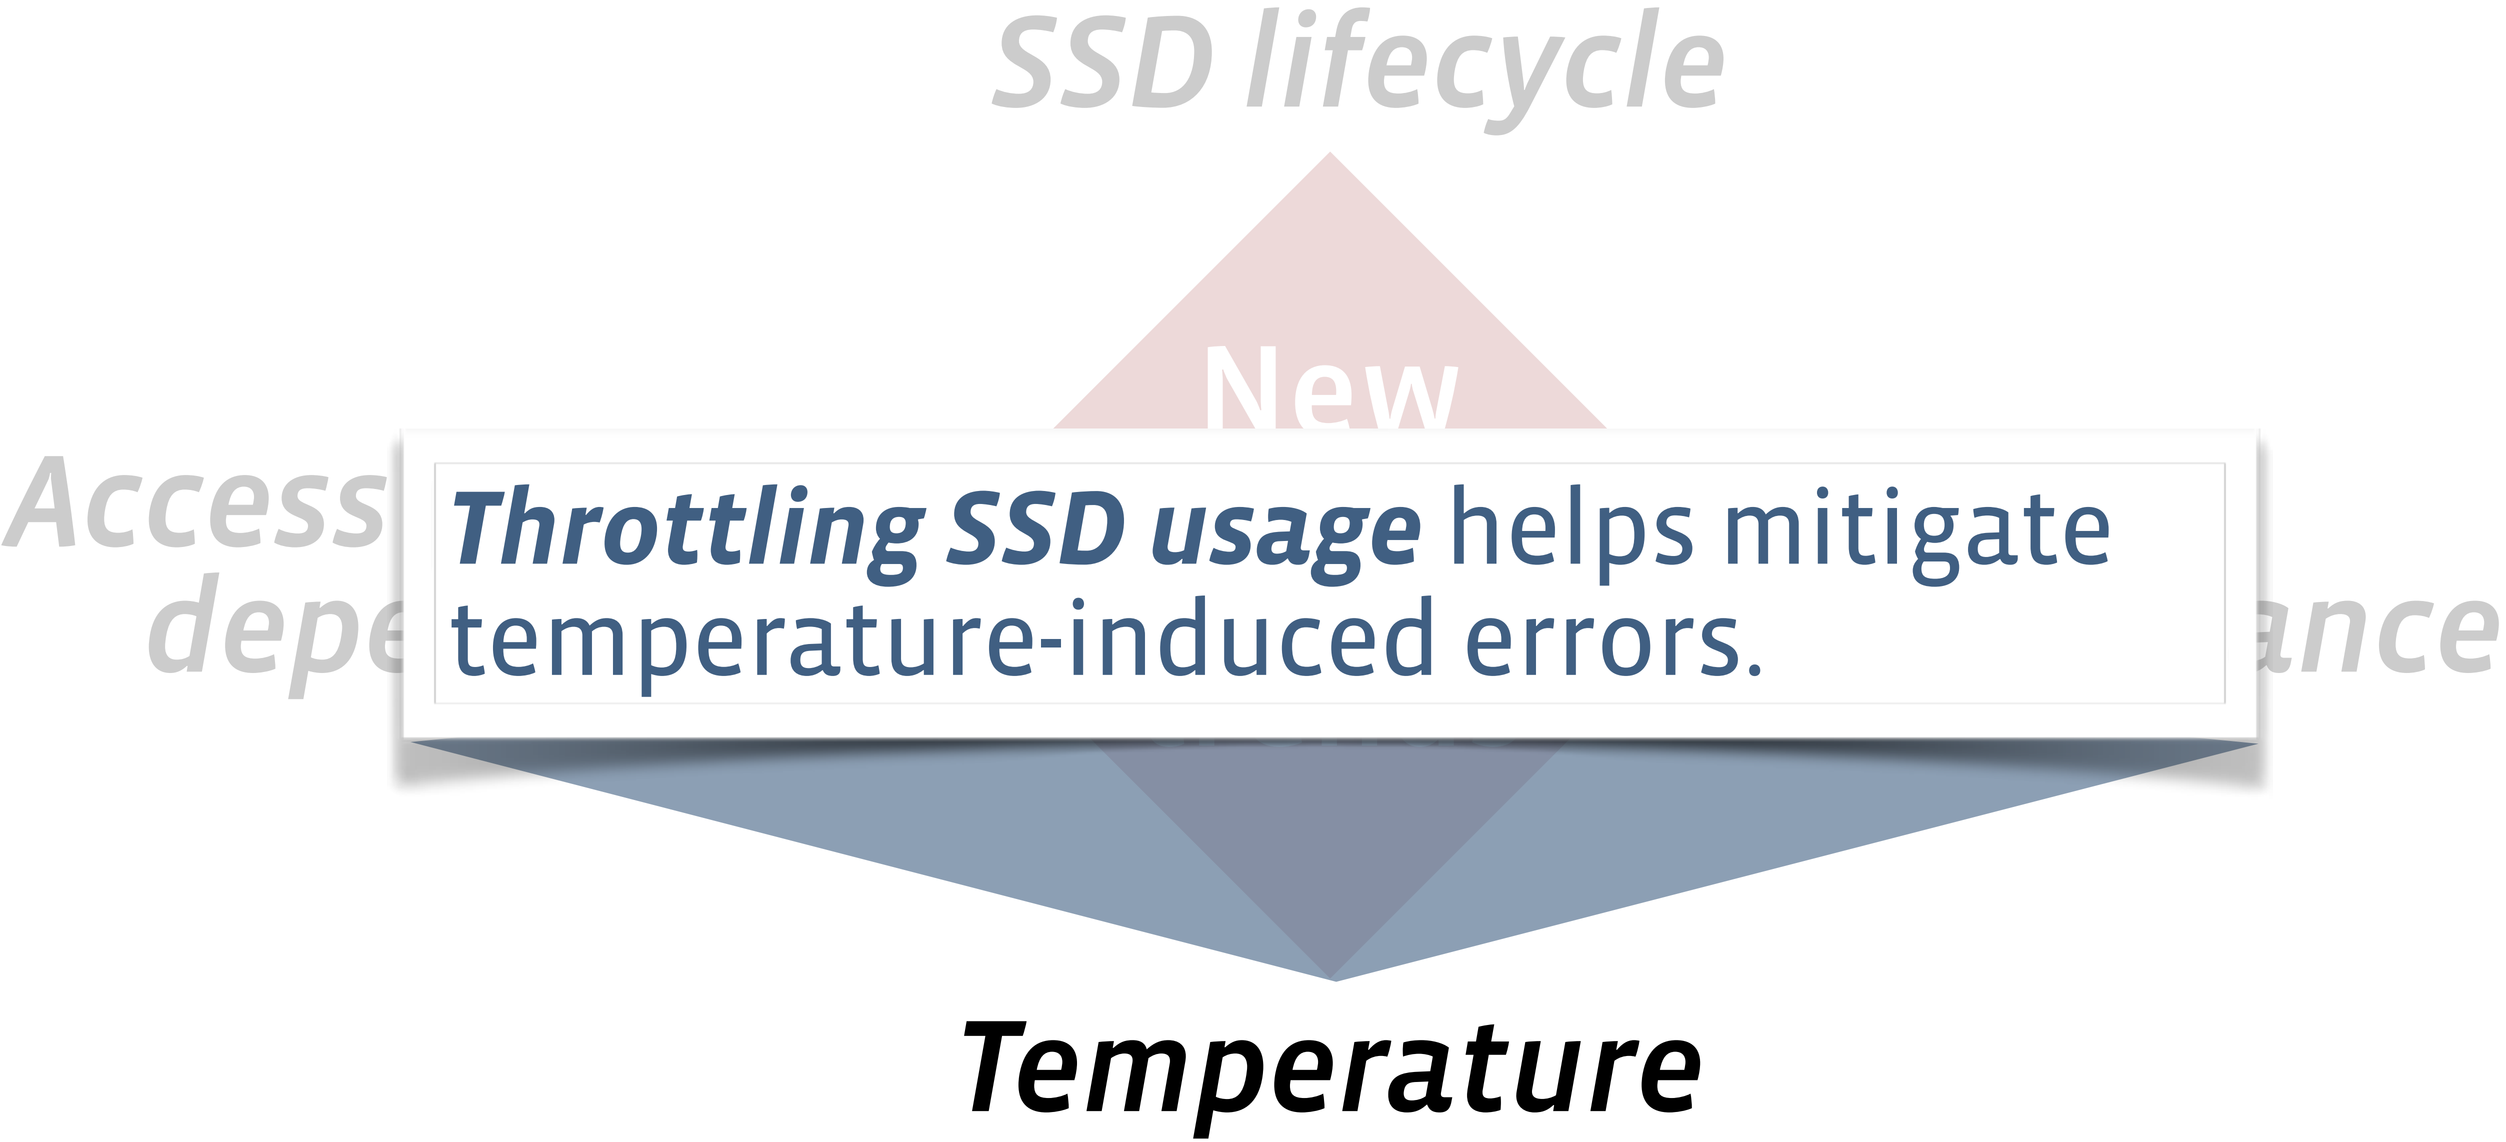

Throttling SSD usage helps mitigate temperature-induced errors.

Overview

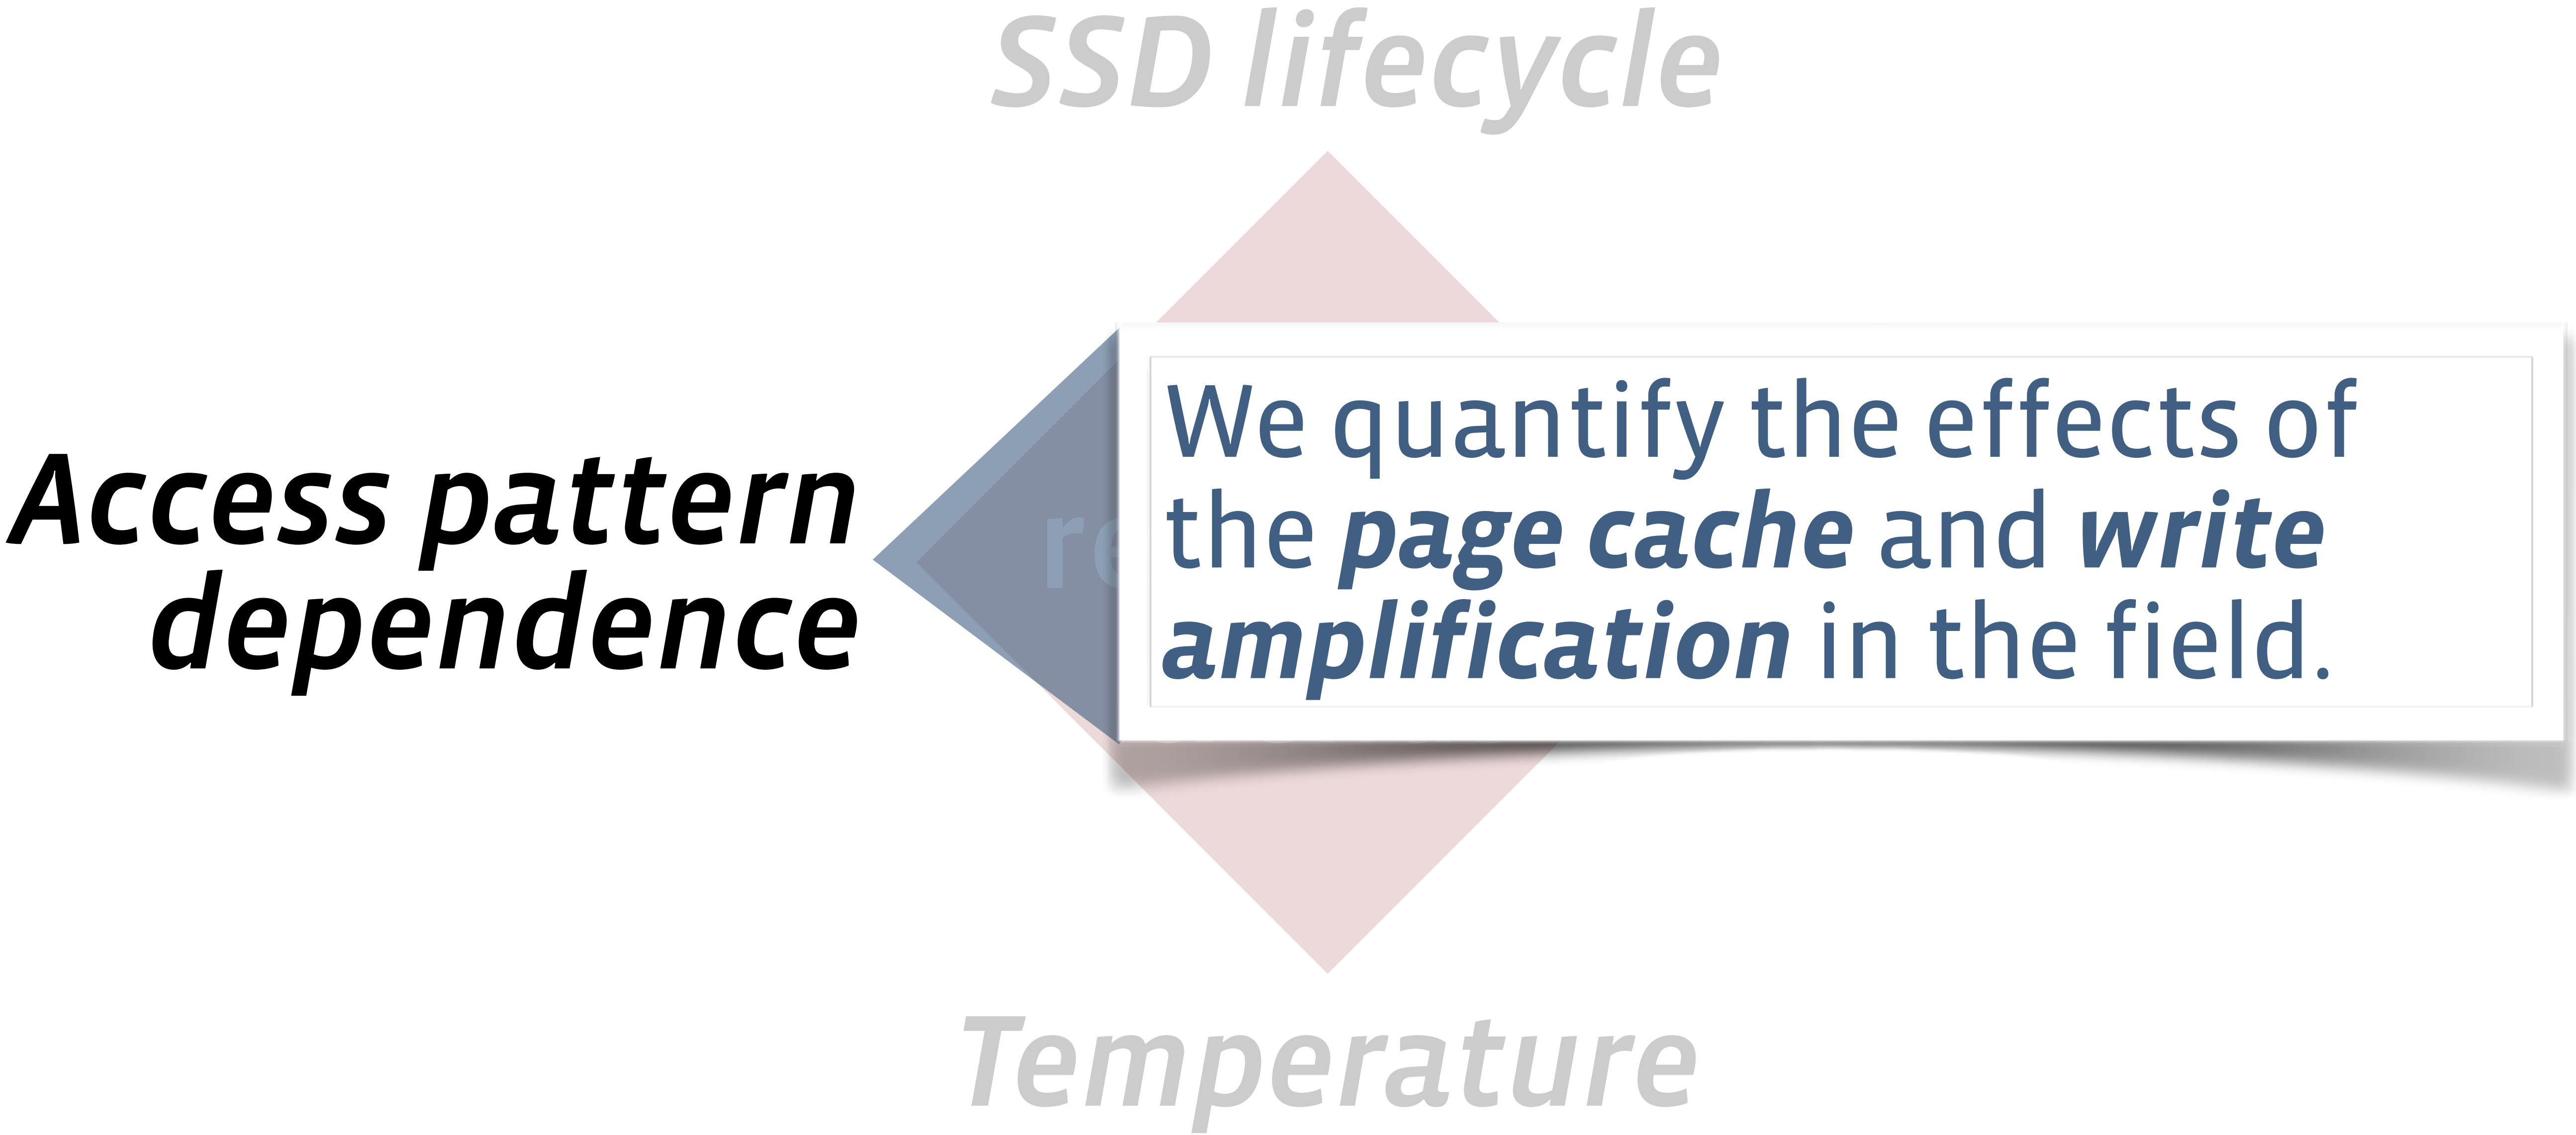

SSD lifecycle

Access pattern dependence

We quantify the effects of the ***page cache*** and ***write amplification*** in the field.

Temperature

Outline

- background and motivation
- server SSD architecture
- error collection/analysis methodology
- SSD reliability trends
- summary

Background and motivation

Flash memory

- persistent
- high performance
- hard disk alternative
- used in solid-state drives (SSDs)

Flash memory

- persistent
- high performance
- hard disk alternative
- used in solid-state drives (SSDs)
- **prone to a variety of errors**
  - wearout, disturbance, retention

Prior Flash Error Studies (I)

1. Overall flash error analysis

- Yu Cai, Erich F. Haratsch, Onur Mutlu, and Ken Mai,
  Error Patterns in MLC NAND Flash Memory: Measurement, Characterization, and Analysis, DATE 2012.

- Yu Cai, Gulay Yalcin, Onur Mutlu, Erich F. Haratsch, Adrian Cristal, Osman Unsal, and Ken Mai,
  Error Analysis and Retention-Aware Error Management for NAND Flash Memory, Intel Technology Journal 2013.

2. Program and erase cycling noise analysis

- Yu Cai, Erich F. Haratsch, Onur Mutlu, and Ken Mai,
  Threshold Voltage Distribution in MLC NAND Flash Memory: Characterization, Analysis and Modeling, DATE 2013.

Prior Flash Error Studies (II)

3. Retention noise analysis and management

- Yu Cai, Gulay Yalcin, Onur Mutlu, Erich F. Haratsch, Adrian Cristal, Osman Unsal, and Ken Mai,
  Flash Correct-and-Refresh: Retention-Aware Error Management for Increased Flash Memory Lifetime, ICCD 2012.

- Yu Cai, Yixin Luo, Erich F. Haratsch, Ken Mai, and Onur Mutlu,
  Data Retention in MLC NAND Flash Memory: Characterization, Optimization and Recovery, HPCA 2015.

- Yixin Luo, Yu Cai, Saugata Ghose, Jongmoo Choi, and Onur Mutlu,
  WARM: Improving NAND Flash Memory Lifetime with Write-hotness Aware Retention Management, MSST 2015.

Prior Flash Error Studies (III)

4. Cell-to-cell interference analysis and management

- Yu Cai, Onur Mutlu, Erich F. Haratsch, and Ken Mai, **Program Interference in MLC NAND Flash Memory: Characterization, Modeling, and Mitigation**, ICCD 2013.

- Yu Cai, Gulay Yalcin, Onur Mutlu, Erich F. Haratsch, Osman Unsal, Adrian Cristal, and Ken Mai, **Neighbor-Cell Assisted Error Correction for MLC NAND Flash Memories**, SIGMETRICS 2014.

5. Read disturb noise study

- Yu Cai, Yixin Luo, Saugata Ghose, Erich F. Haratsch, Ken Mai, and Onur Mutlu, **Read Disturb Errors in MLC NAND Flash Memory: Characterization and Mitigation**, DSN 2015.

Some Prior Talks on Flash Errors

- Saugata Ghose, *Write-hotness Aware Retention Management*, FMS 2016.
- Onur Mutlu, *Read Disturb Errors in MLC NAND Flash Memory*, FMS 2015.
- Yixin Luo, *Data Retention in MLC NAND Flash Memory*, FMS 2015.
- Onur Mutlu,
  Error Analysis and Management for MLC NAND Flash Memory, FMS 2014.

- FMS 2016 posters:
  - WARM: Improving NAND Flash Memory Lifetime with Write-hotness Aware Retention Management
  - Read Disturb Errors in MLC NAND Flash Memory
  - Data Retention in MLC NAND Flash Memory

Prior Works on Flash Error Analysis

- Controlled studies that experimentally analyze many types of error
  - retention, program interference, read disturb, wear

- Conducted on raw flash chips, not full SSD-based systems

- Use synthetic access patterns, not real workloads in production systems

- Do not account for the storage software stack

- Small number of chips and small amount of time

Prior Lower-Level Flash Error Studies

- Provide a lot of insight

- Lead to new reliability and performance techniques
  - E.g., to manage errors in a controller

- But they do not provide information on
  - errors that appear during real-system operation
  - beyond the correction capability of the controller

In-The-Field Operation Effects

- Access patterns **not** controlled
- Real applications access SSDs over years
- Through the storage software stack (employs buffering)
- Through the SSD controller (employs ECC and wear leveling)
- Factors in platform design (e.g., number of SSDs) can affect access patterns
- Many SSDs and flash chips in a real data center

Our goal

Understand SSD reliability:

- at a large scale
  - millions of device-days, across four years
- in the field
  - realistic workloads and systems

Server SSD architecture

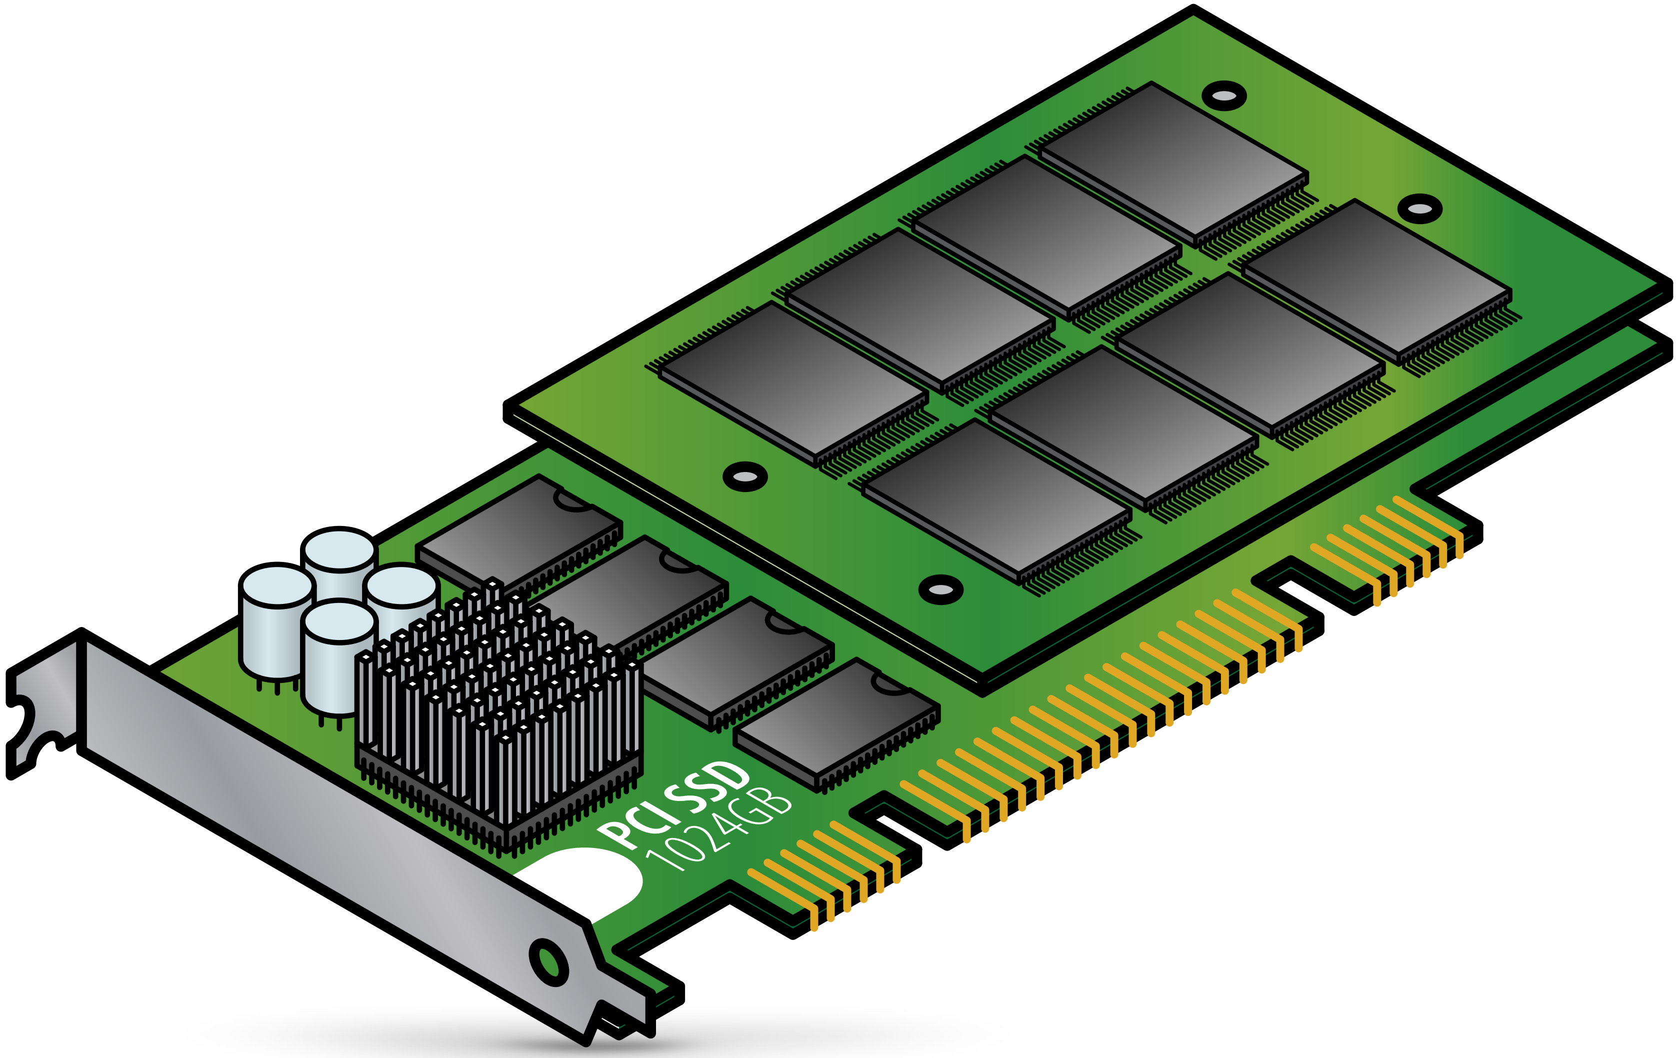
PCI SSD
1024GB

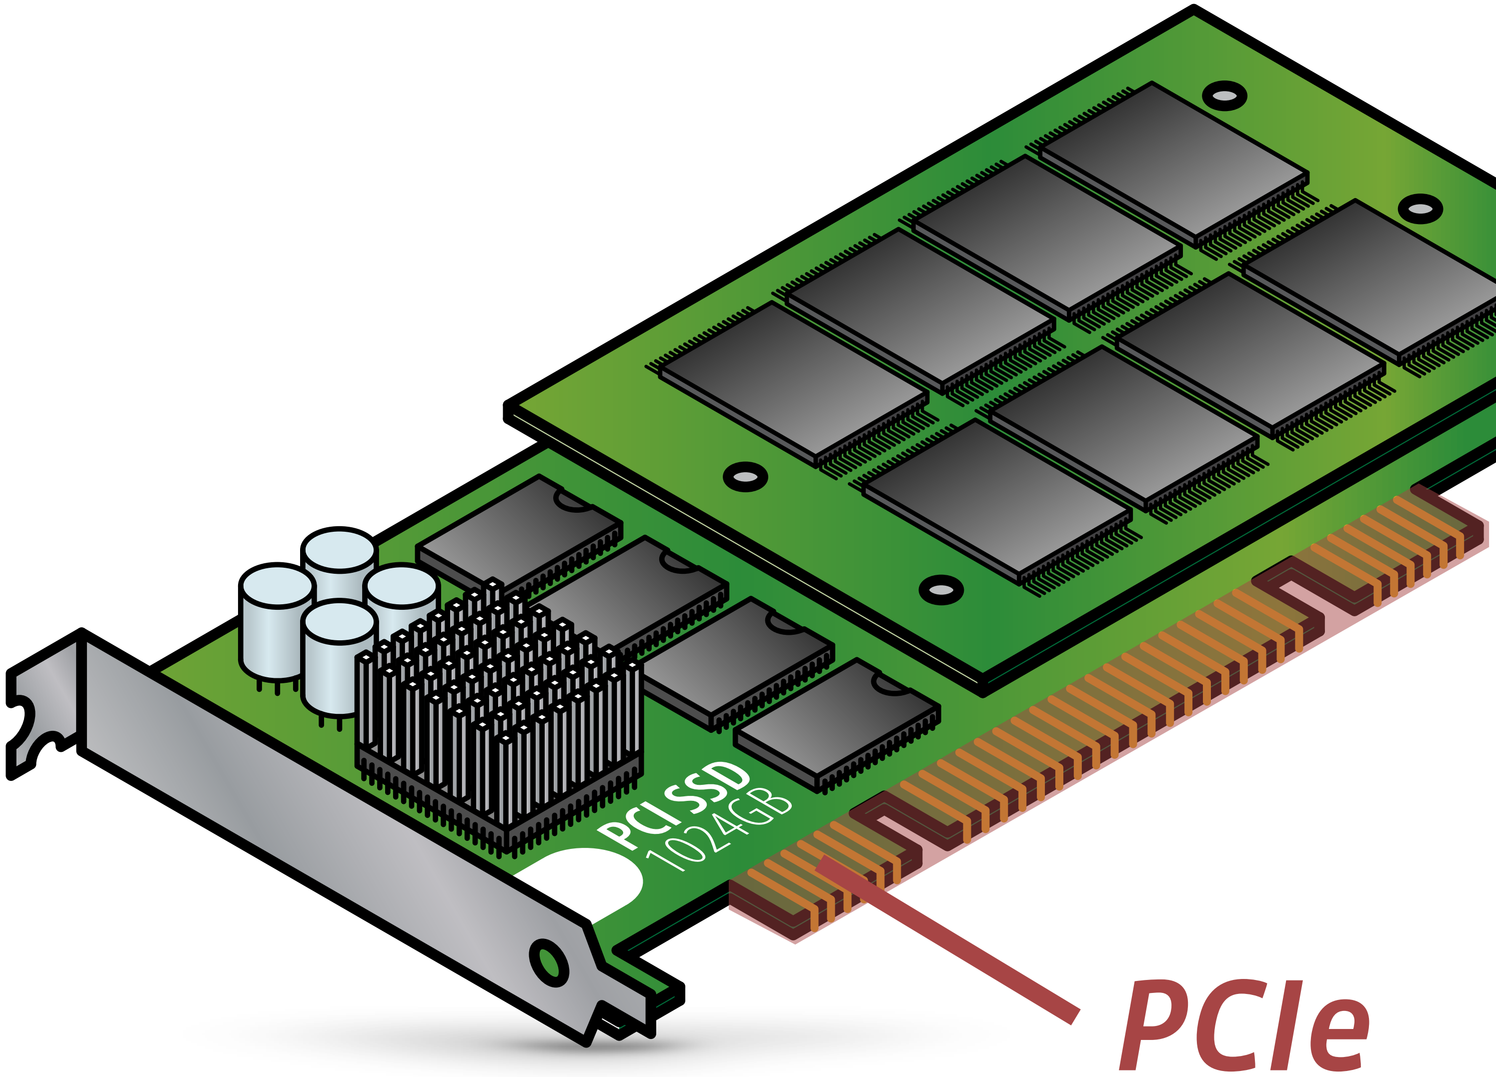
PCI SSD
1024GB
PCIe

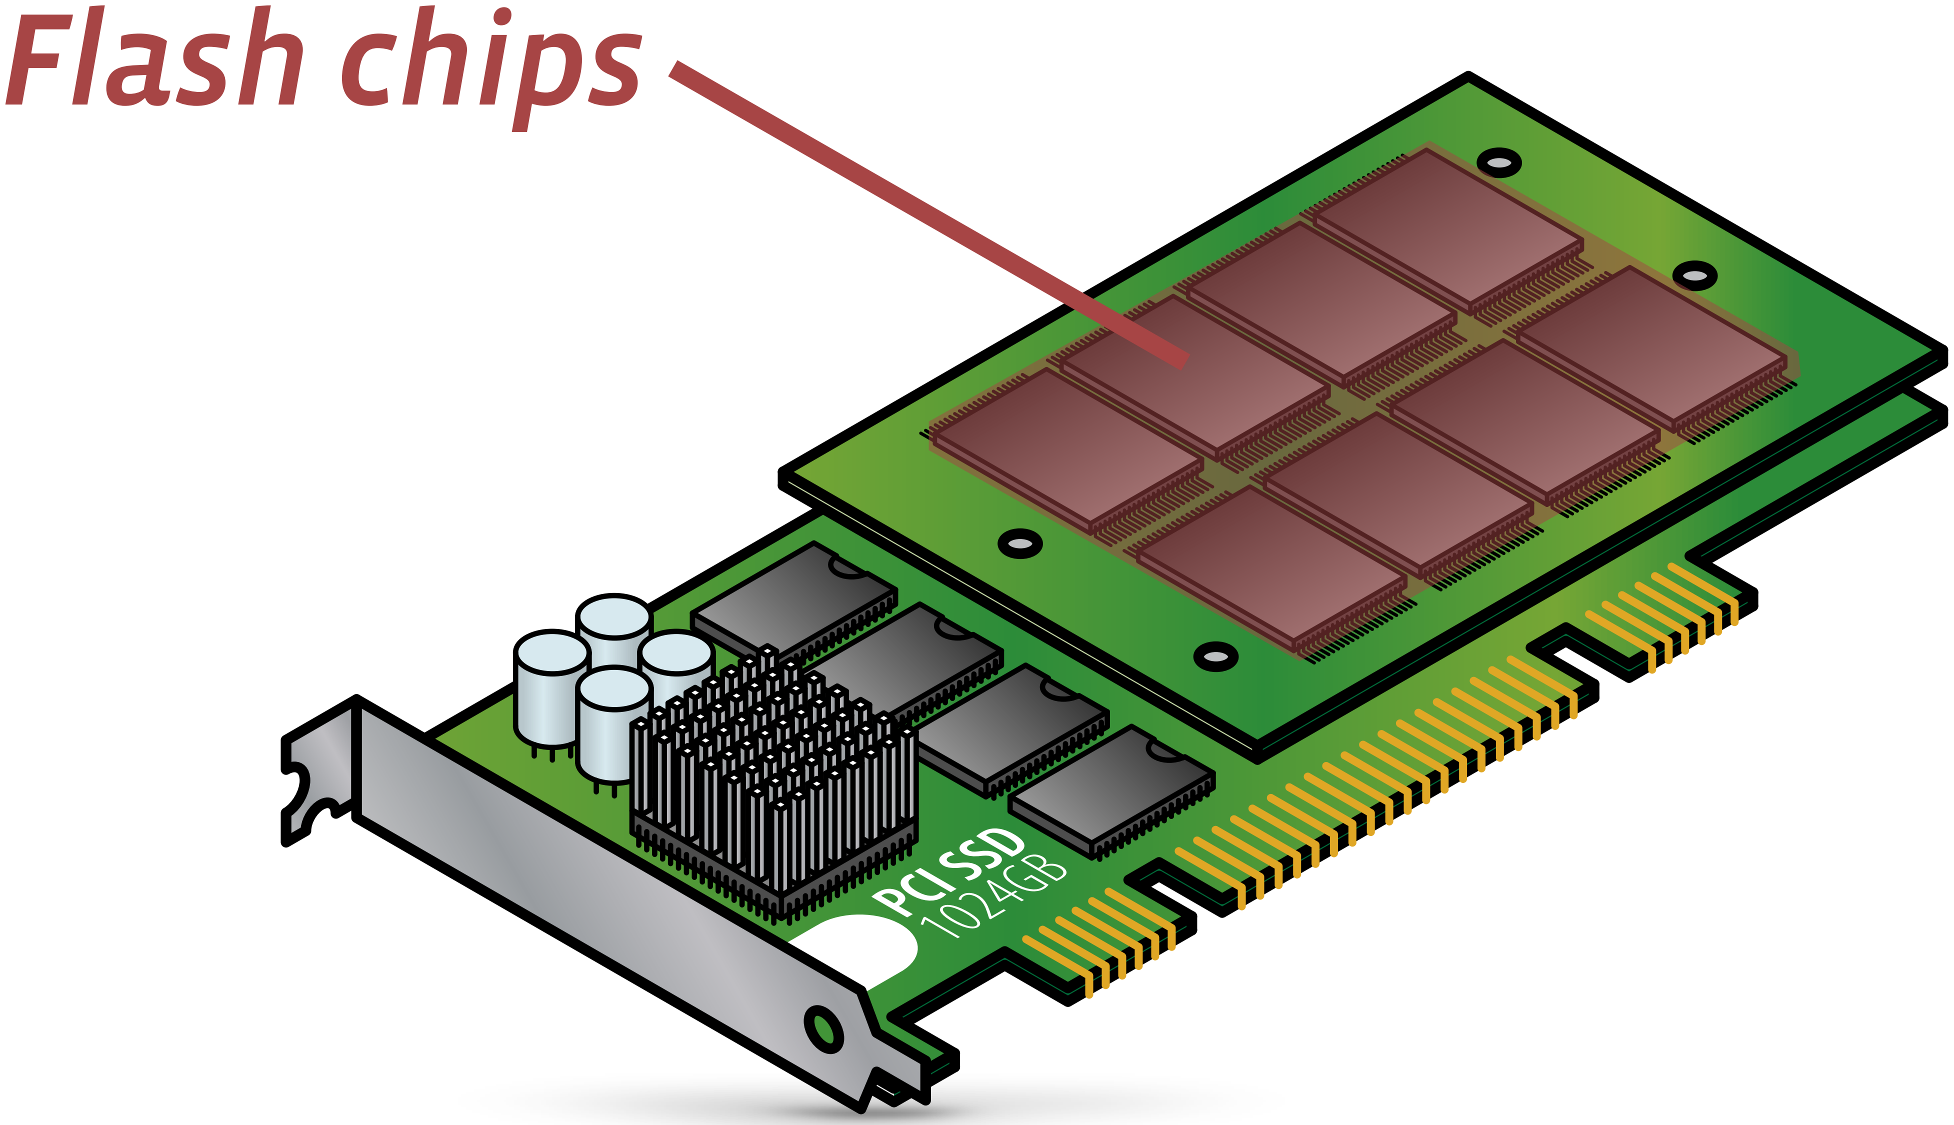
Flash chips
PCI SSD
1024GB

SSD controller

- translates addresses
- schedules accesses
- performs wear leveling

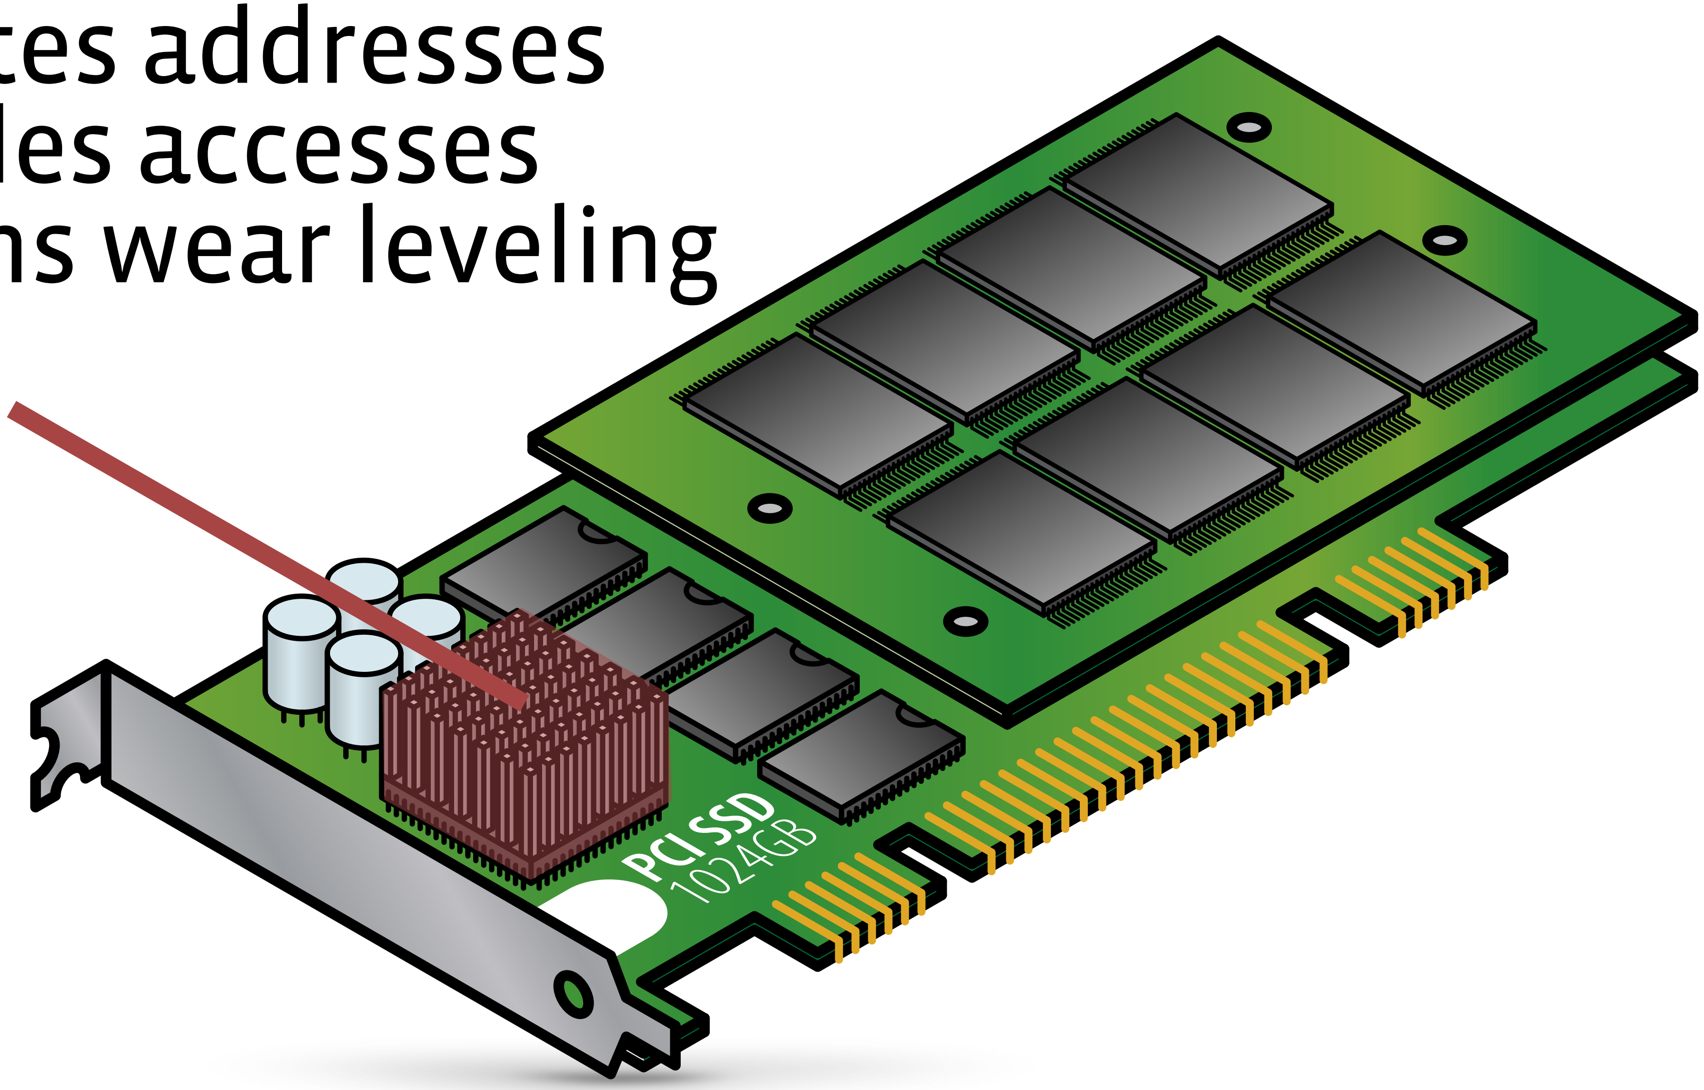

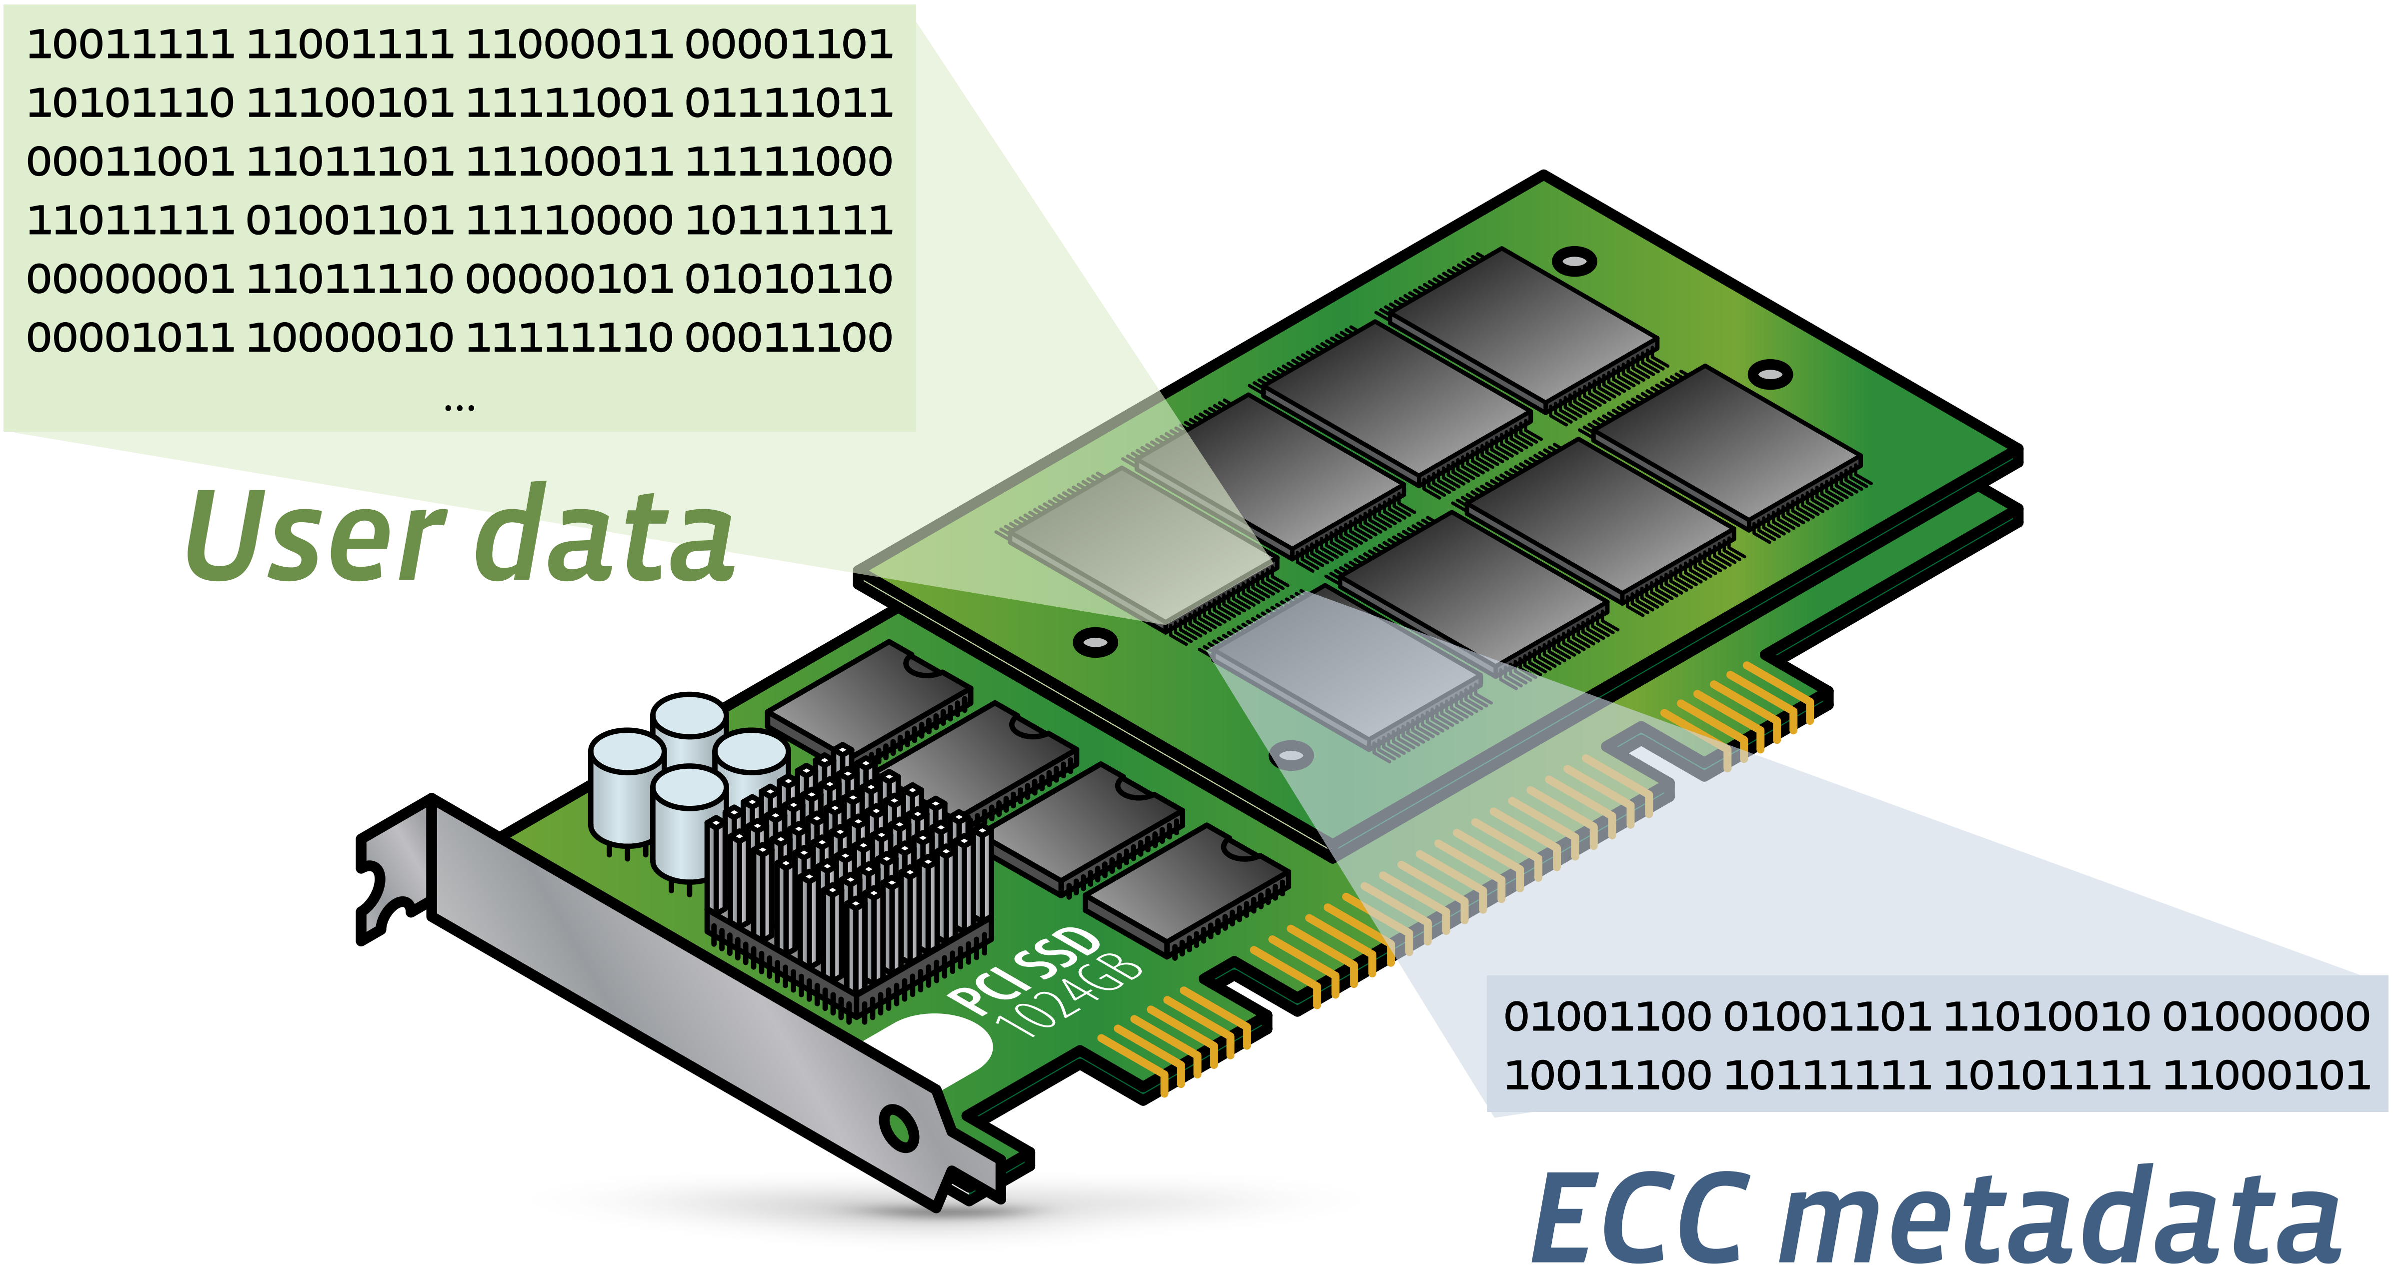
10011111 11001111 11000011 00001101
10101110 11100101 11111001 01111011
00011001 11011101 11100011 11111000
11011111 01001101 11110000 10111111
00000001 11011110 00000101 01010110
00001011 10000010 11111110 00011100
...
User data
PCI SSD
1024GB
01001100 01001101 11010010 01000000
10011100 10111111 10101111 11000101
ECC metadata

Types of errors

Small errors

- 10's of flipped bits per KB
- silently corrected by SSD controller

Large errors

- 100's of flipped bits per KB
- corrected by host using driver
- referred to as SSD failure

Types of errors

Small errors

We examine ***large errors (SSD failures)*** in this study.

er

Large errors

- ~100's of flipped bits per KB
- corrected by host using driver
- refer to as SSD failure

Error collection/ analysis methodology

SSD data measurement

- metrics stored on SSDs
- measured across SSD lifetime

SSD characteristics

- 6 different system configurations
  - 720GB, 1.2TB, and 3.2TB SSDs
  - servers have 1 or 2 SSDs
  - this talk: representative systems
- 6 months to 4 years of operation
- 15TB to 50TB read and written

Platform and SSD Characteristics

- Six different platforms
- Spanning a majority of SSDs at Facebook's production servers

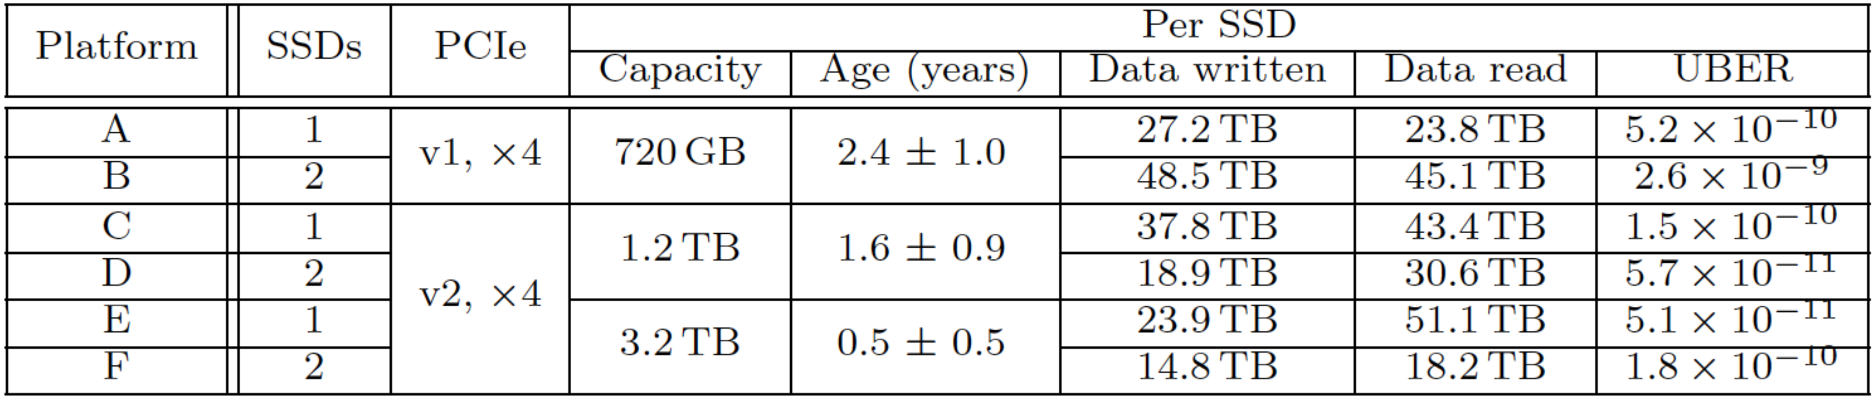

Platform	SSDs	PCIe	Per SSD				
			Capacity	Age (years)	Data written	Data read	UBER
A	1	v1, ×4	720 GB	2.4 ± 1.0	27.2 TB	23.8 TB	5.2×10^{-10}
B	2				48.5 TB	45.1 TB	2.6×10^{-9}
C	1	v2, ×4	1.2 TB	1.6 ± 0.9	37.8 TB	43.4 TB	1.5×10^{-10}
D	2				18.9 TB	30.6 TB	5.7×10^{-11}
E	1		3.2 TB	0.5 ± 0.5	23.9 TB	51.1 TB	5.1×10^{-11}
F	2				14.8 TB	18.2 TB	1.8×10^{-10}

Table 1: The platforms examined in our study. PCIe technology is denoted by vX, ×Y where X = version and Y = number of lanes. Data was collected over the entire age of the SSDs. Data written and data read are to/from the physical storage over an SSD's lifetime. UBER = uncorrectable bit error rate (Section 3.2).

Bit error rates (BER)

- BER = bit errors per bits transmitted
- ***1 error per 385M bits*** transmitted to ***1 error per 19.6B bits*** transmitted
  - averaged across all SSDs in each system type
- ***10x to 1000x lower*** than prior studies
  - large errors, SSD performs wear leveling

Some Definitions

- Uncorrectable error
  - Cannot be corrected by the SSD
  - But corrected by the host CPU driver

- SSD failure rate
  - Fraction of SSDs in a “bucket” that have had at least one uncorrectable error

Different Platforms, Different Failure Rates

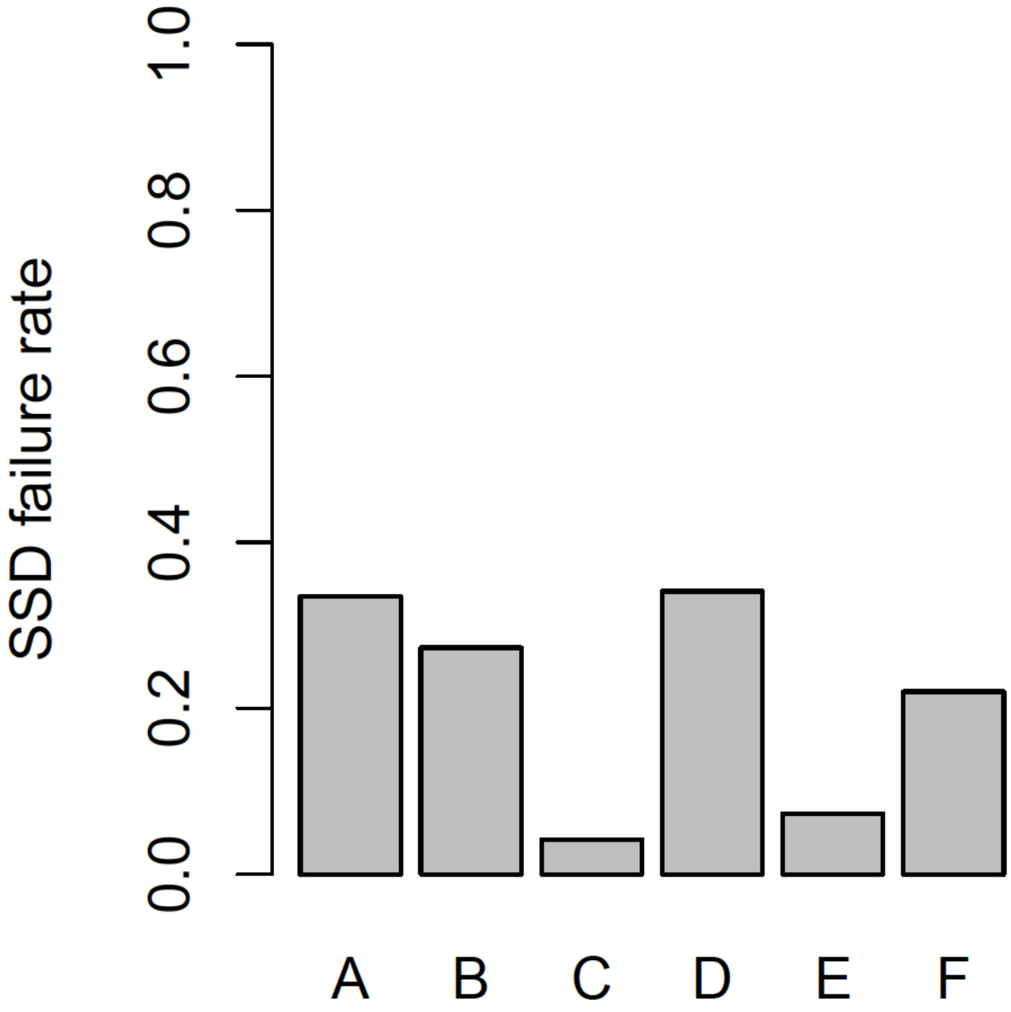

Older Platforms → Higher SSD Error Rates

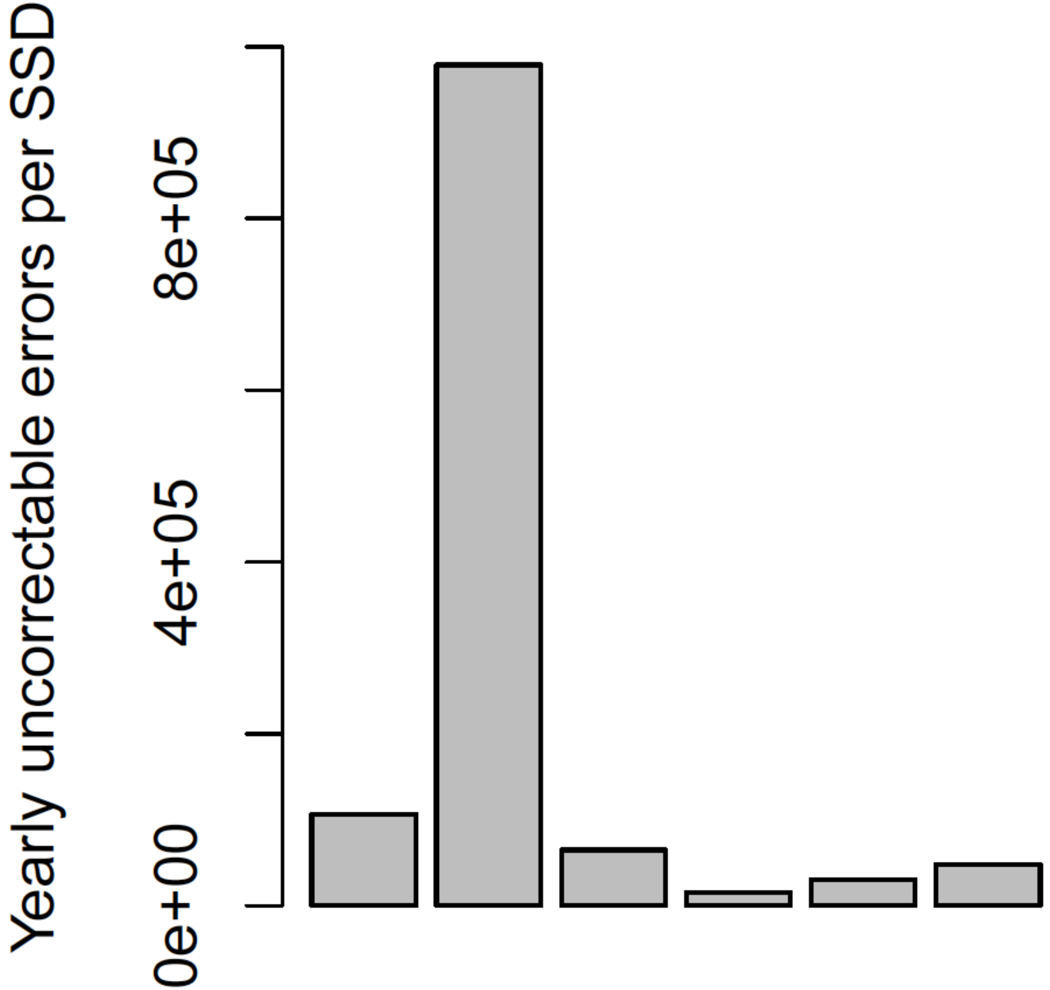

Platforms with Multiple SSDs

- Failures of SSDs in the same platform are correlated
  - Multiple SSDs in one host

- Conclusion: Operational conditions related to platform affect SSD failure trends

A few SSDs cause most errors

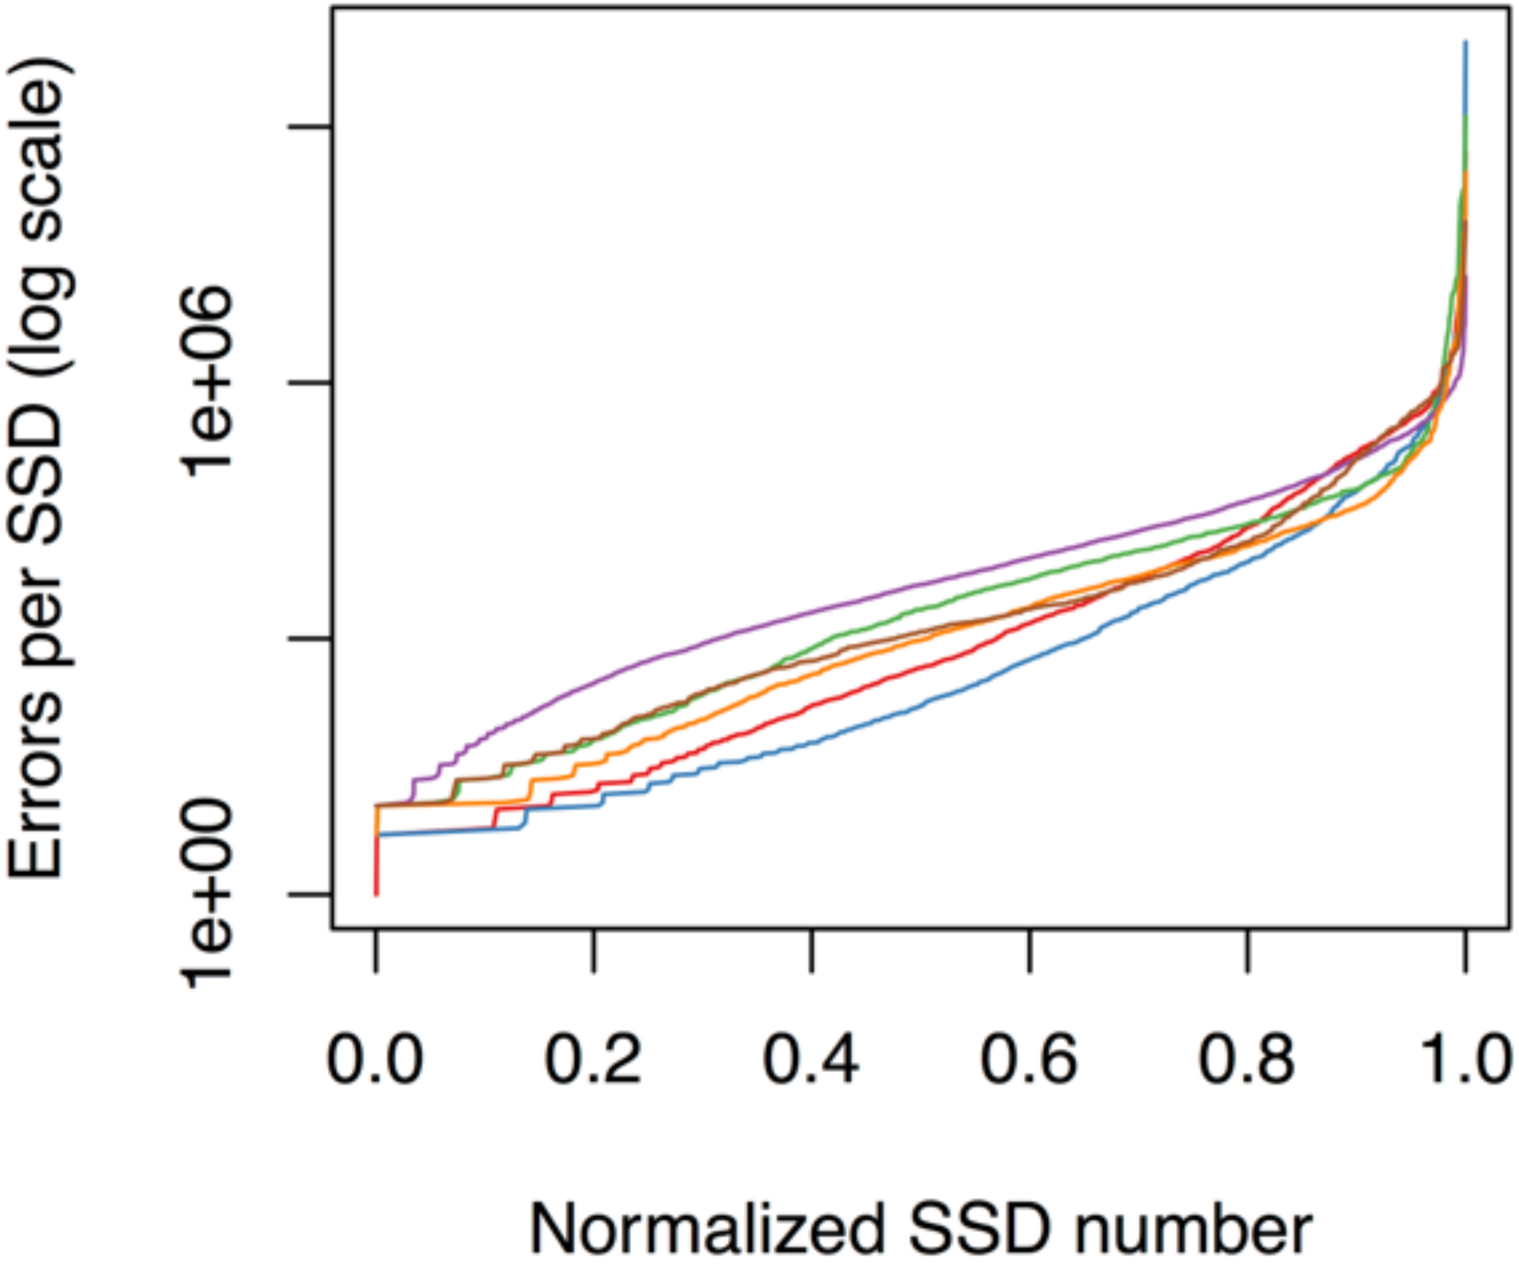

A few SSDs cause most errors

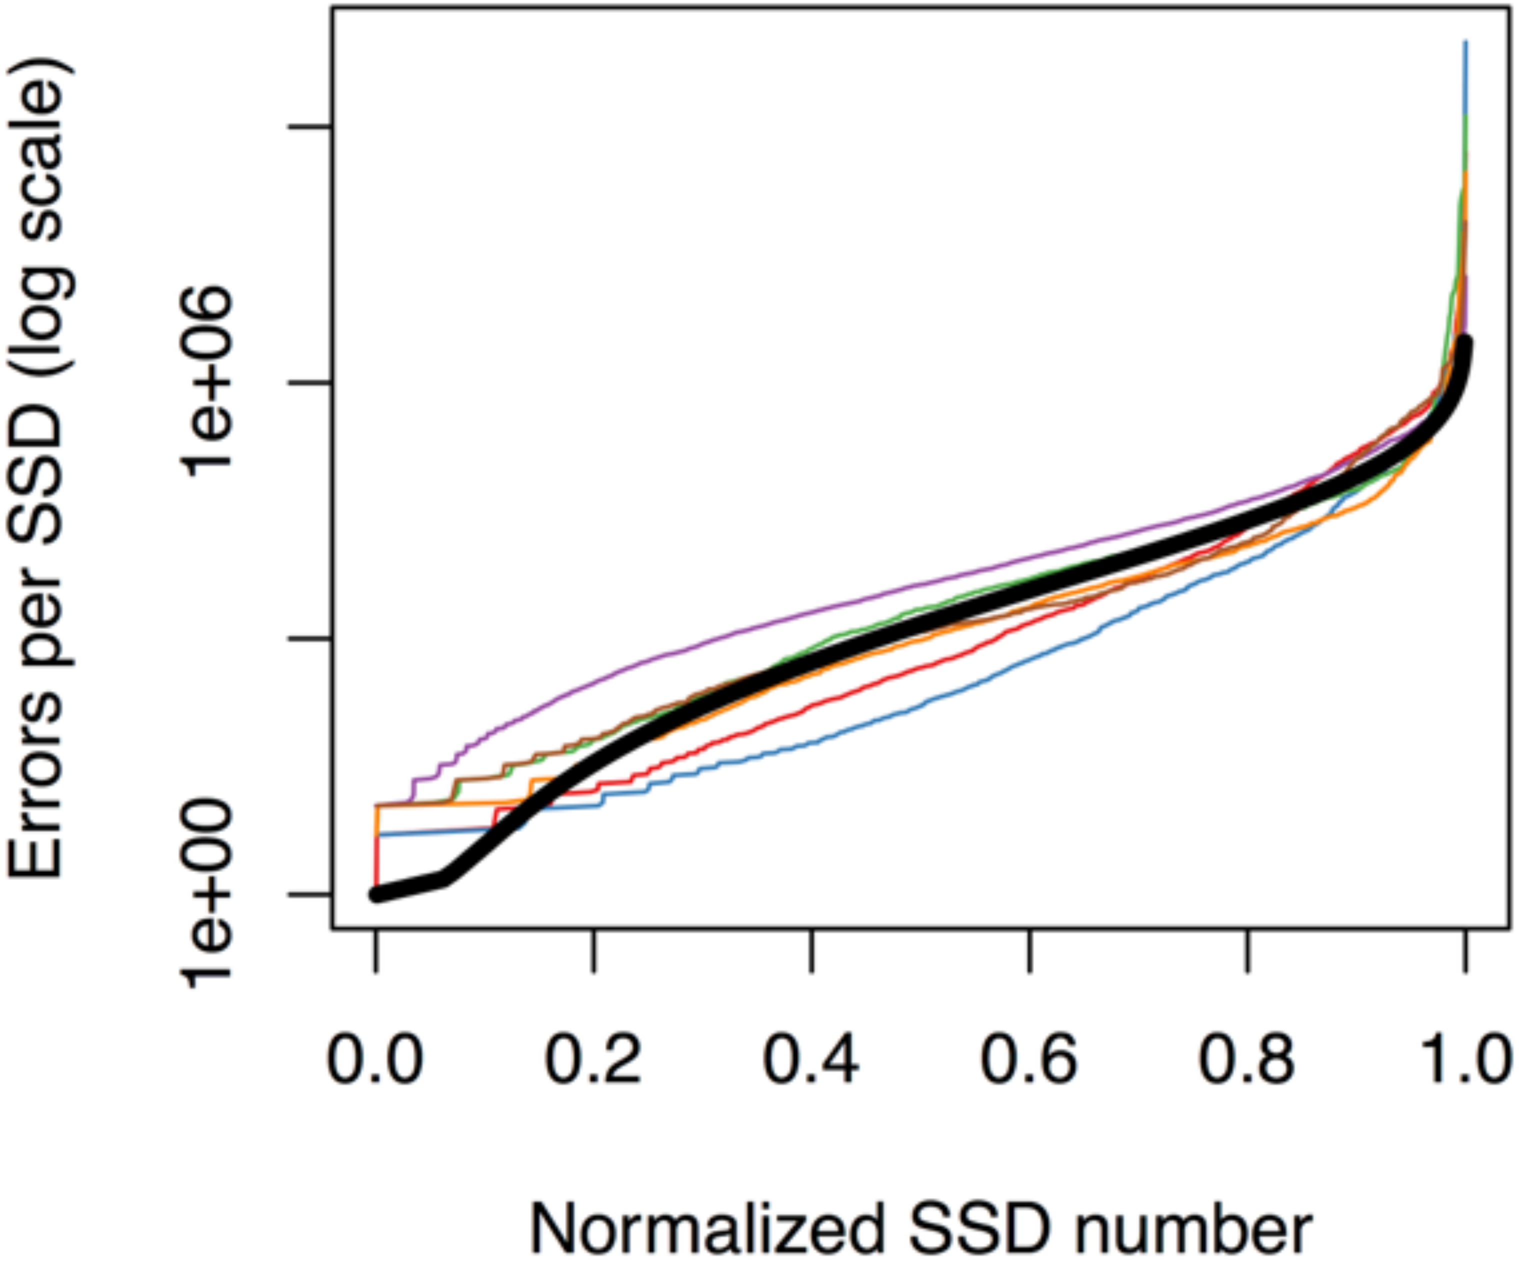

10% of SSDs have >80% of errors

Errors follow Weibull distribution

A few SSDs cause most errors

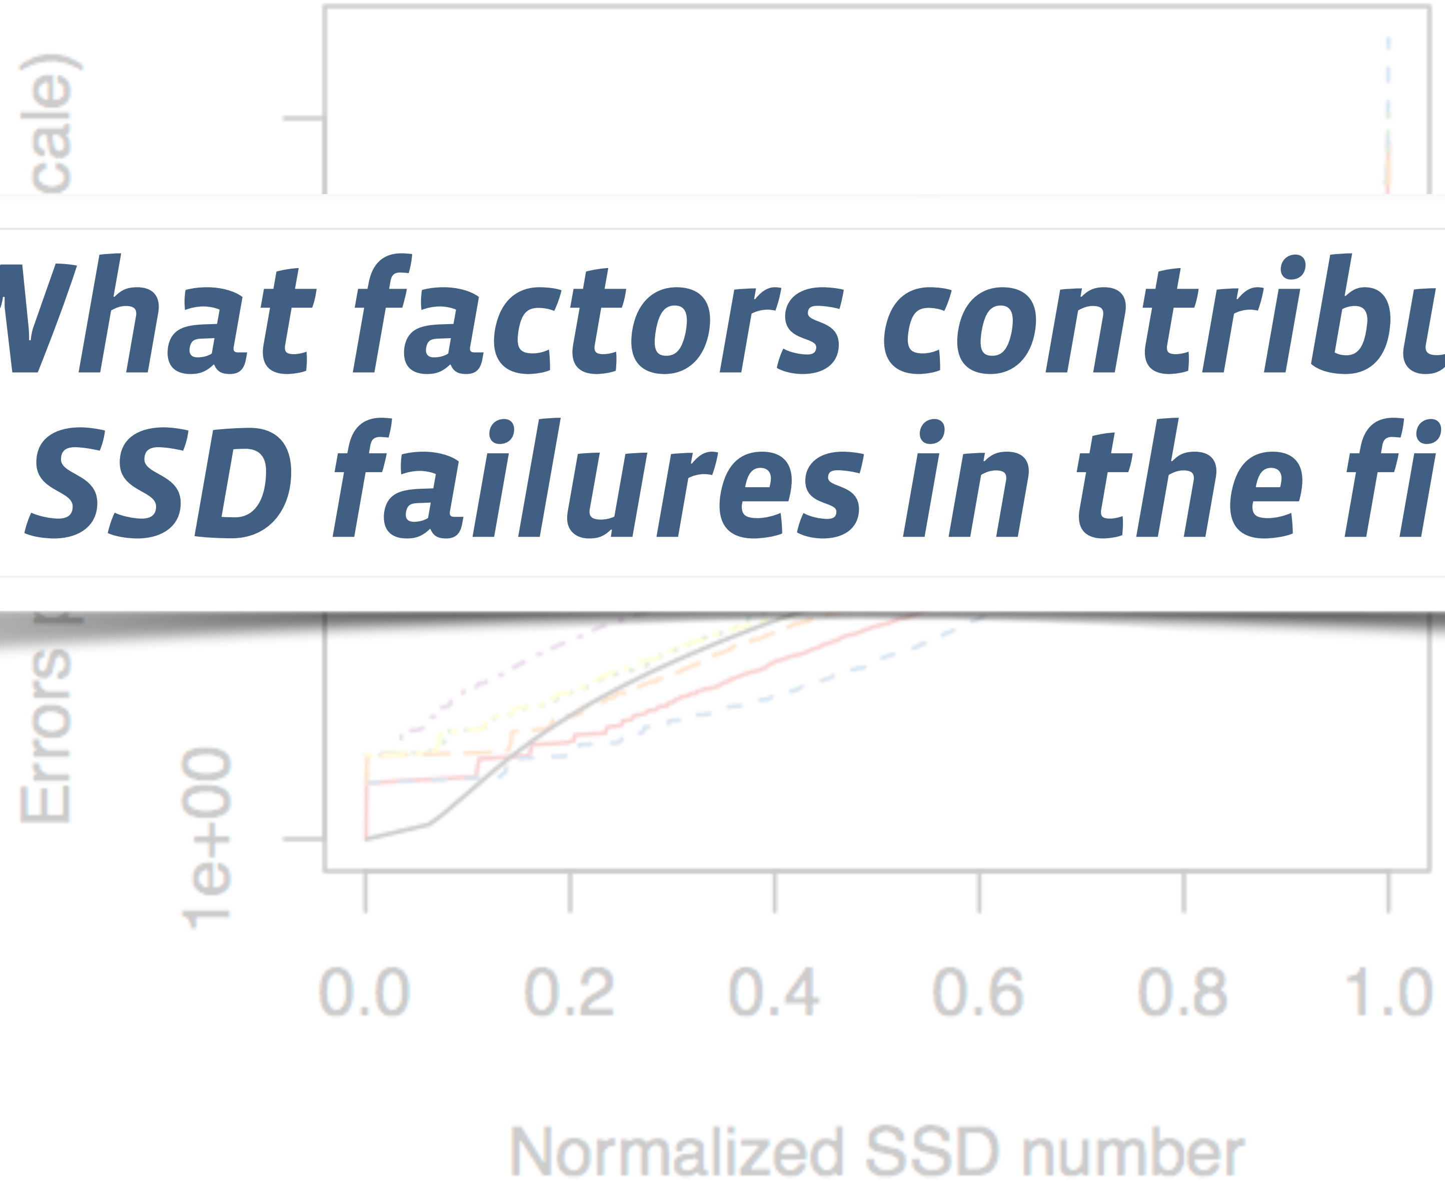

What factors contribute to SSD failures in the field?

Analytical methodology

- not feasible to log every error
- instead, analyze **lifetime counters**
- **snapshot-based** analysis

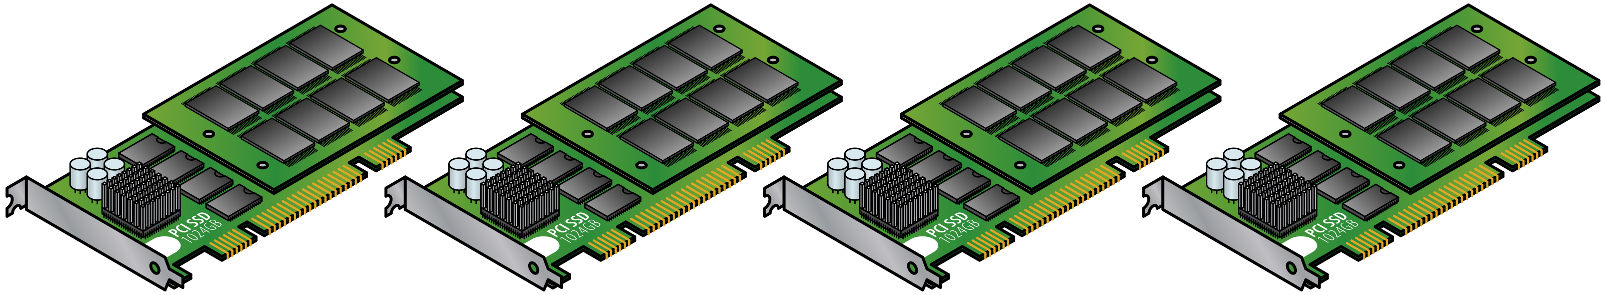
PCI SSD
1024GB
PCI SSD
1024GB
PCI SSD
1024GB
PCI SSD
1024GB

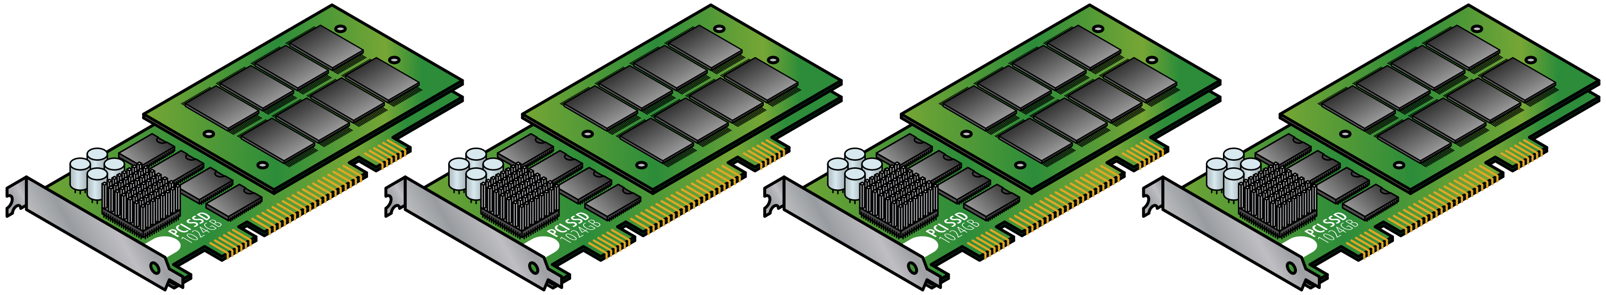

Errors	54,326	0	2	10
Data written	10TB	2TB	5TB	6TB

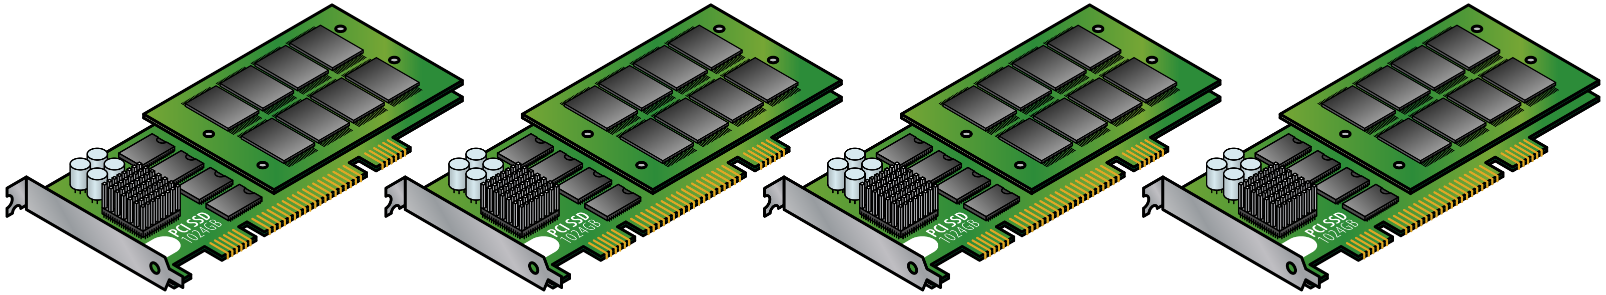

Errors	54,326	0	2	10
Data written	10TB	2TB	5TB	6TB

2014-11-1

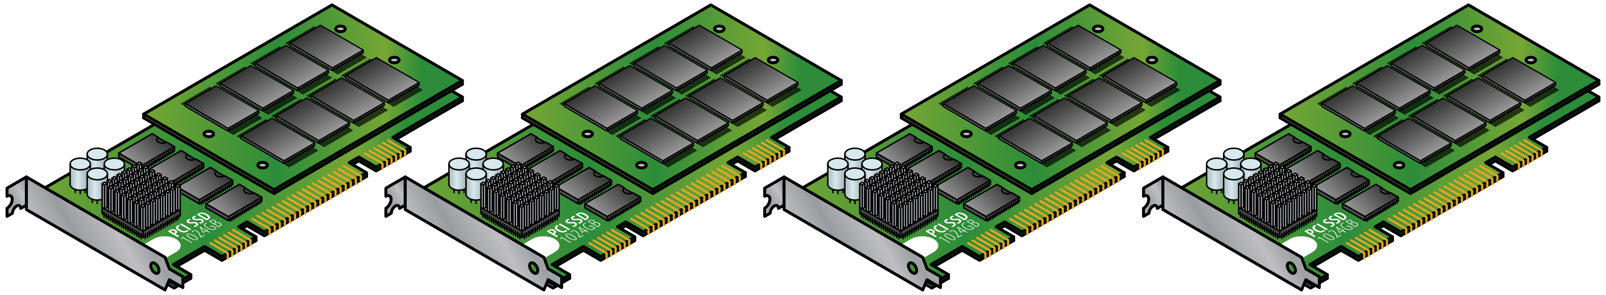

Errors	54,326	0	2	10
Data written	10TB	2TB	5TB	6TB

2014-11-1

Errors

Data written

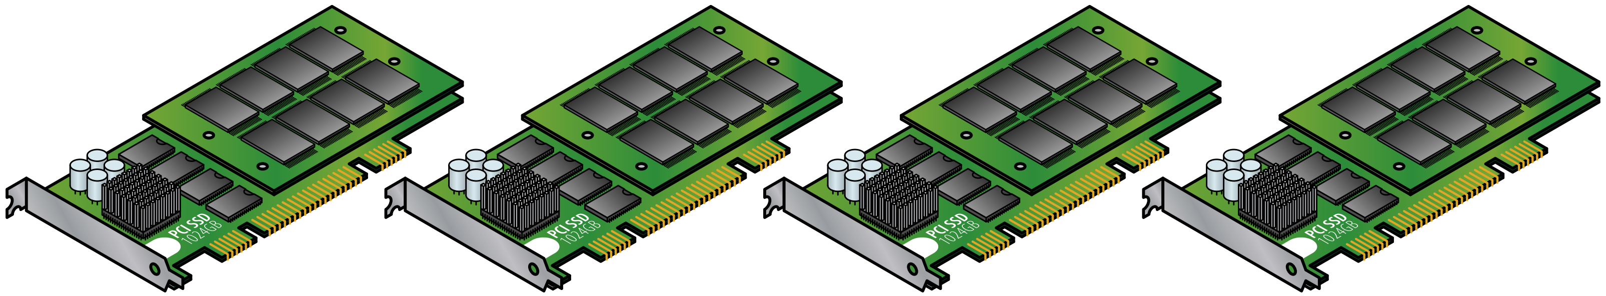

Errors	54,326	0	2	10
Data written	10TB	2TB	5TB	6TB

2014-11-1

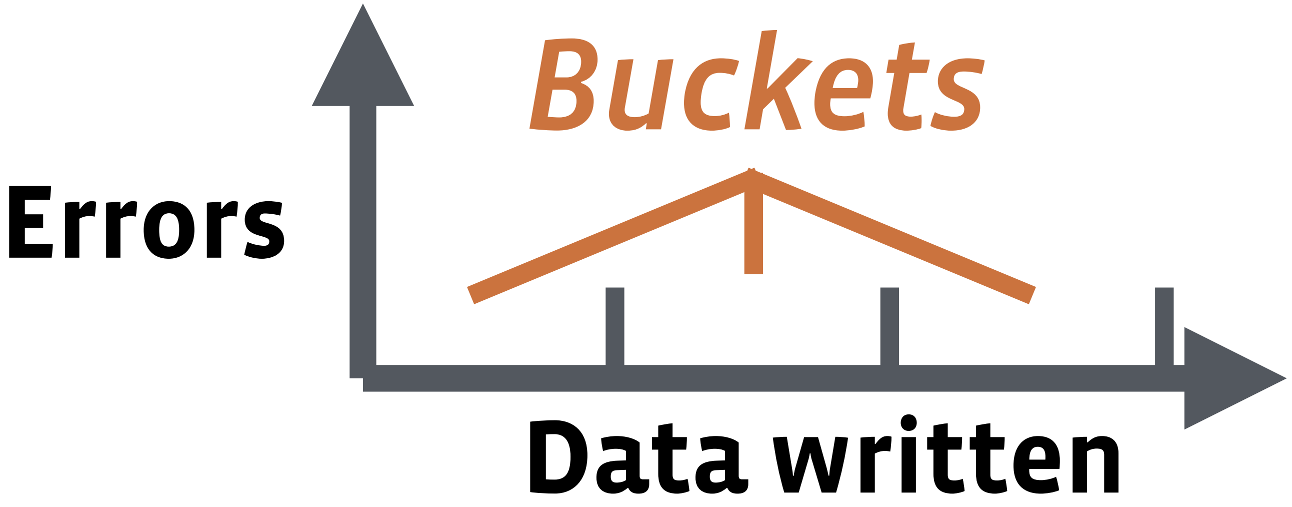

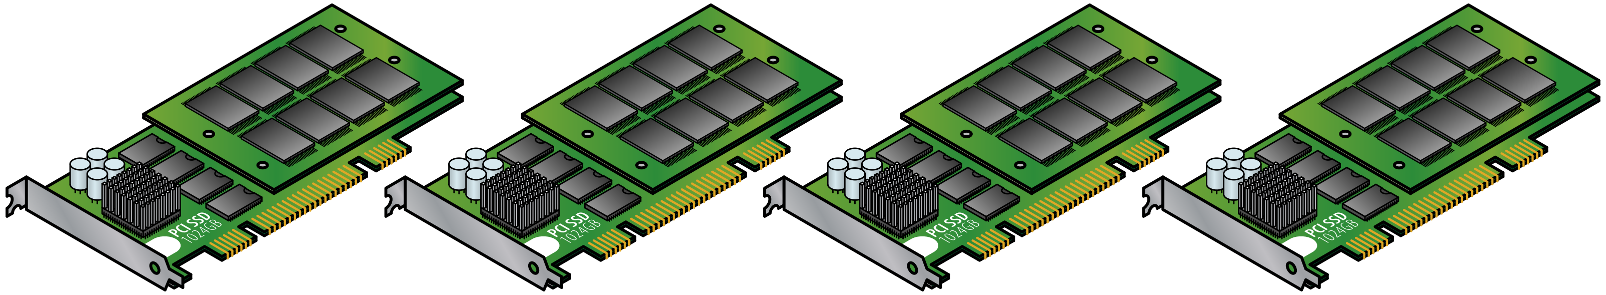

Errors	54,326	0	2	10
Data written	10TB	2TB	5TB	6TB

2014-11-1

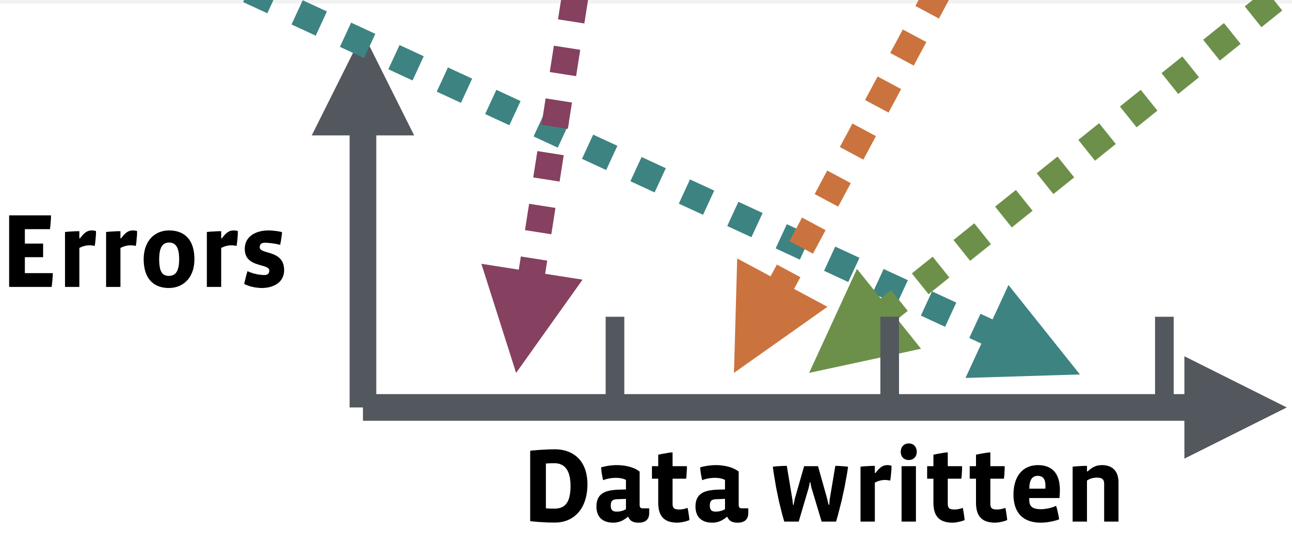

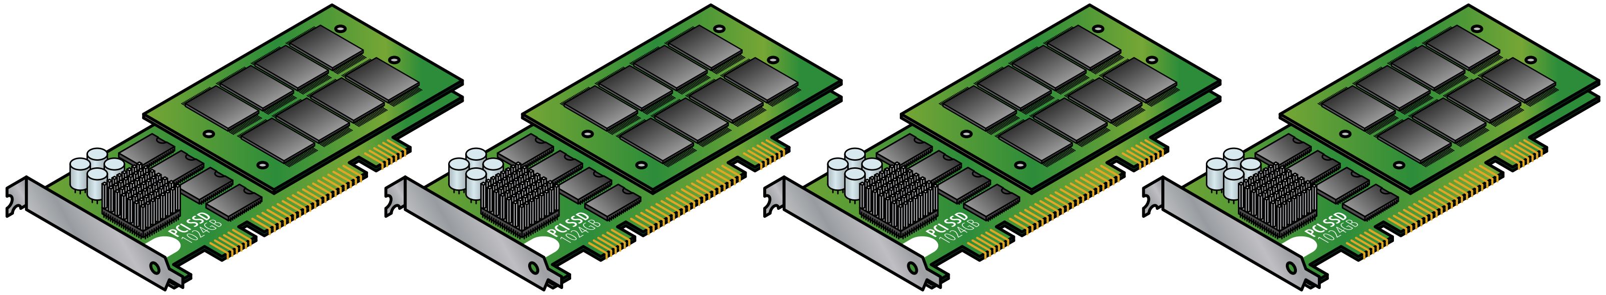

Errors	54,326	0	2	10
Data written	10TB	2TB	5TB	6TB

2014-11-1

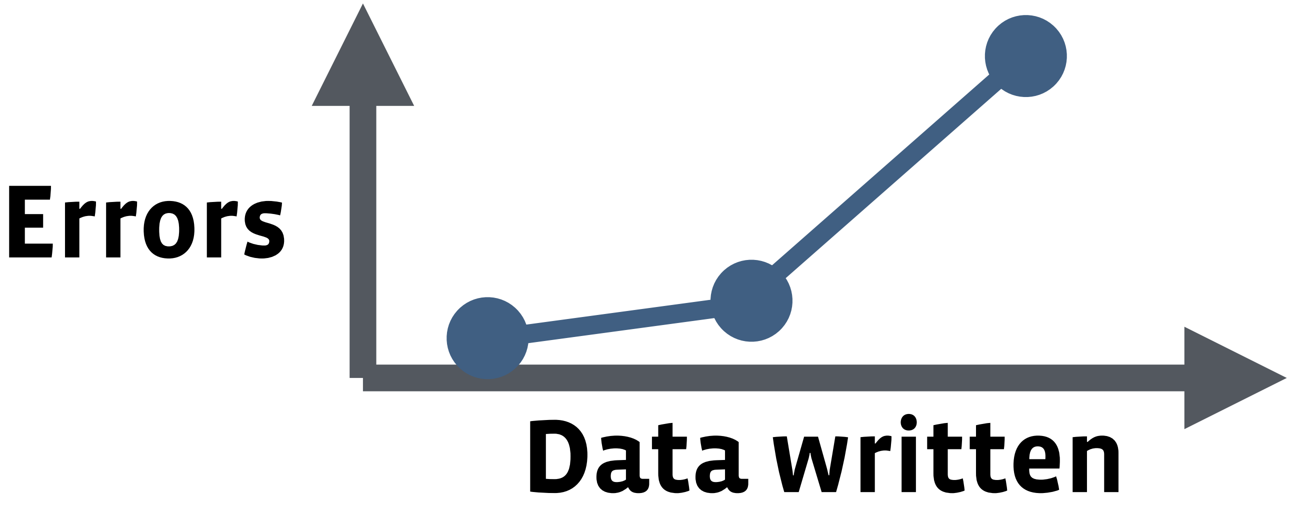

Some Definitions

- Uncorrectable error
  - Cannot be corrected by the SSD
  - But corrected by the host CPU driver

- SSD failure rate
  - Fraction of SSDs in a “bucket” that have had at least one uncorrectable error

SSD reliability trends

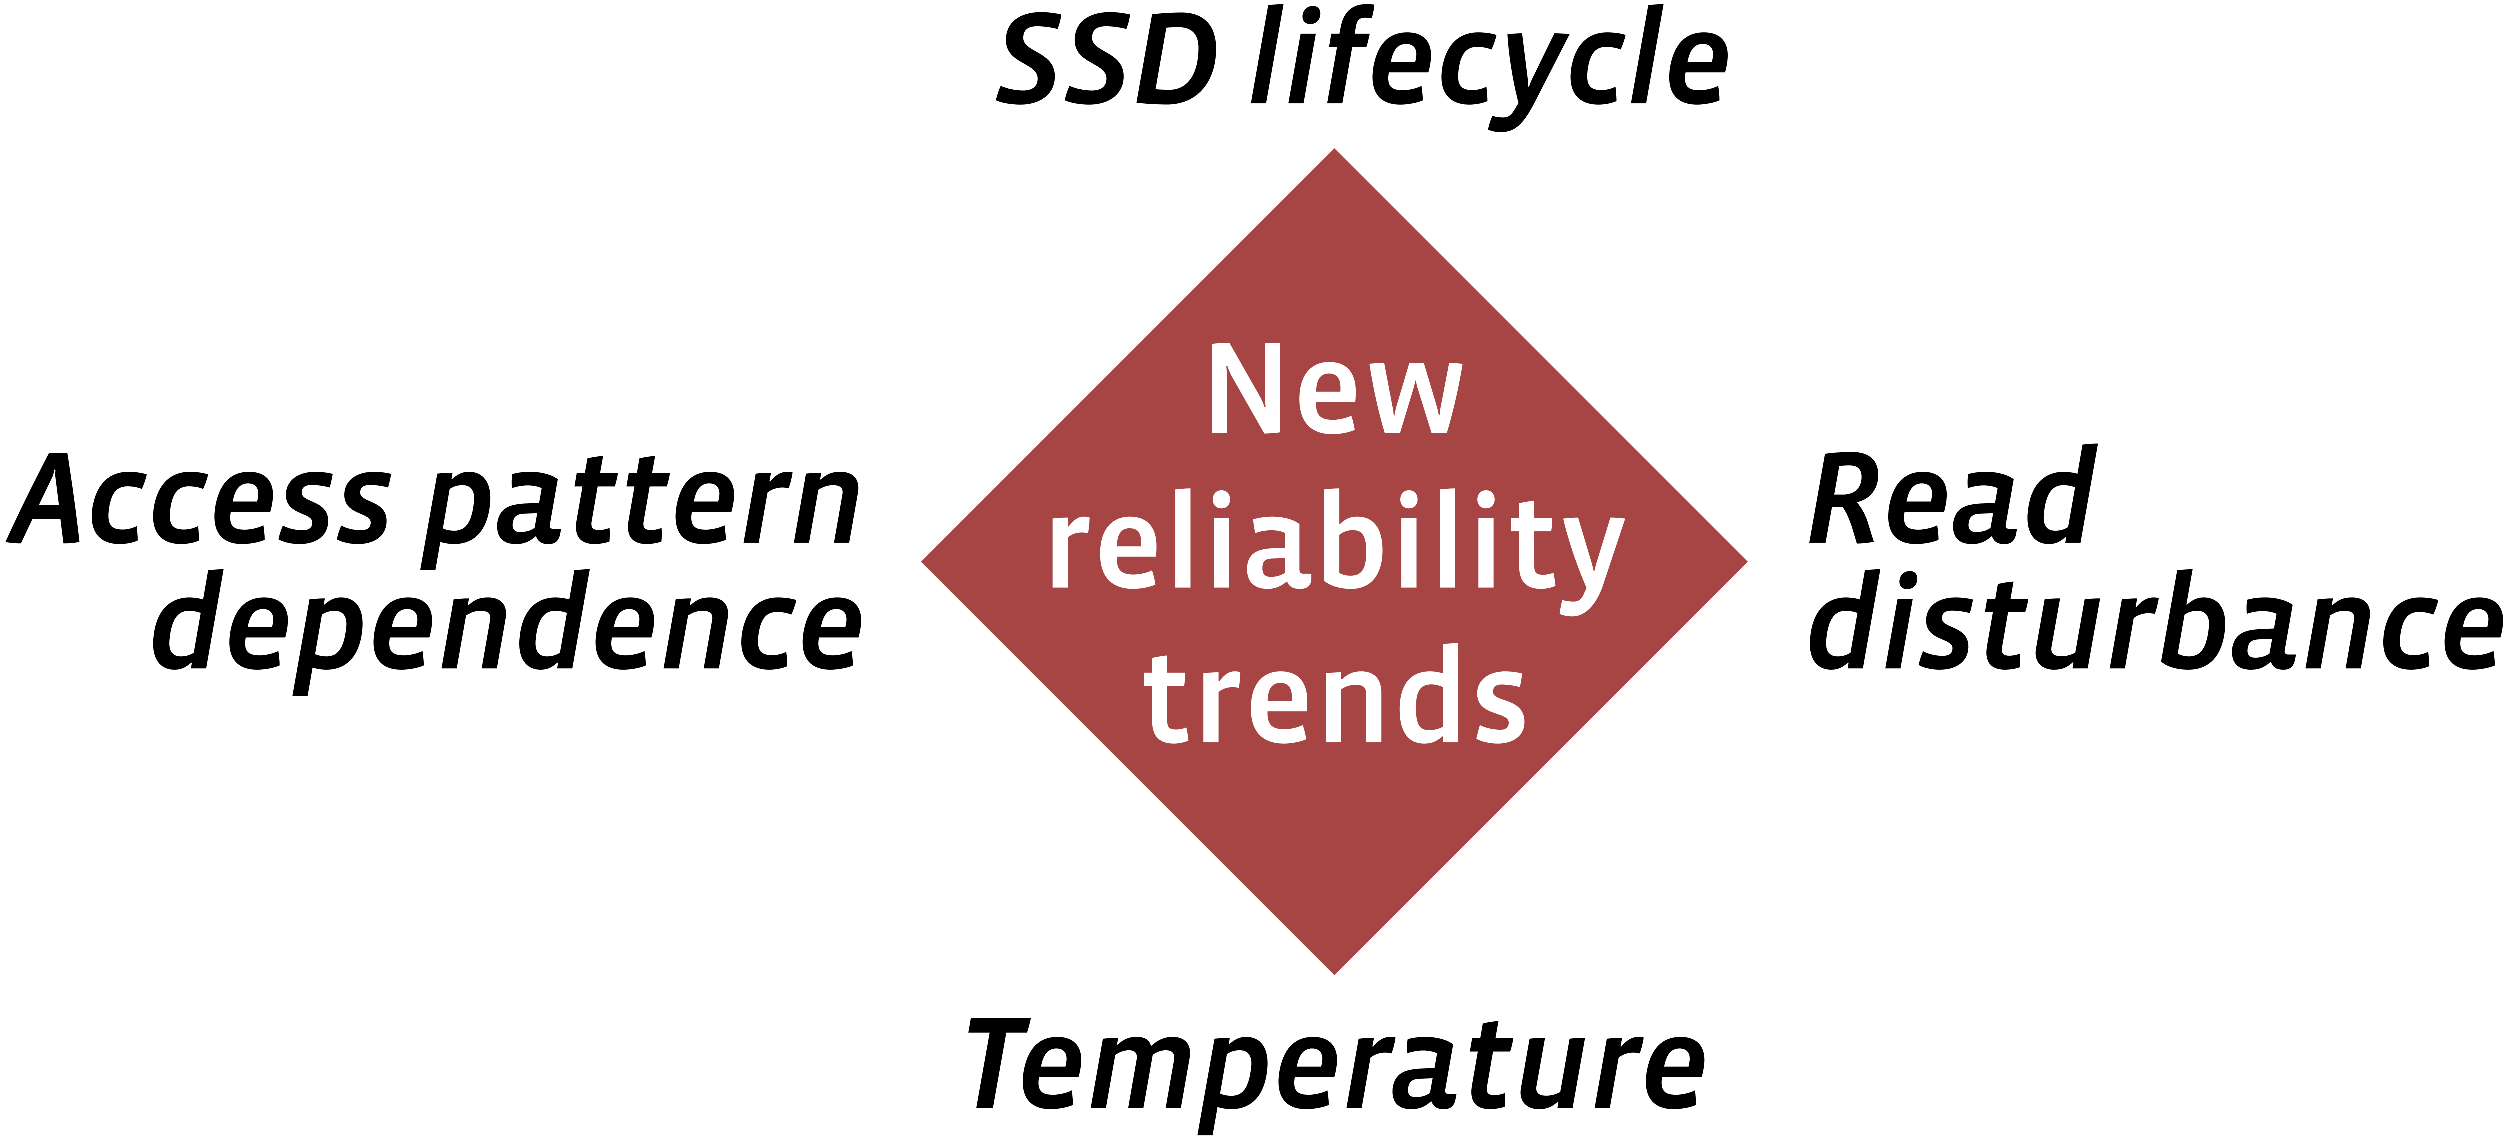

SSD lifecycle

Access pattern dependence

New reliability trends

Read disturbance

Temperature

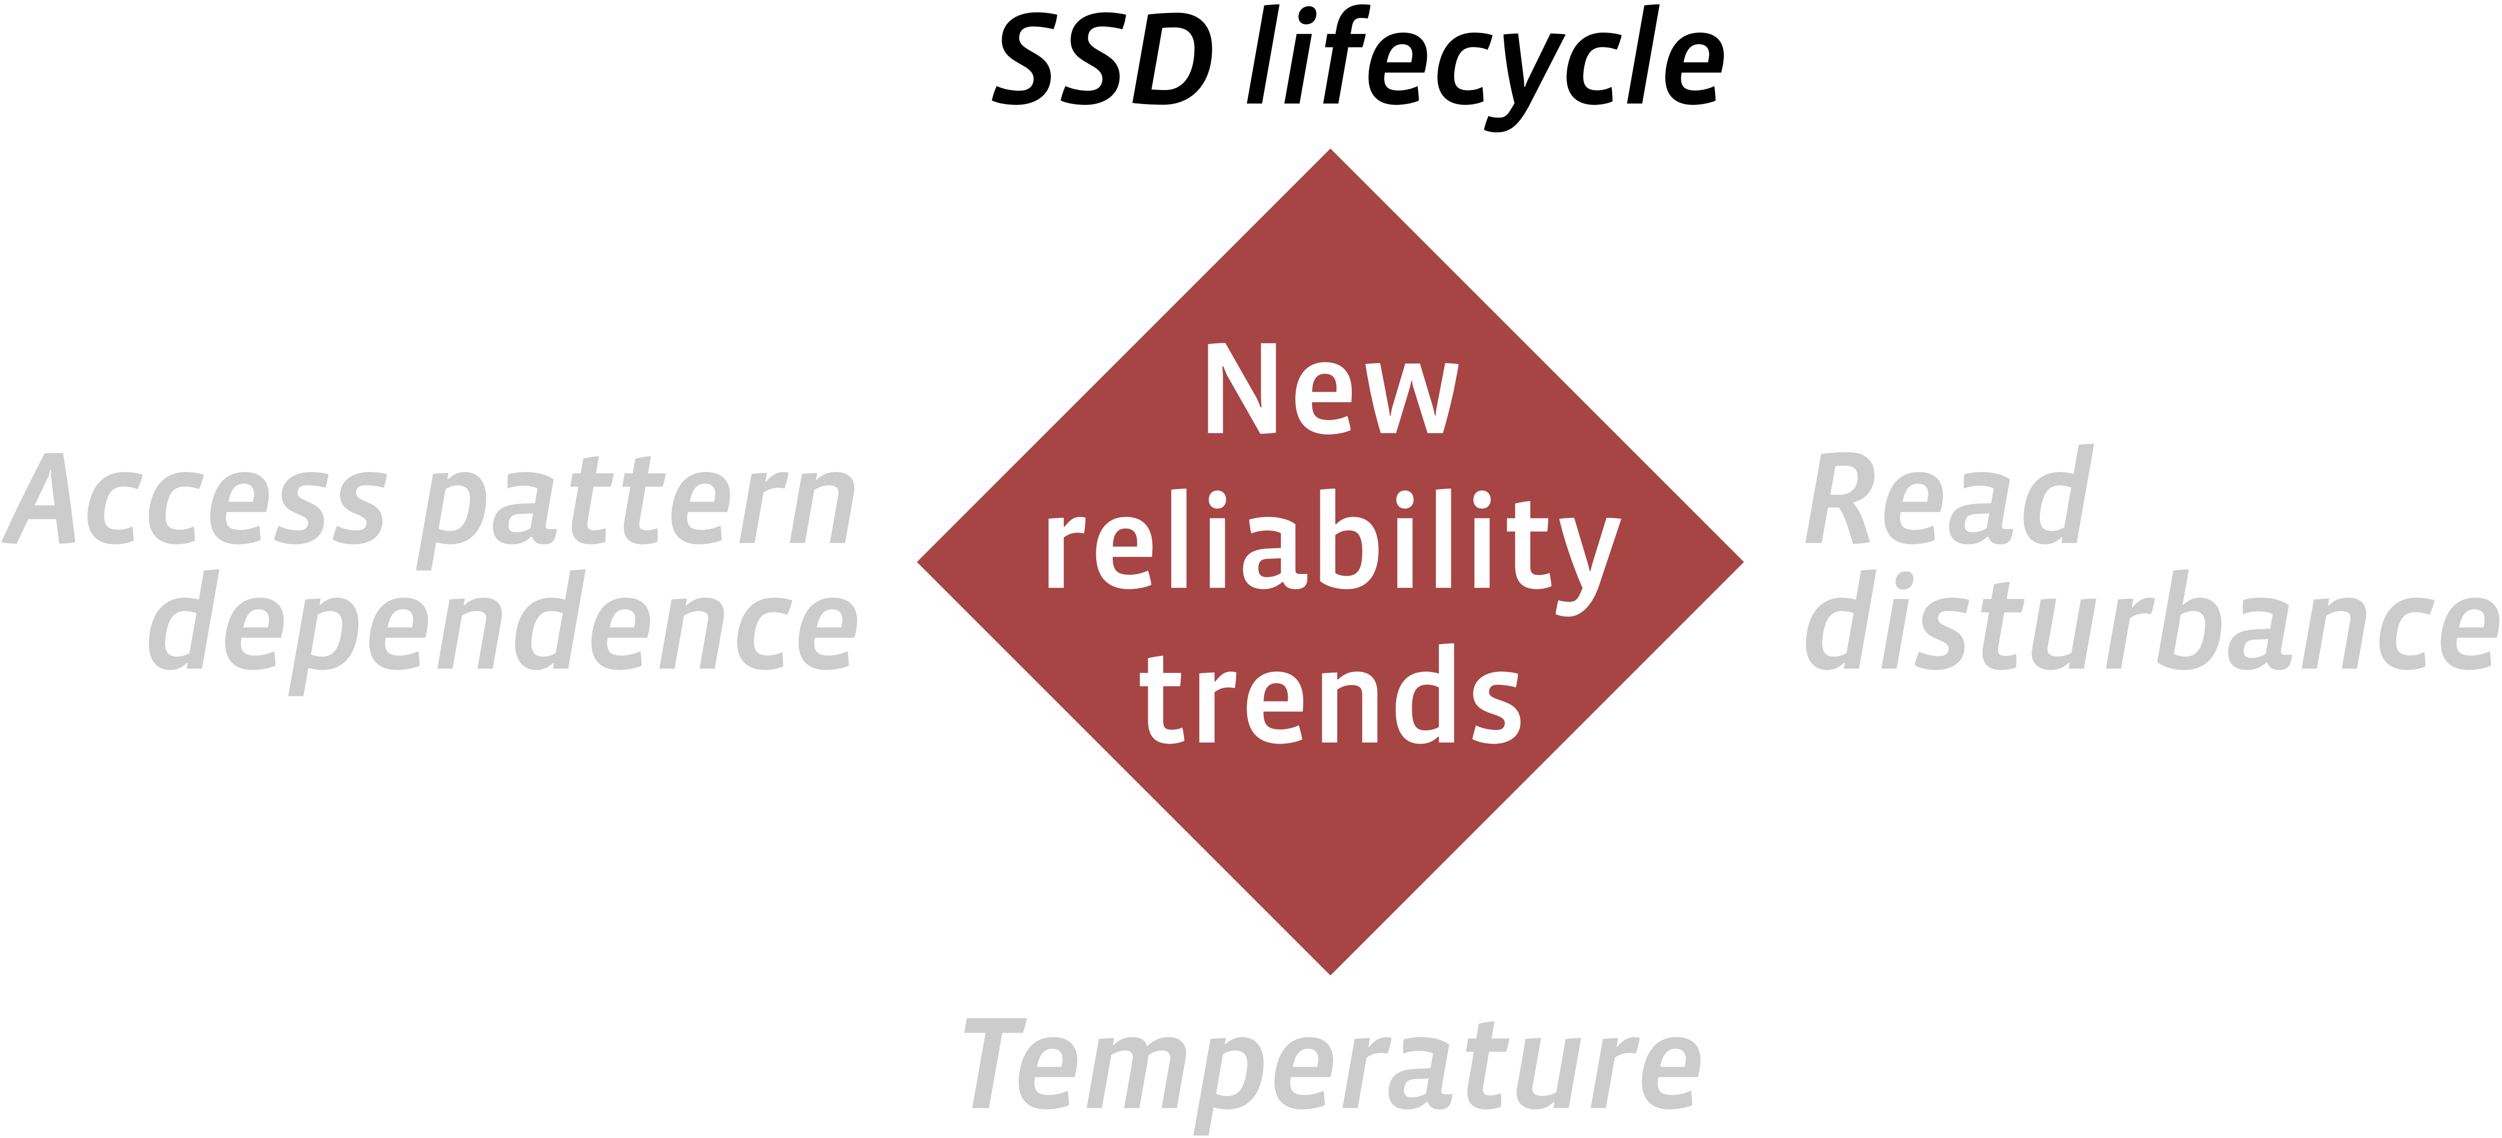

SSD lifecycle

Access pattern dependence

New reliability trends

Read disturbance

Temperature

Storage lifecycle background: *the **bathtub curve*** for disk drives

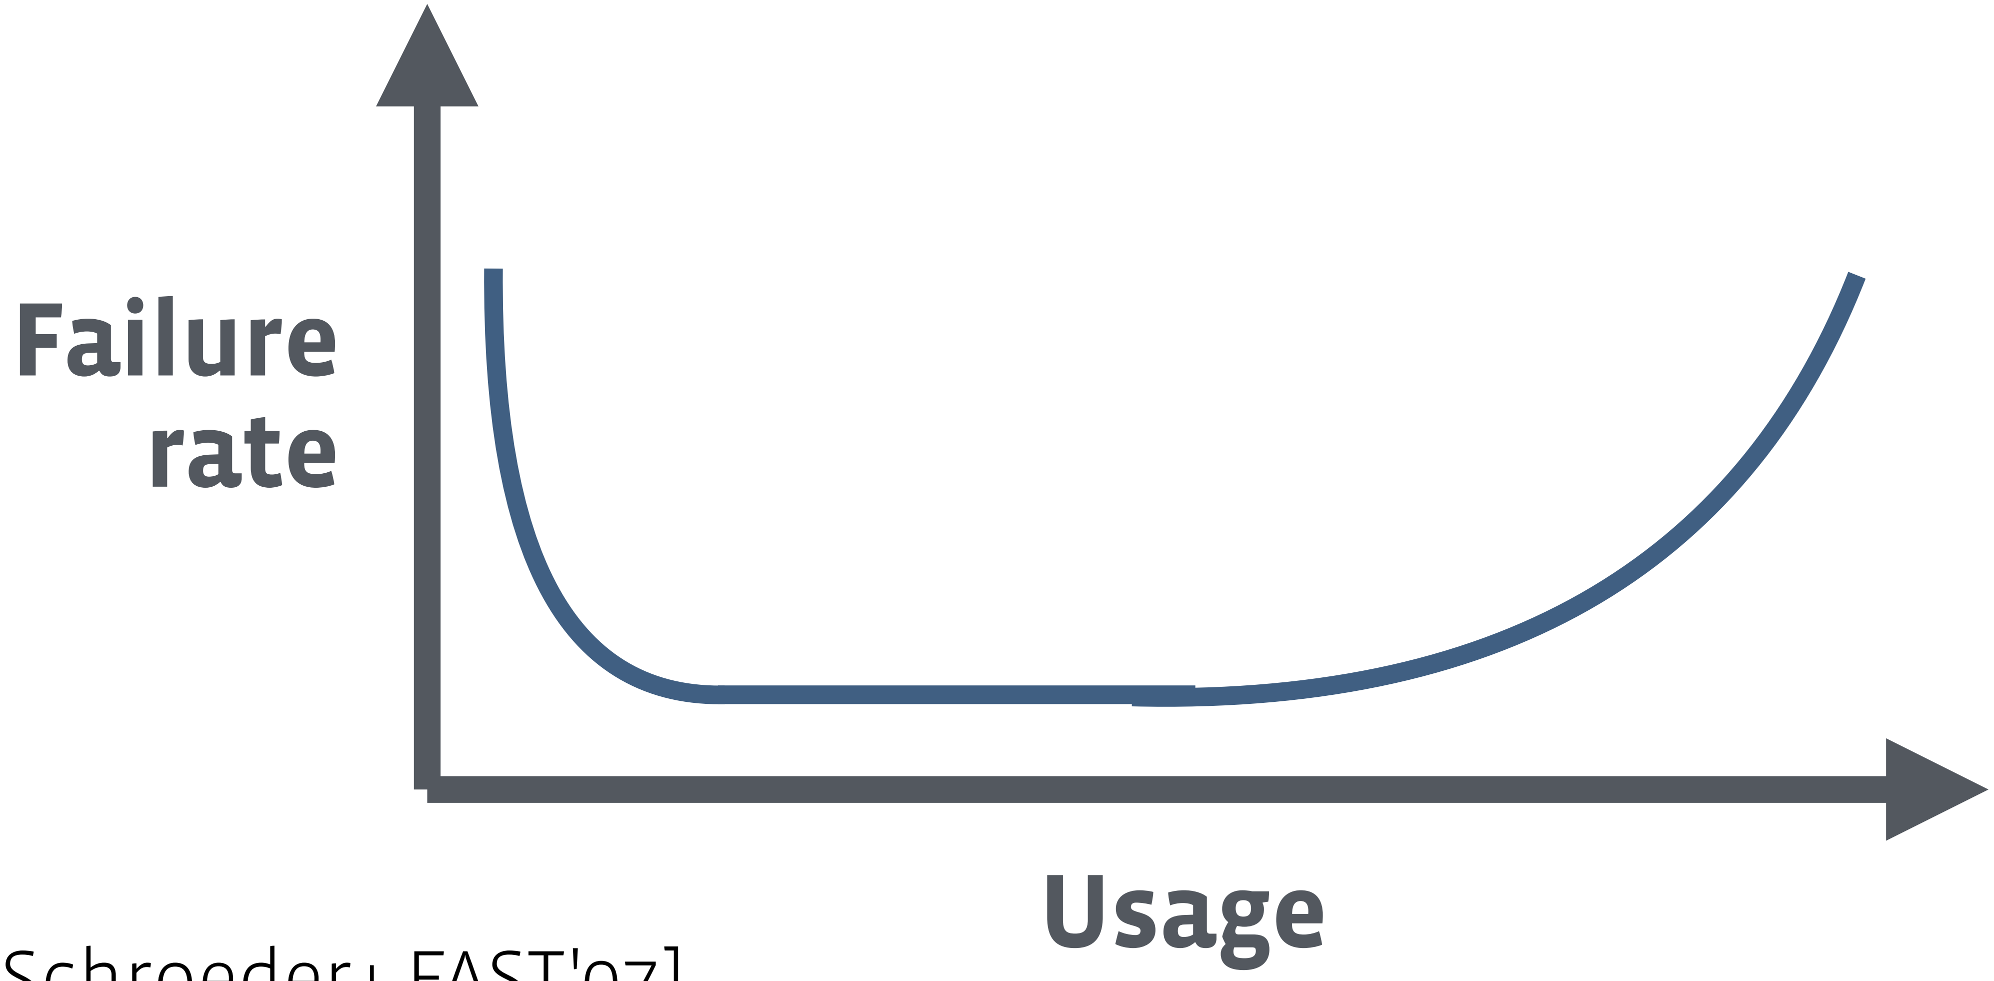

[Schroeder+,FAST'07]

Storage lifecycle background: *the* ***bathtub curve*** for disk drives

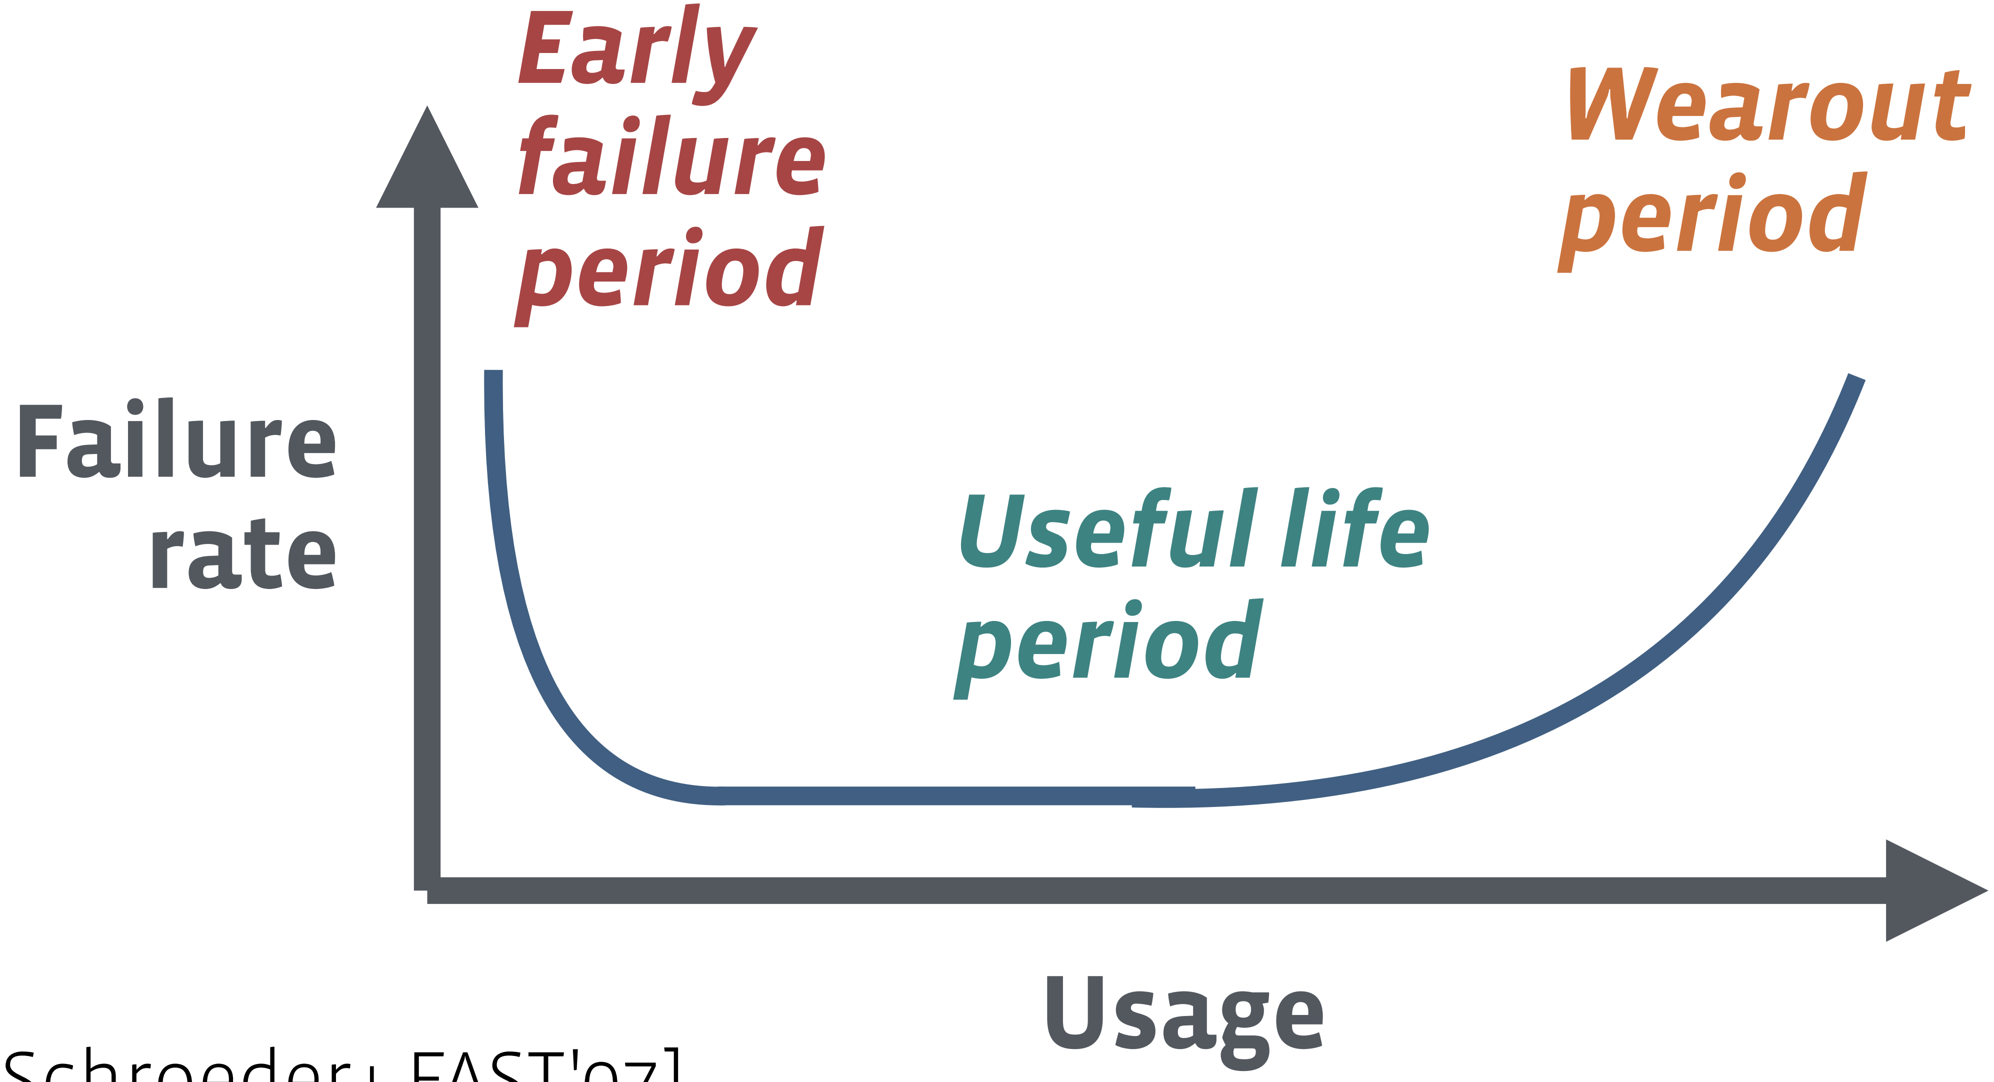

[Schroeder+,FAST'07]

Storage lifecycle background:
*the **bathtub curve** for disk drives*

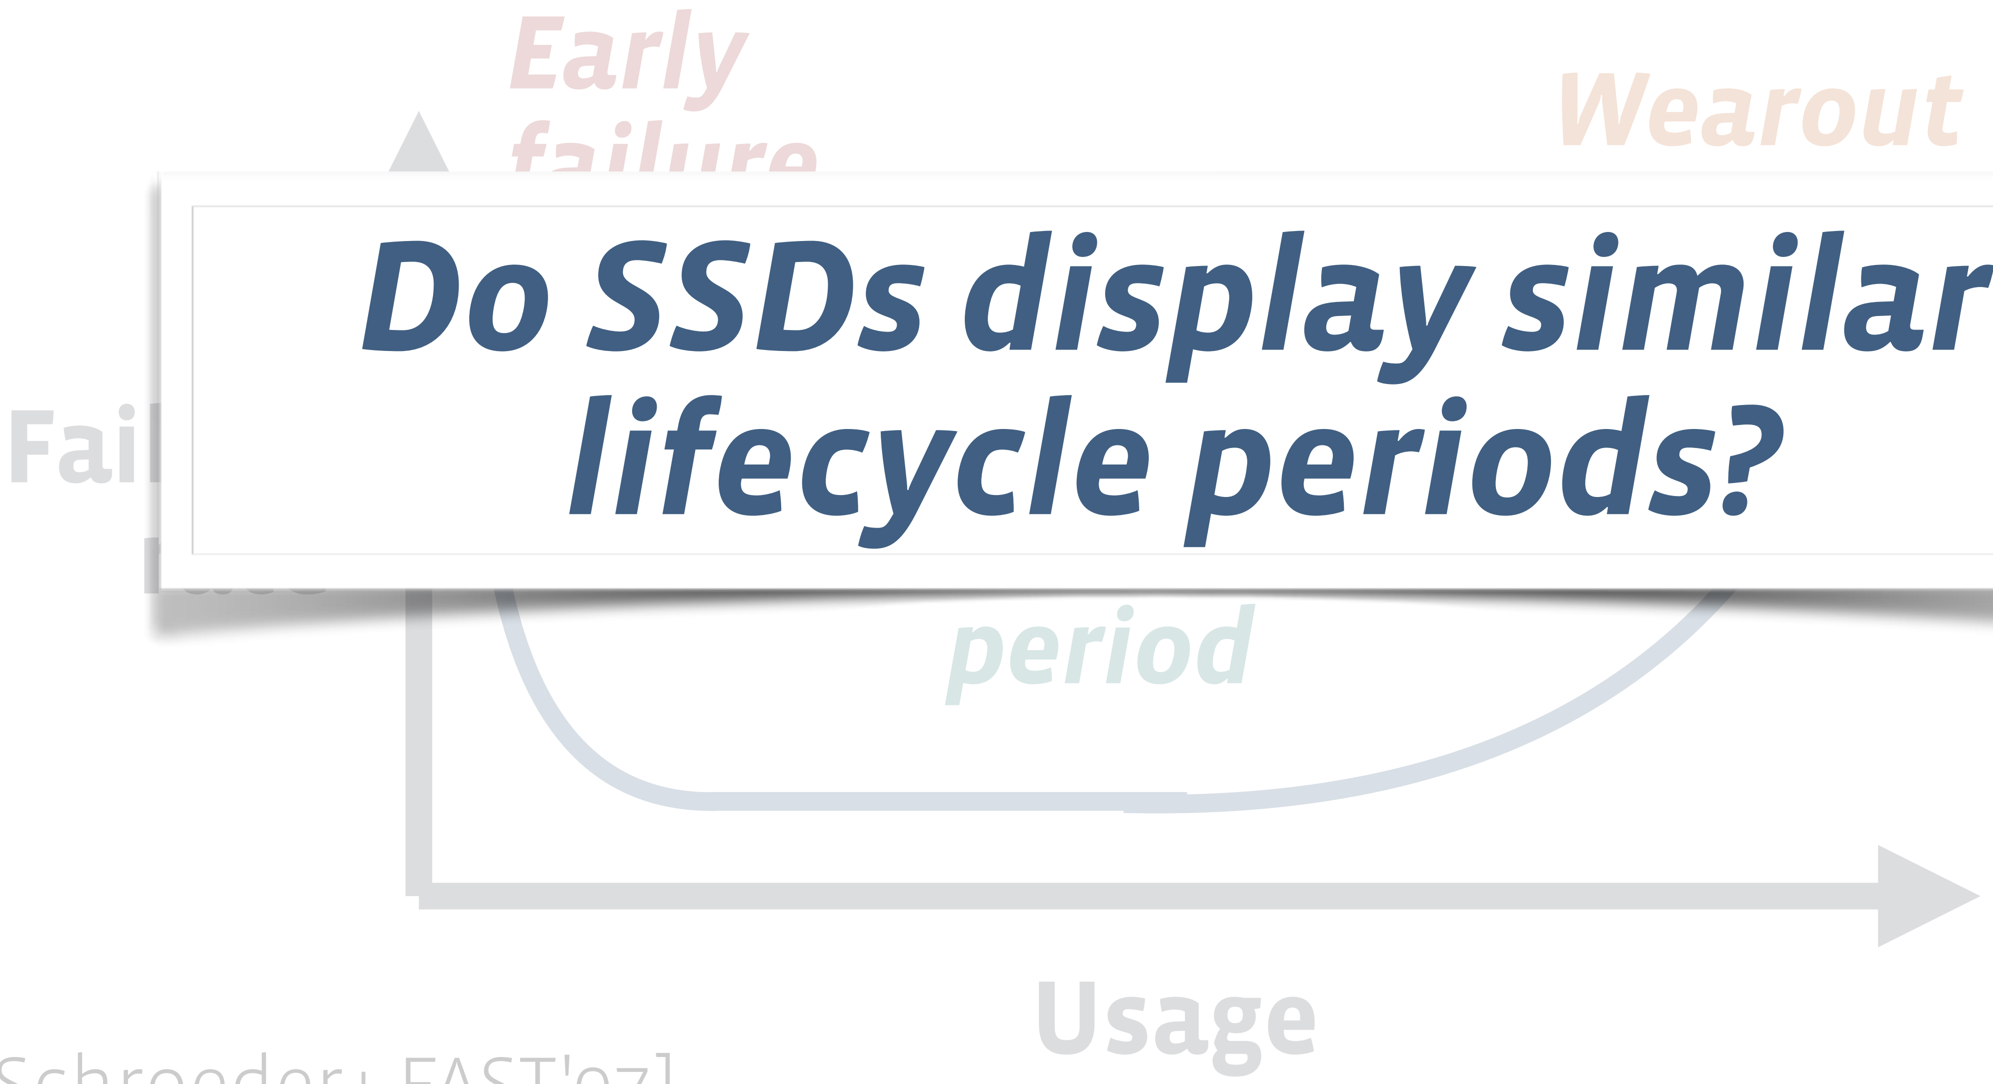

Do SSDs display similar lifecycle periods?

[Schroeder+,FAST'07]

Use **data written to flash**

to examine SSD lifecycle

(time-independent utilization metric)

What We Find

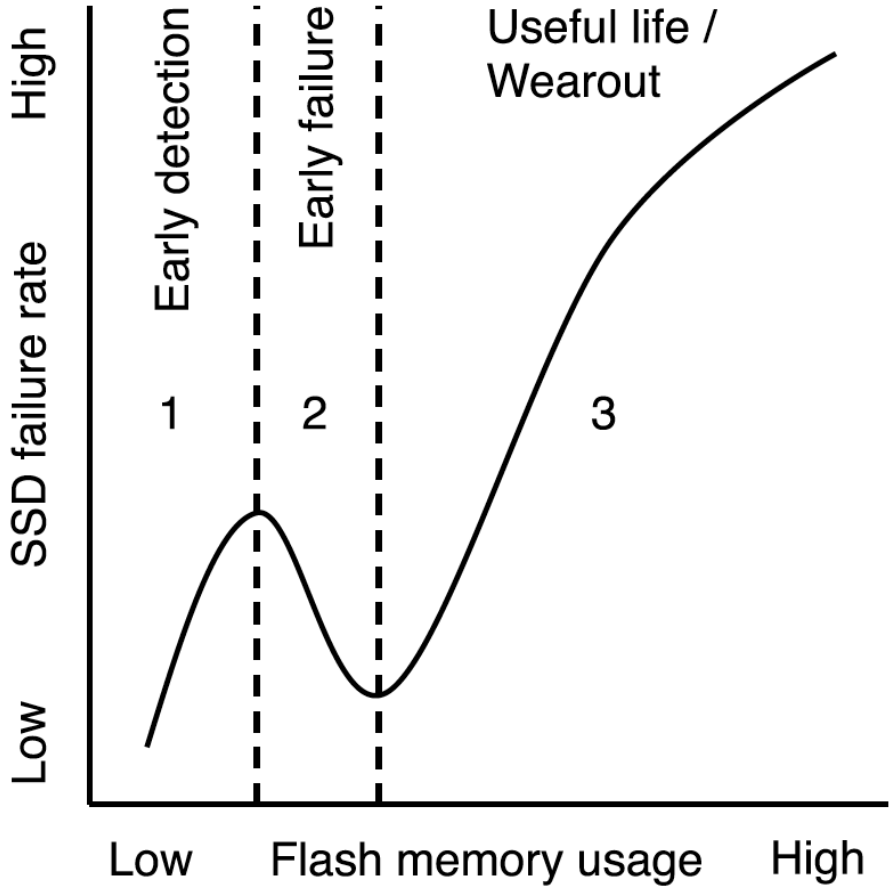

Figure 4: SSD lifecycle failure pattern. SSDs fail at different rates during several distinct periods throughout their lifetime (measured by usage)

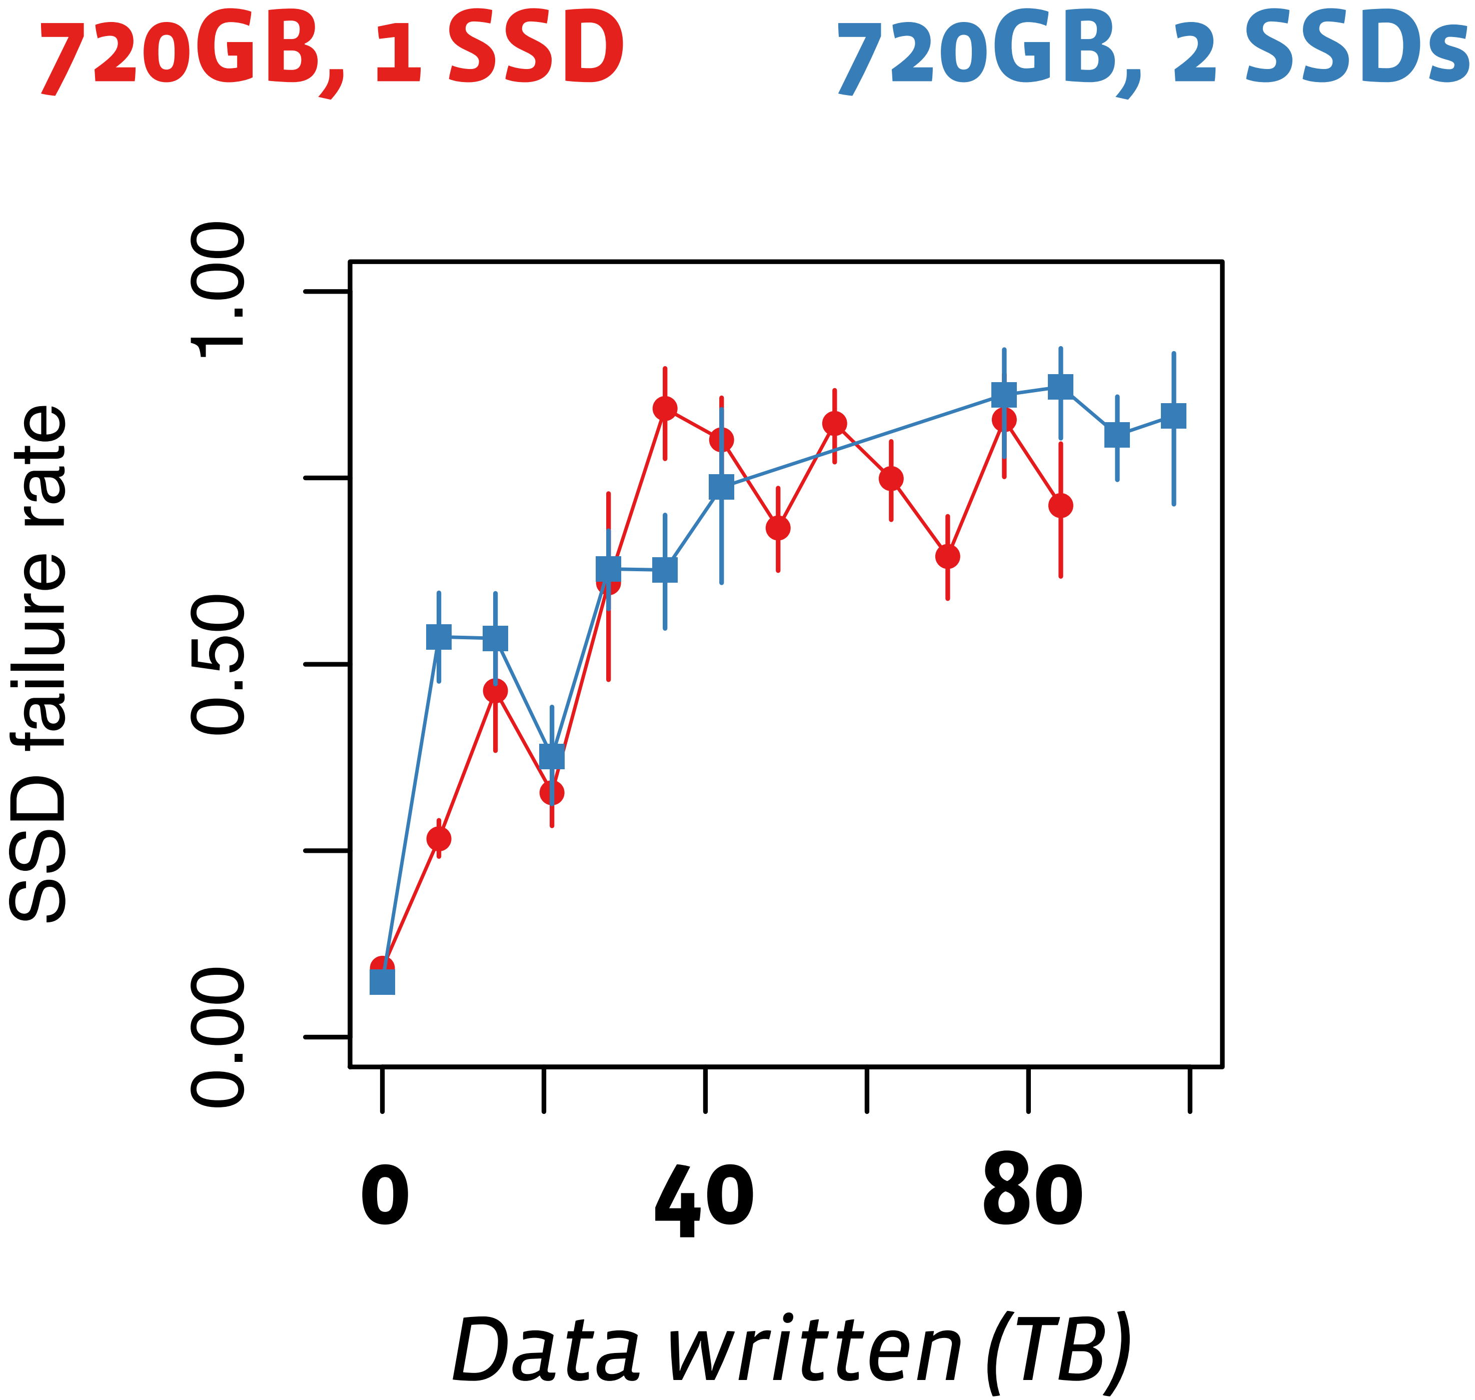
720GB, 1 SSD
720GB, 2 SSDs
SSD failure rate
1.00
0.50
0.00
0
40
80
Data written (TB)

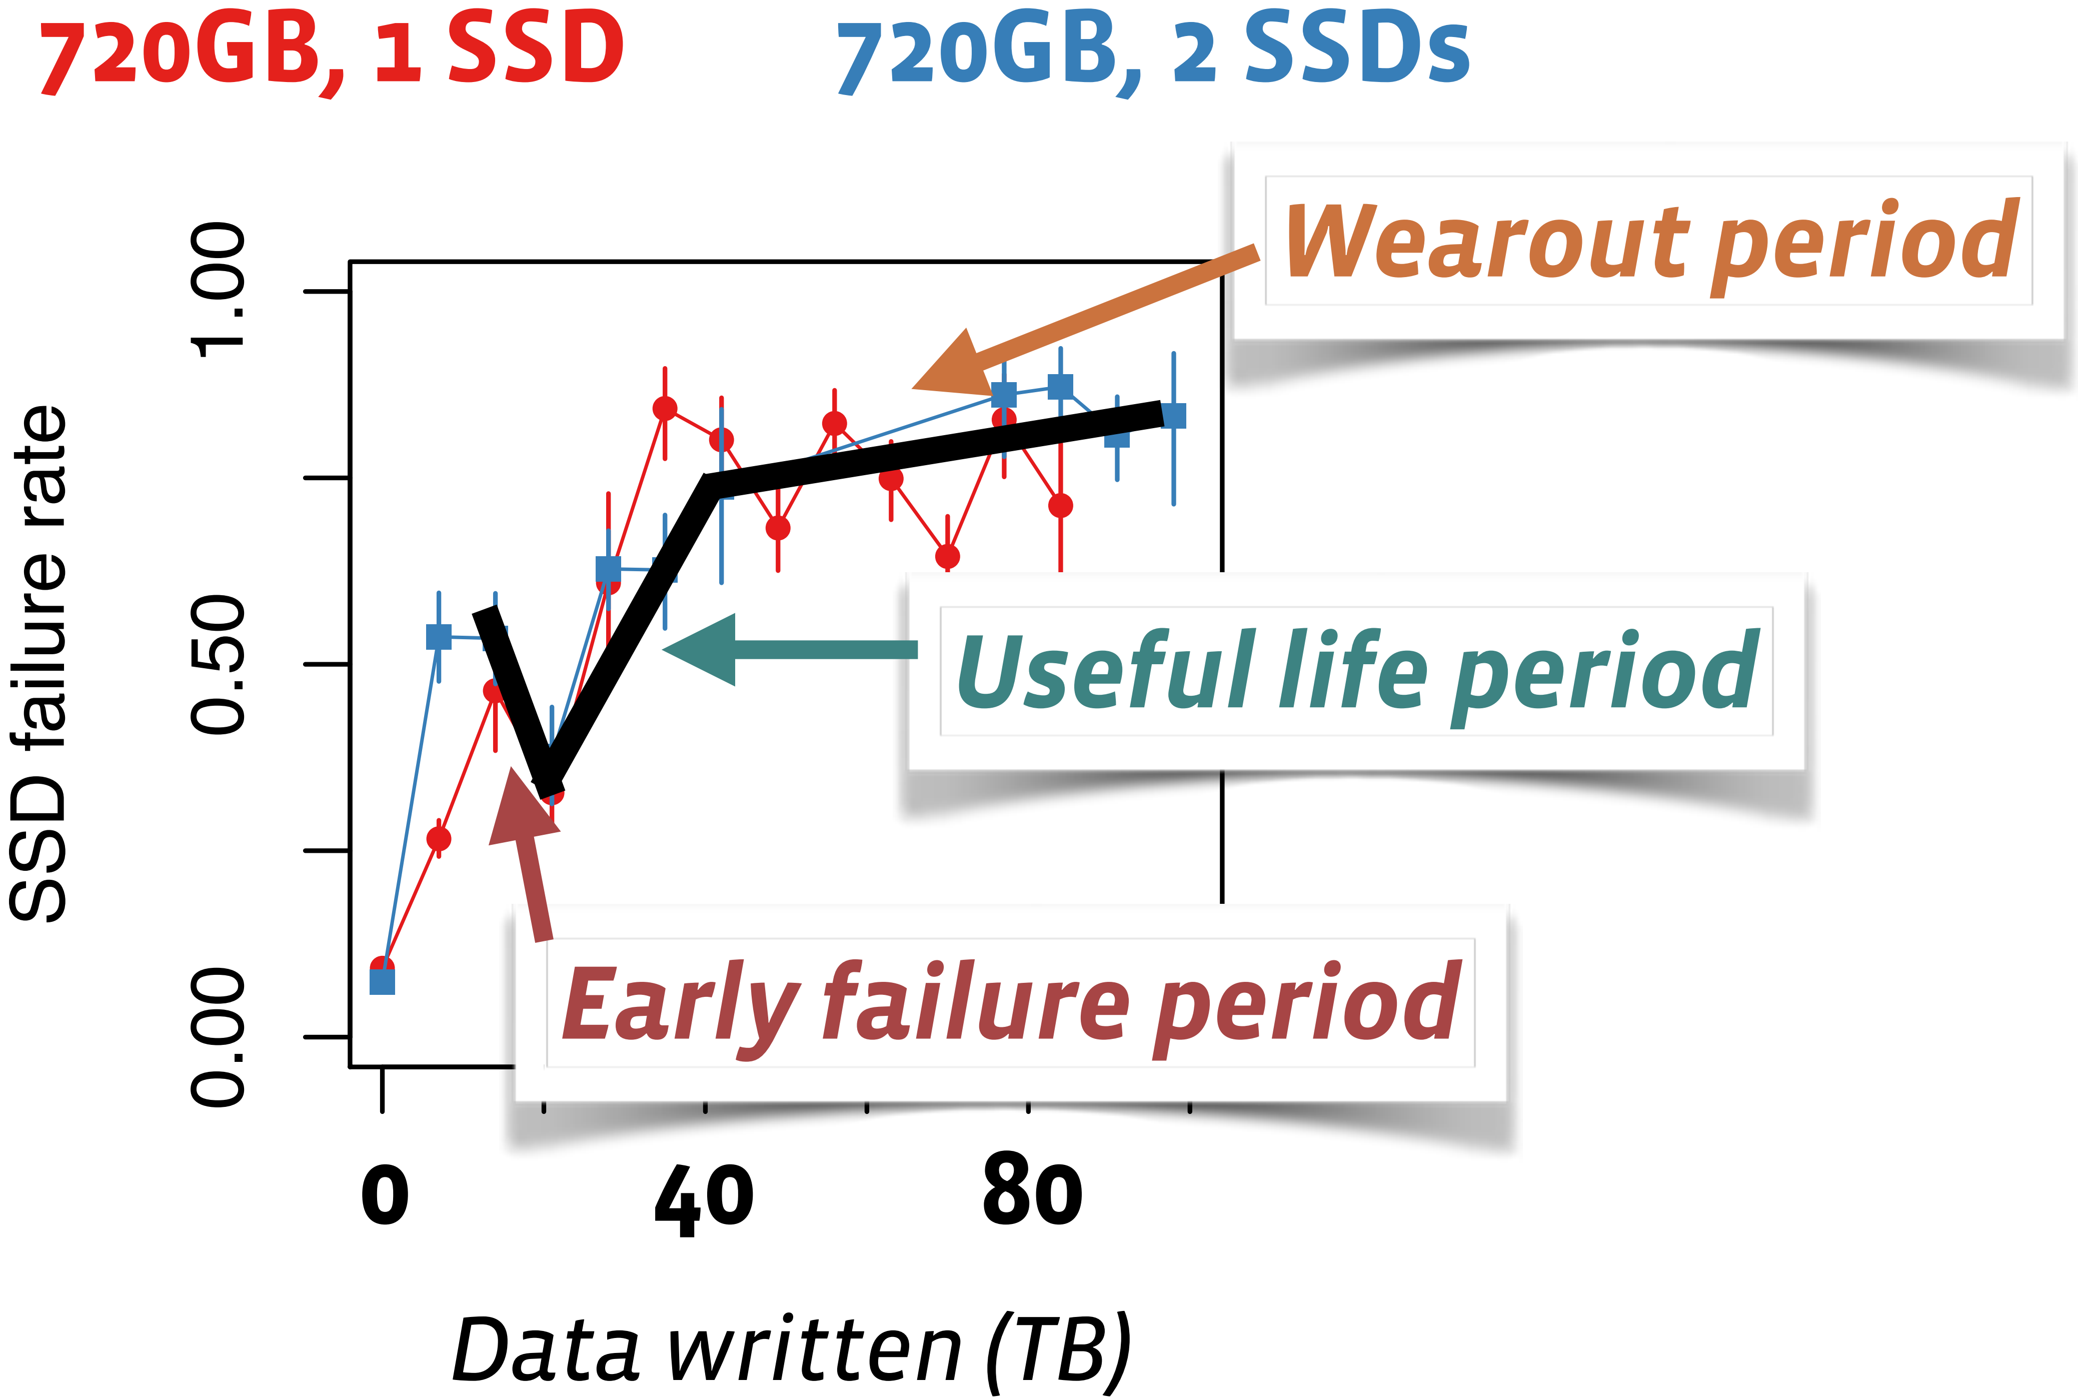

Wearout period

Useful life period

Early failure period

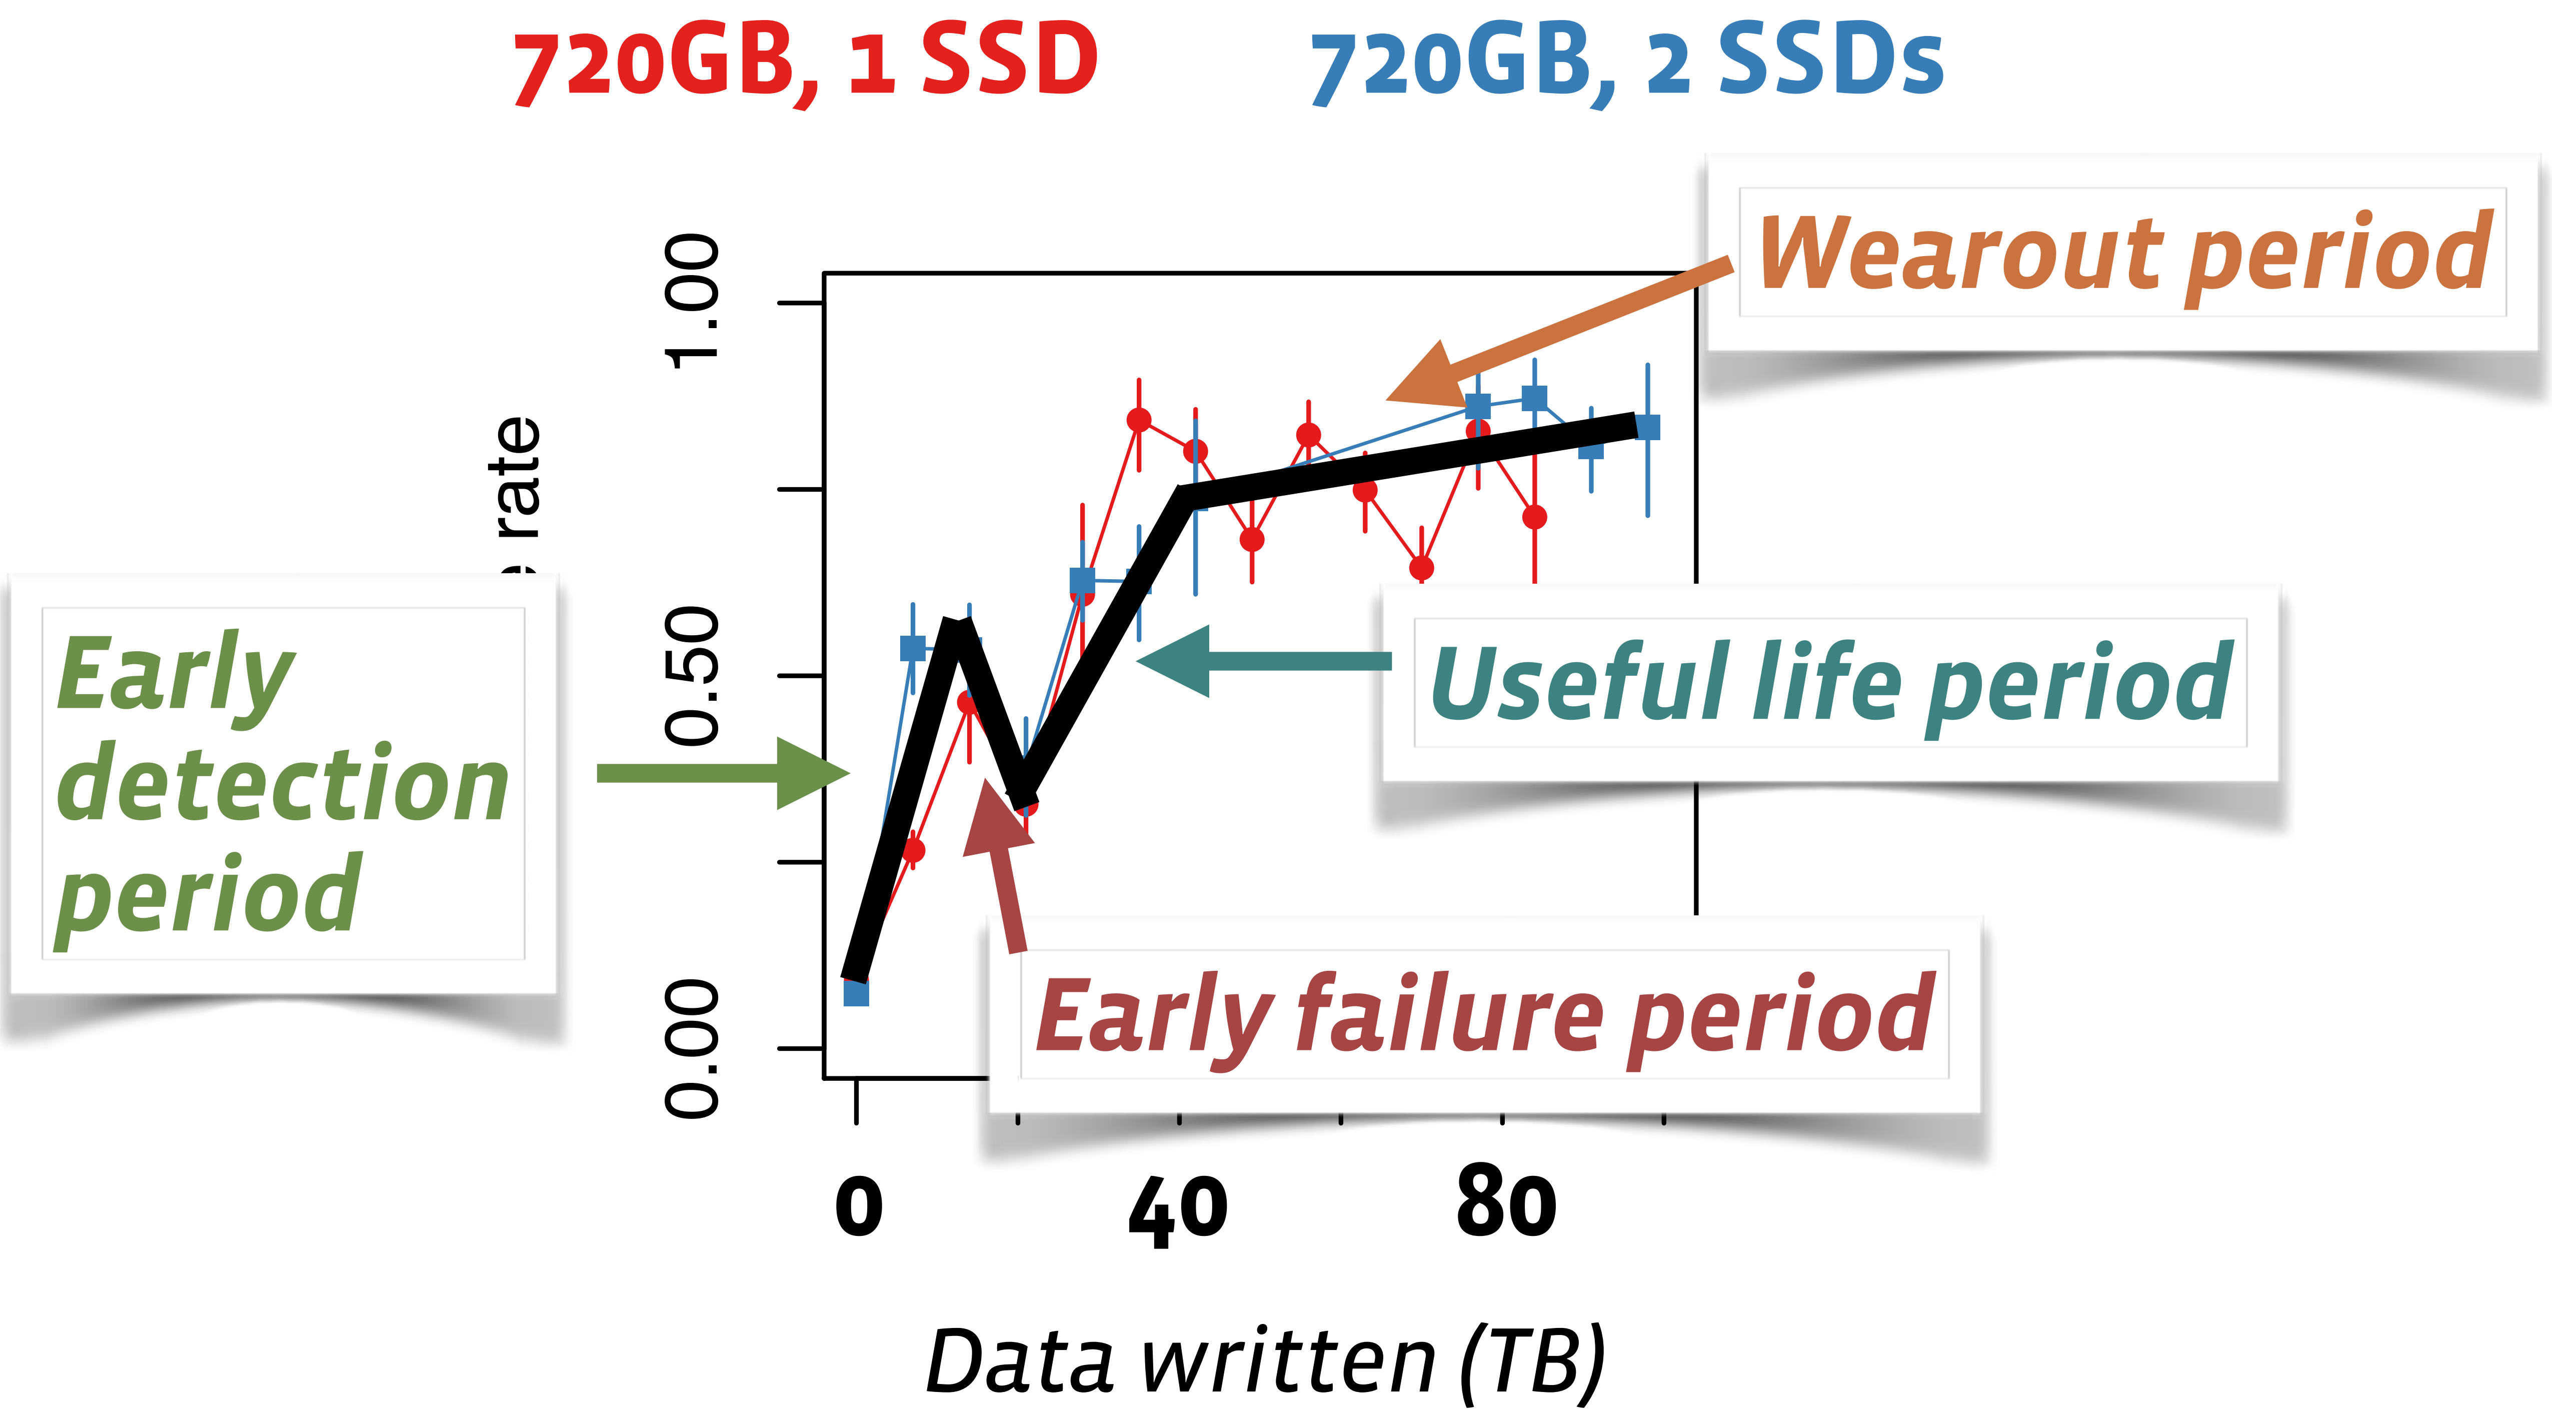

Wearout period

Early detection period

Useful life period

Early failure period

SSD lifecycle

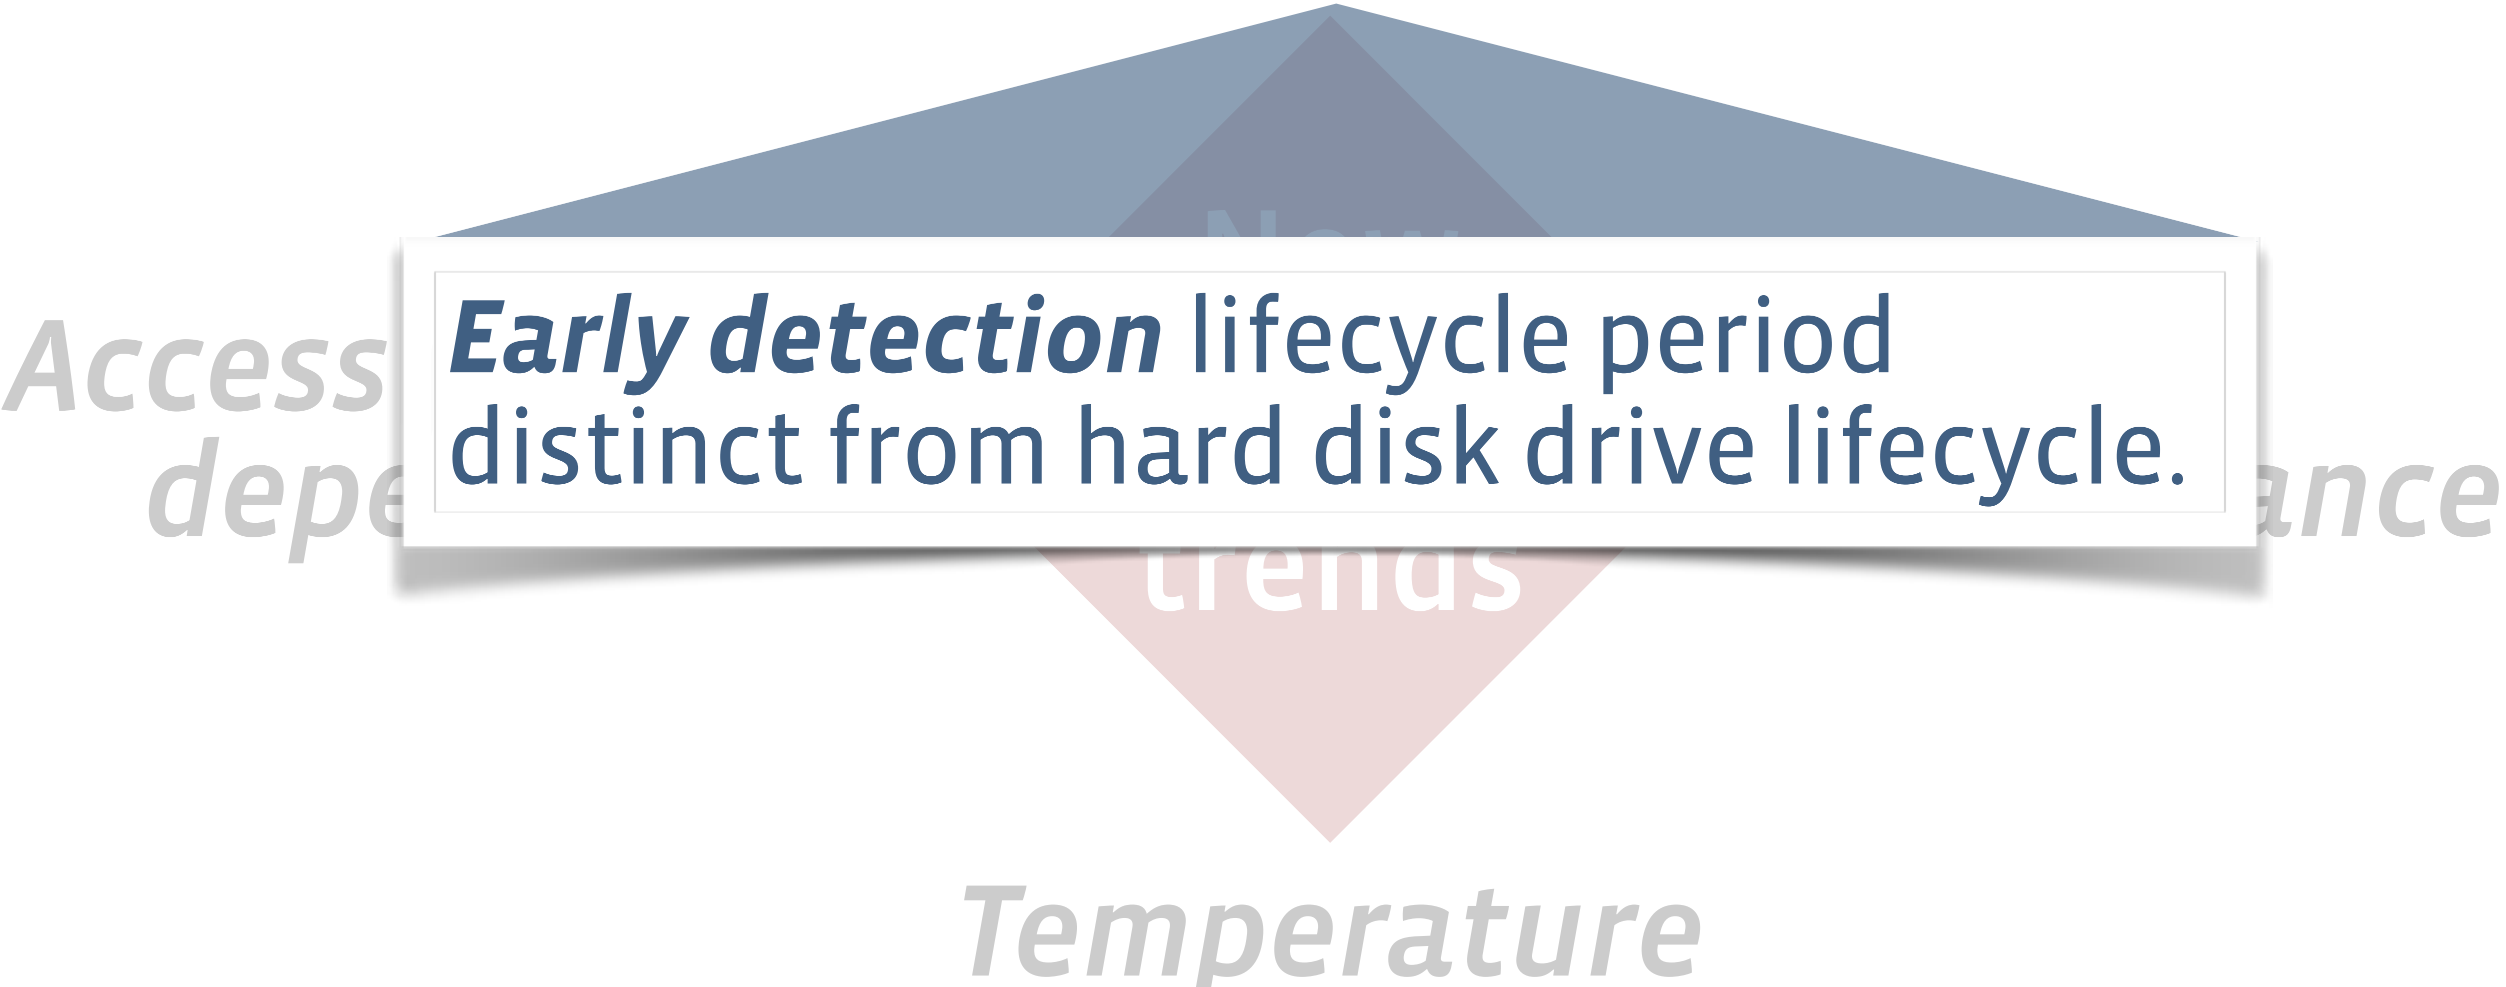

Early detection lifecycle period distinct from hard disk drive lifecycle.

Why Early Detection Period?

- Two pool model of flash blocks: weak and strong

- Weak ones fail early → increasing failure rate early in lifetime
  - SSD takes them offline → lowers the overall failure rate

- Strong ones fail late → increasing failure rate late in lifetime

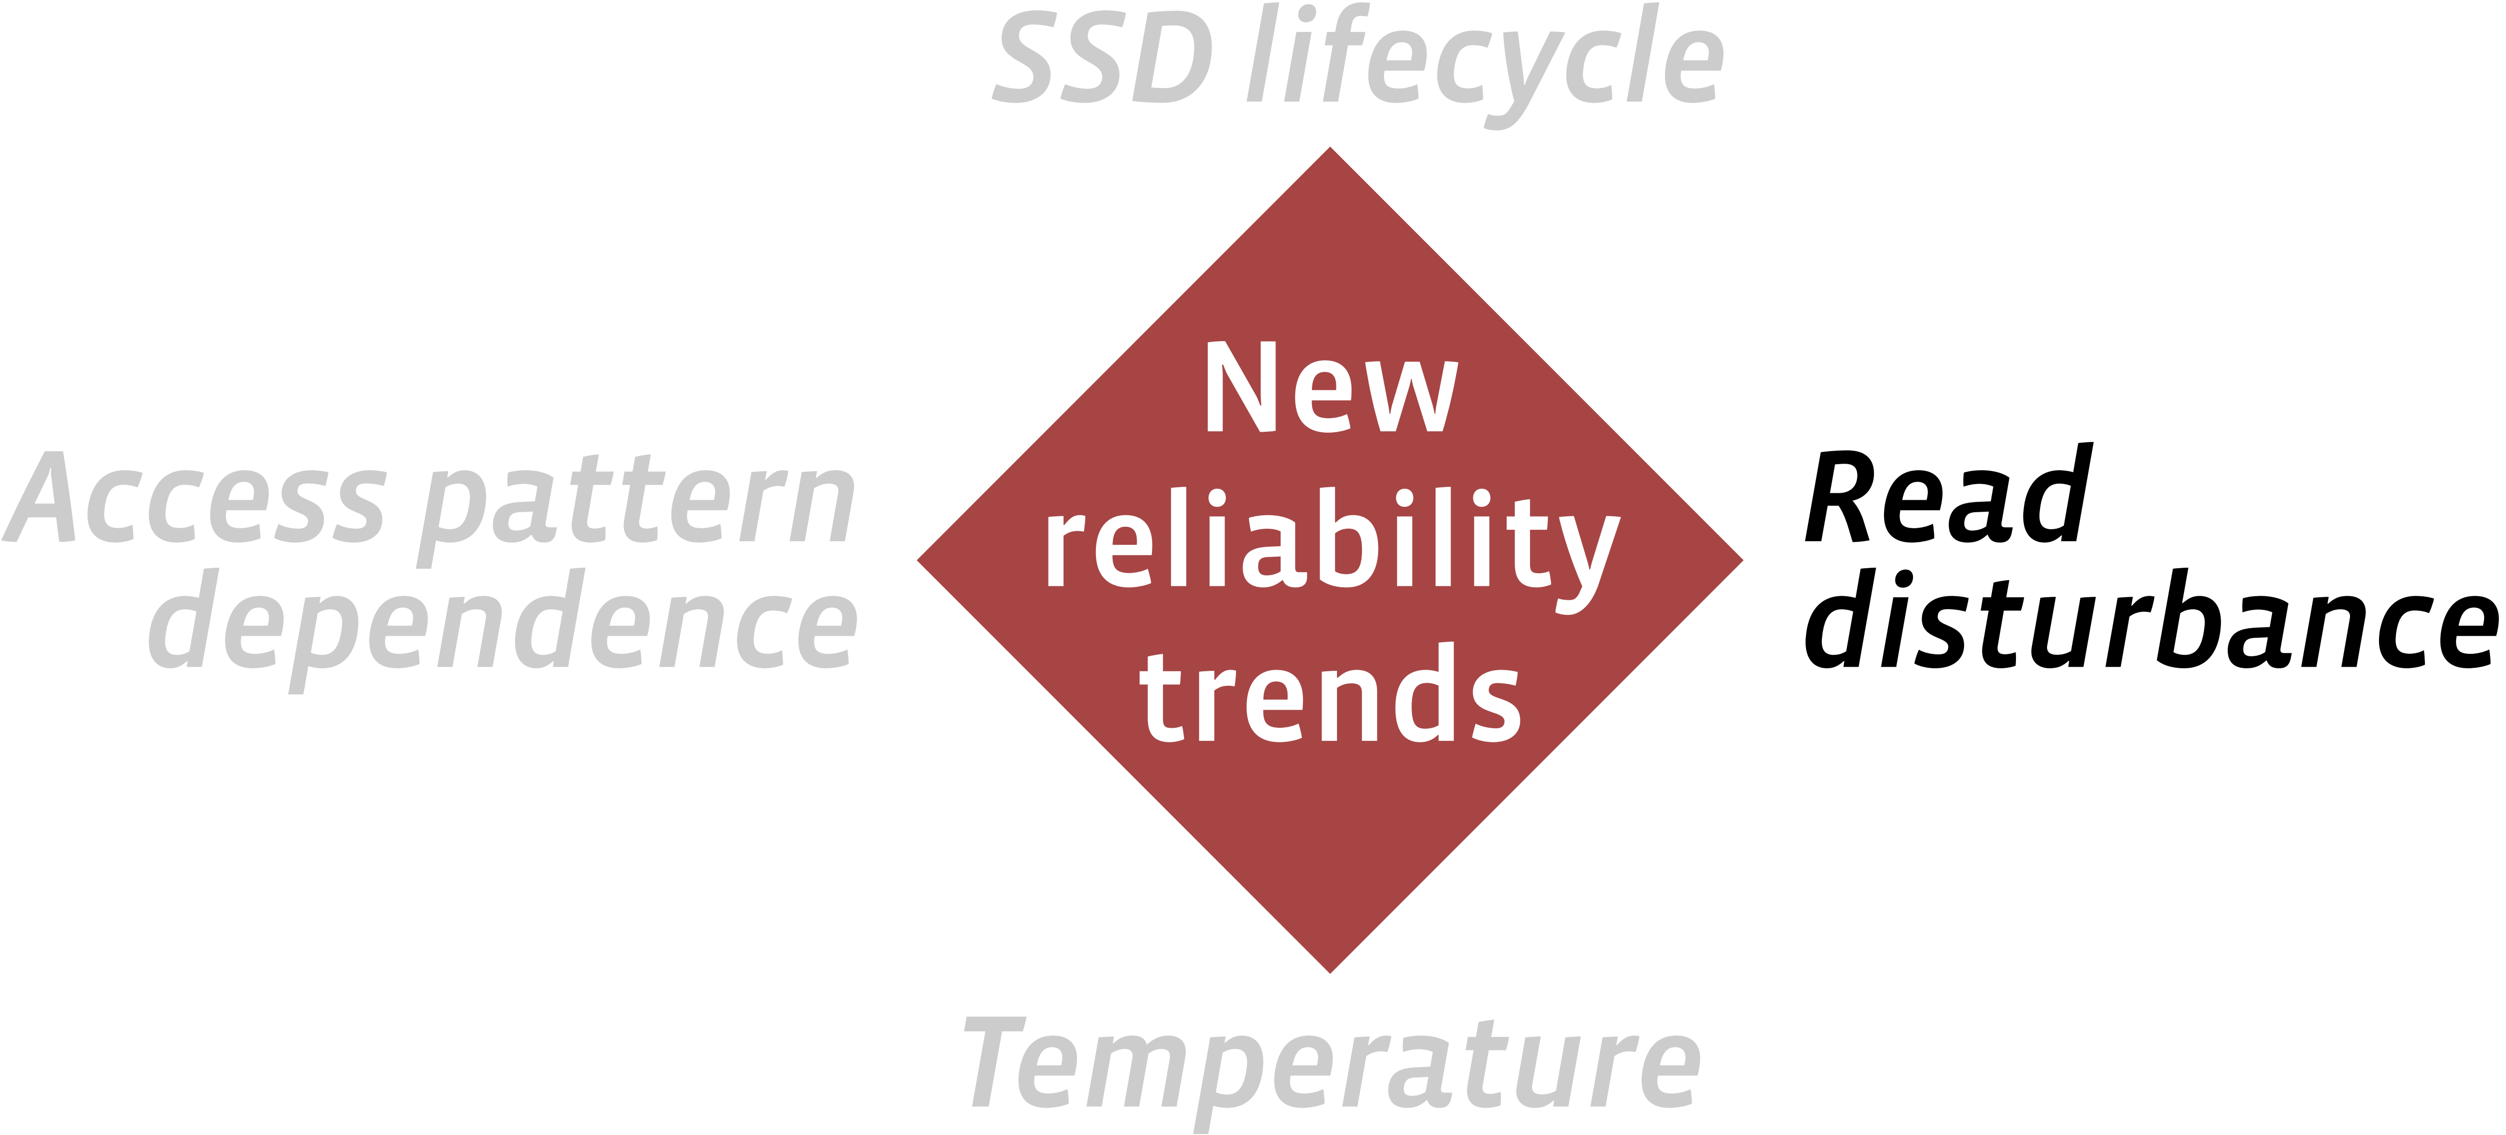

SSD lifecycle

Access pattern dependence

New reliability trends

Read disturbance

Temperature

Read disturbance

- reading data can disturb contents
- failure mode identified in ***lab setting***
- under ***adversarial workloads***

Read from Flash Cell Array

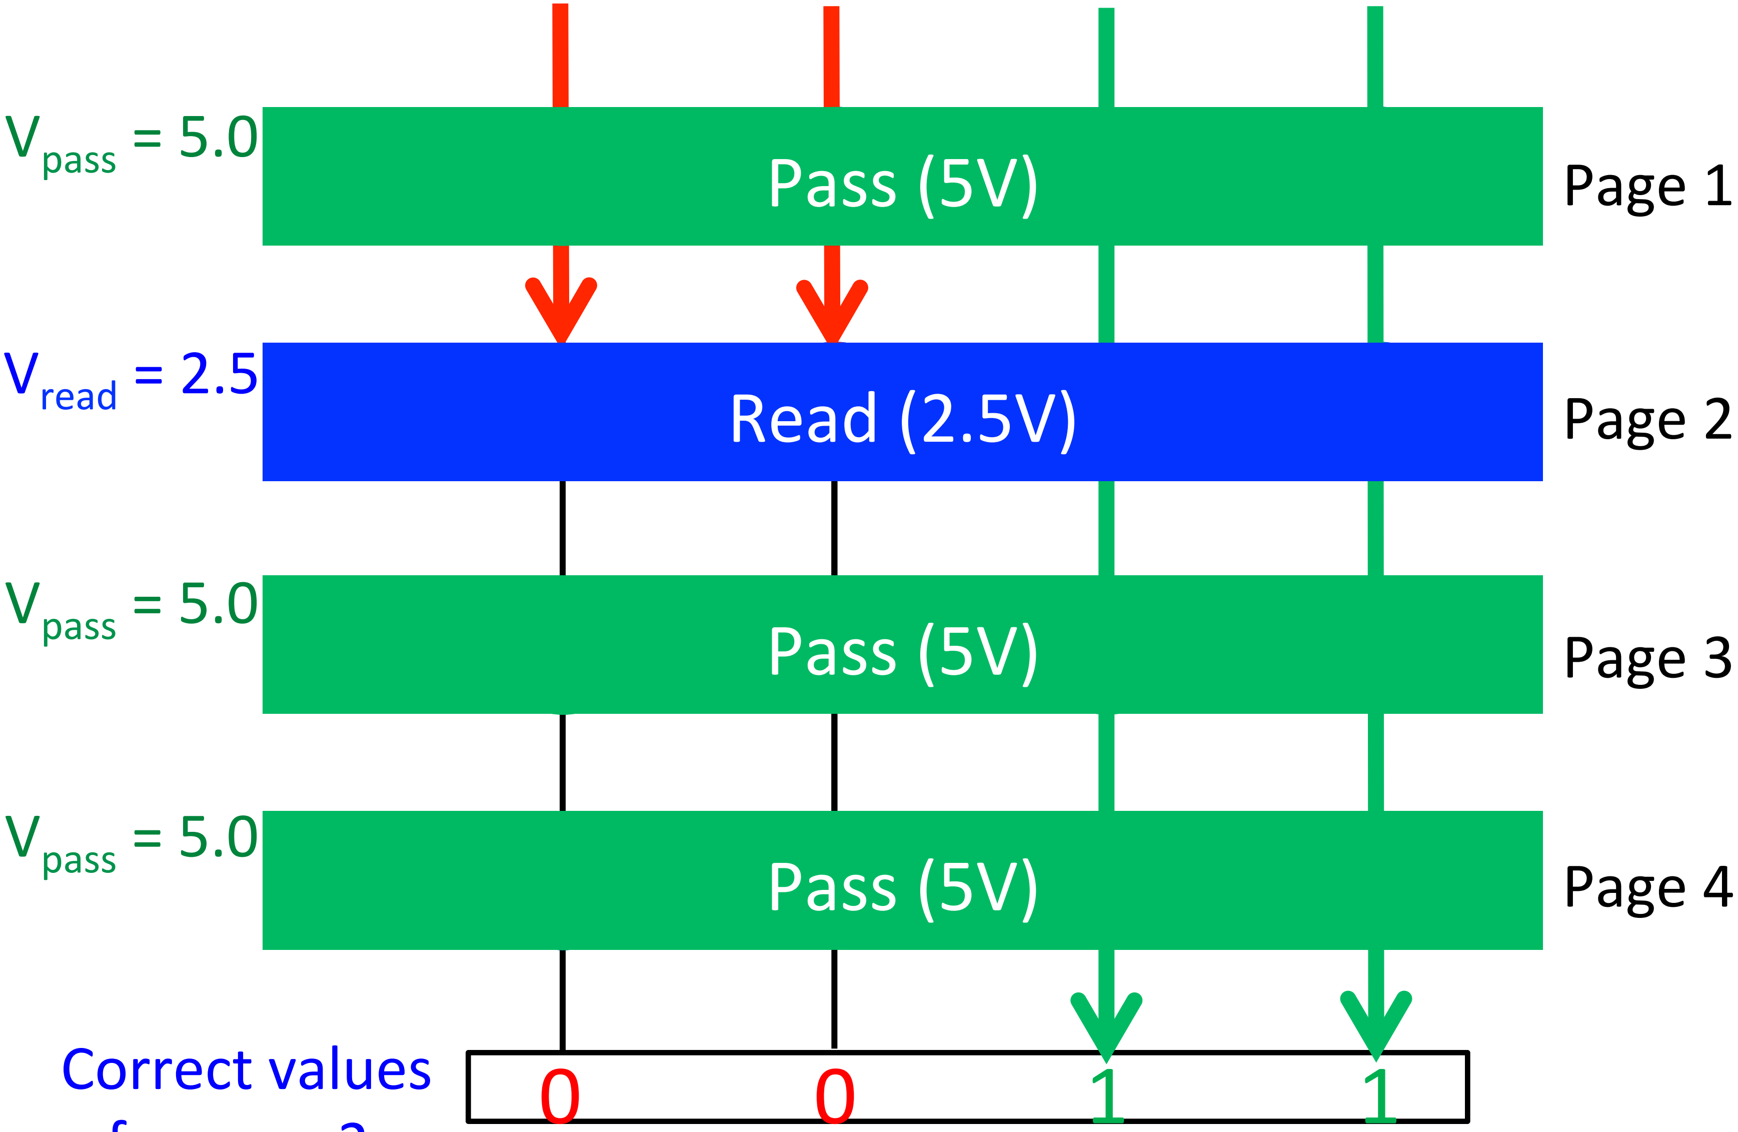

V_{pass} = 5.0

Pass (5V)

Page 1

V_{read} = 2.5

Read (2.5V)

Page 2

V_{pass} = 5.0

Pass (5V)

Page 3

V_{pass} = 5.0

Pass (5V)

Page 4

Correct values for page 2:

0 0 1 1

Read Disturb Problem: “Weak Programming” Effect

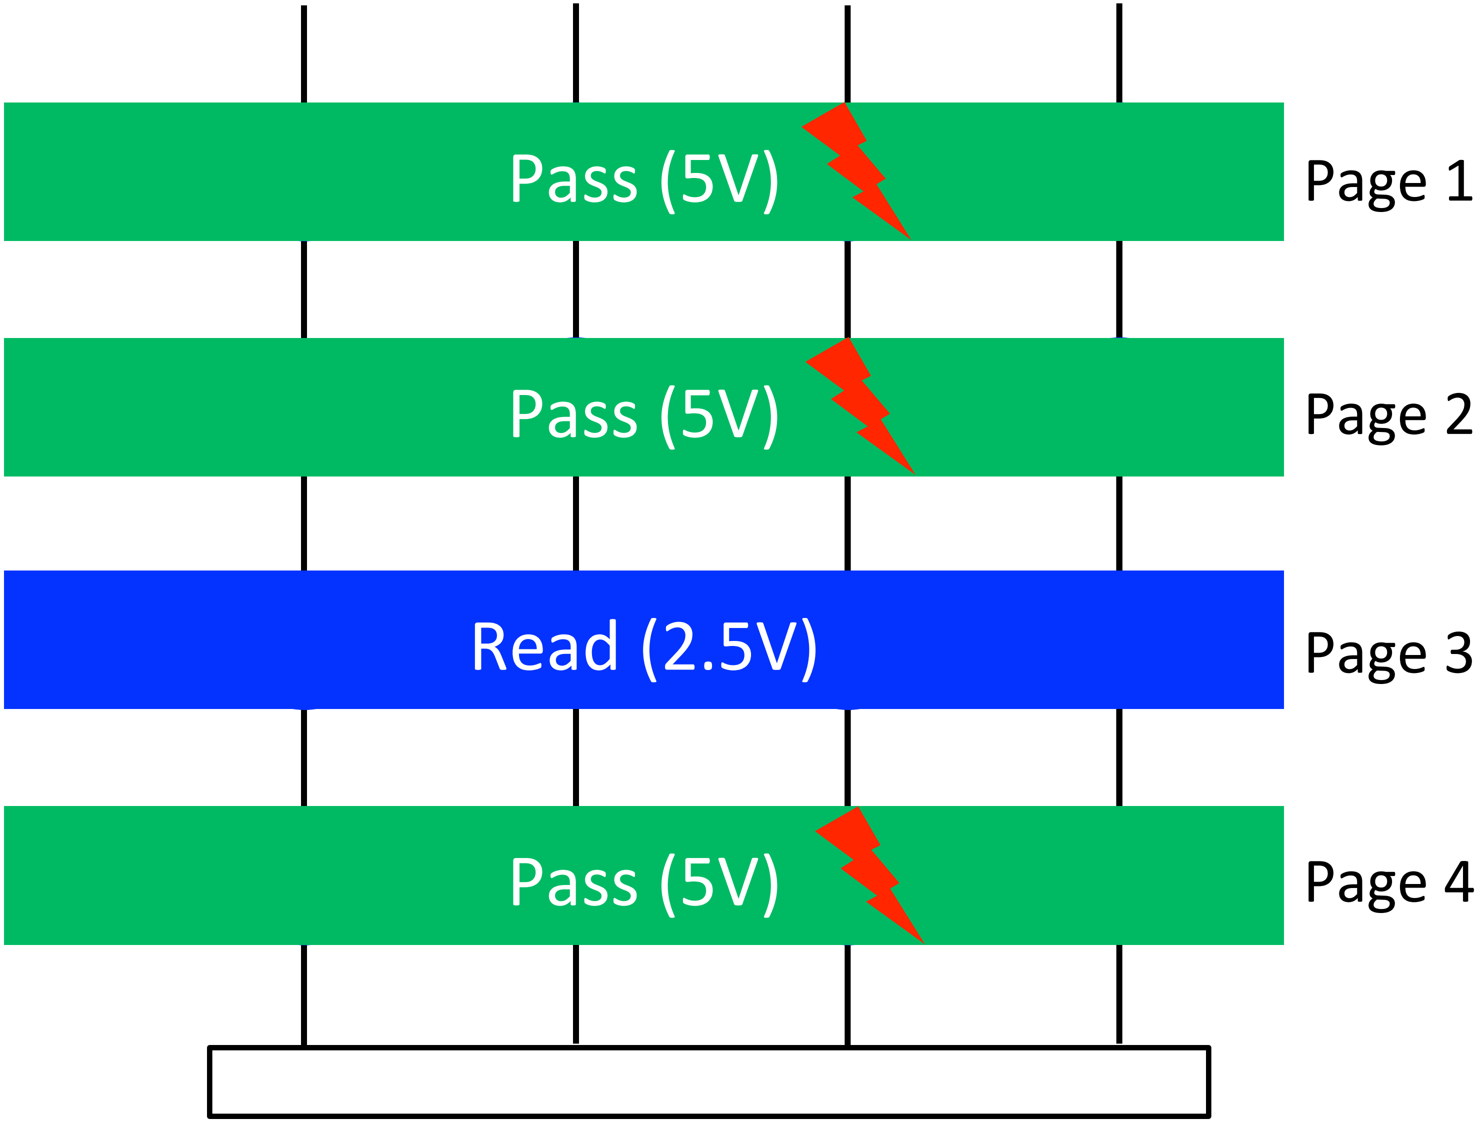

Repeatedly read page 3 (or any page other than page 2)

More on Flash Read Disturb Errors

- Yu Cai, Yixin Luo, Saugata Ghose, Erich F. Haratsch, Ken Mai, and Onur Mutlu,
  "Read Disturb Errors in MLC NAND Flash Memory: Characterization and Mitigation"
  Proceedings of the
  45th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (***DSN***), Rio de Janeiro, Brazil, June 2015.

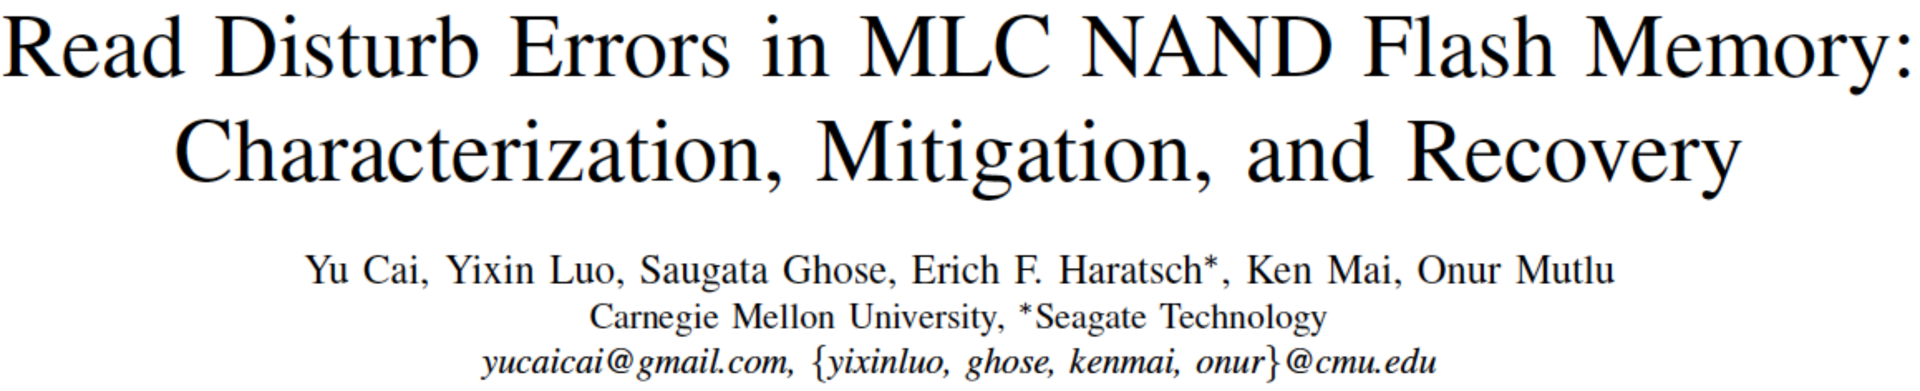

Read Disturb Errors in MLC NAND Flash Memory: Characterization, Mitigation, and Recovery

Yu Cai, Yixin Luo, Saugata Ghose, Erich F. Haratsch*, Ken Mai, Onur Mutlu
Carnegie Mellon University, *Seagate Technology
yucaicai@gmail.com, {yixinluo, ghose, kenmai, onur}@cmu.edu

Read disturbance

Does read disturbance affect SSDs in the field?

-
- *ng*
- under *adversarial workloads*

Examine SSDs with

high **flash R/W** *ratios*

and **most data read**

to understand read effects

(isolate effects of read vs. write errors)

3.2TB, 1 SSD (average R/W = 2.14)

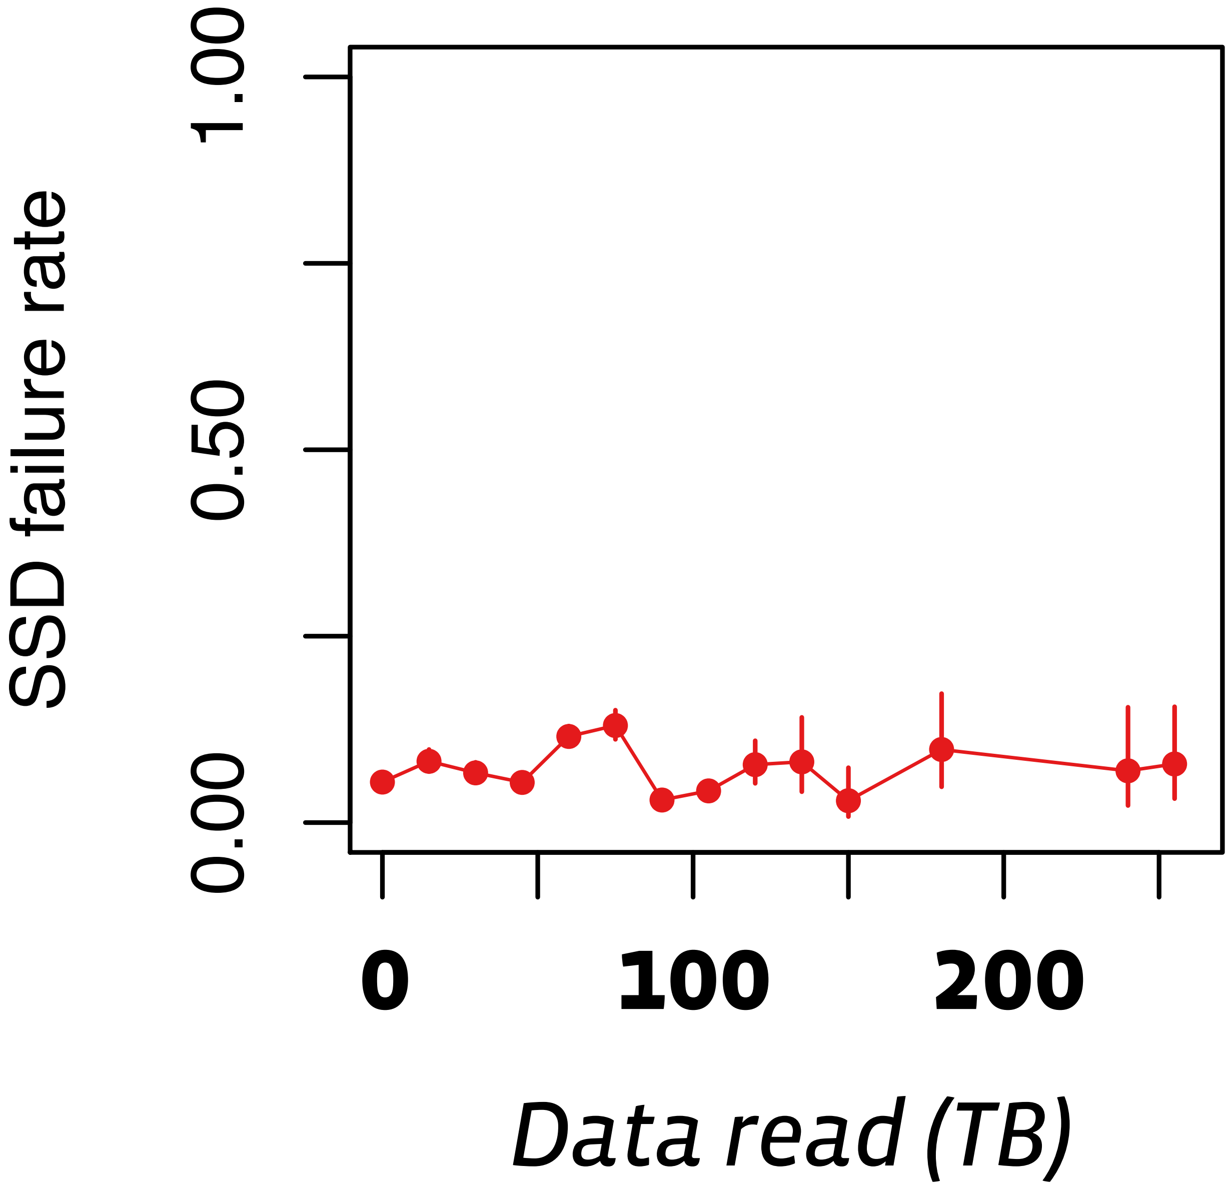

1.2TB, 1 SSD (average R/W = 1.15)

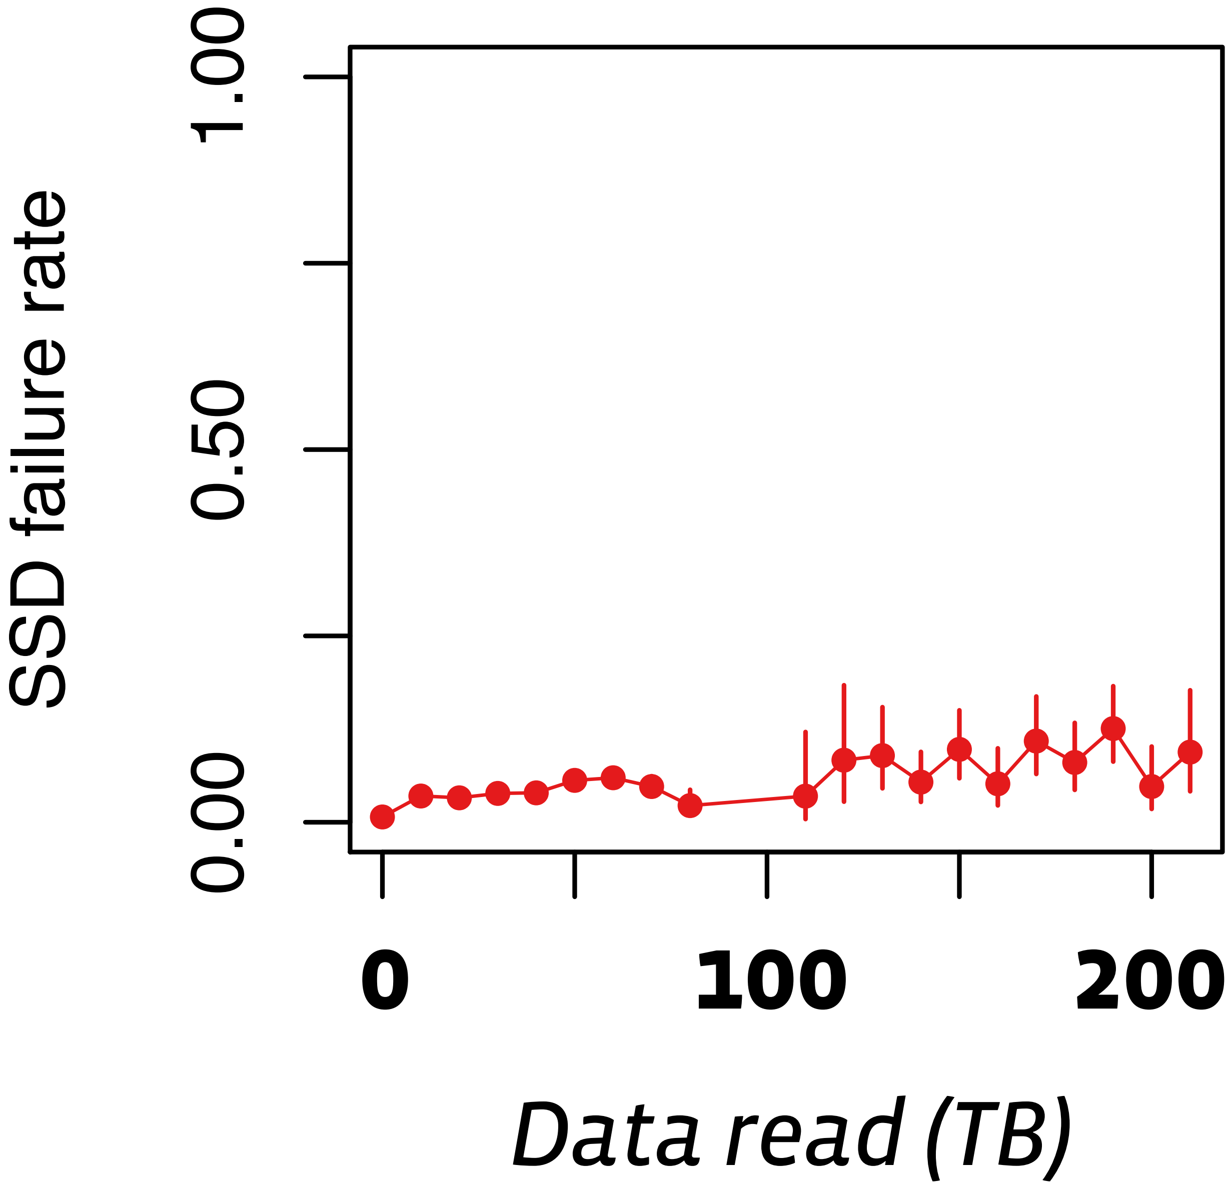

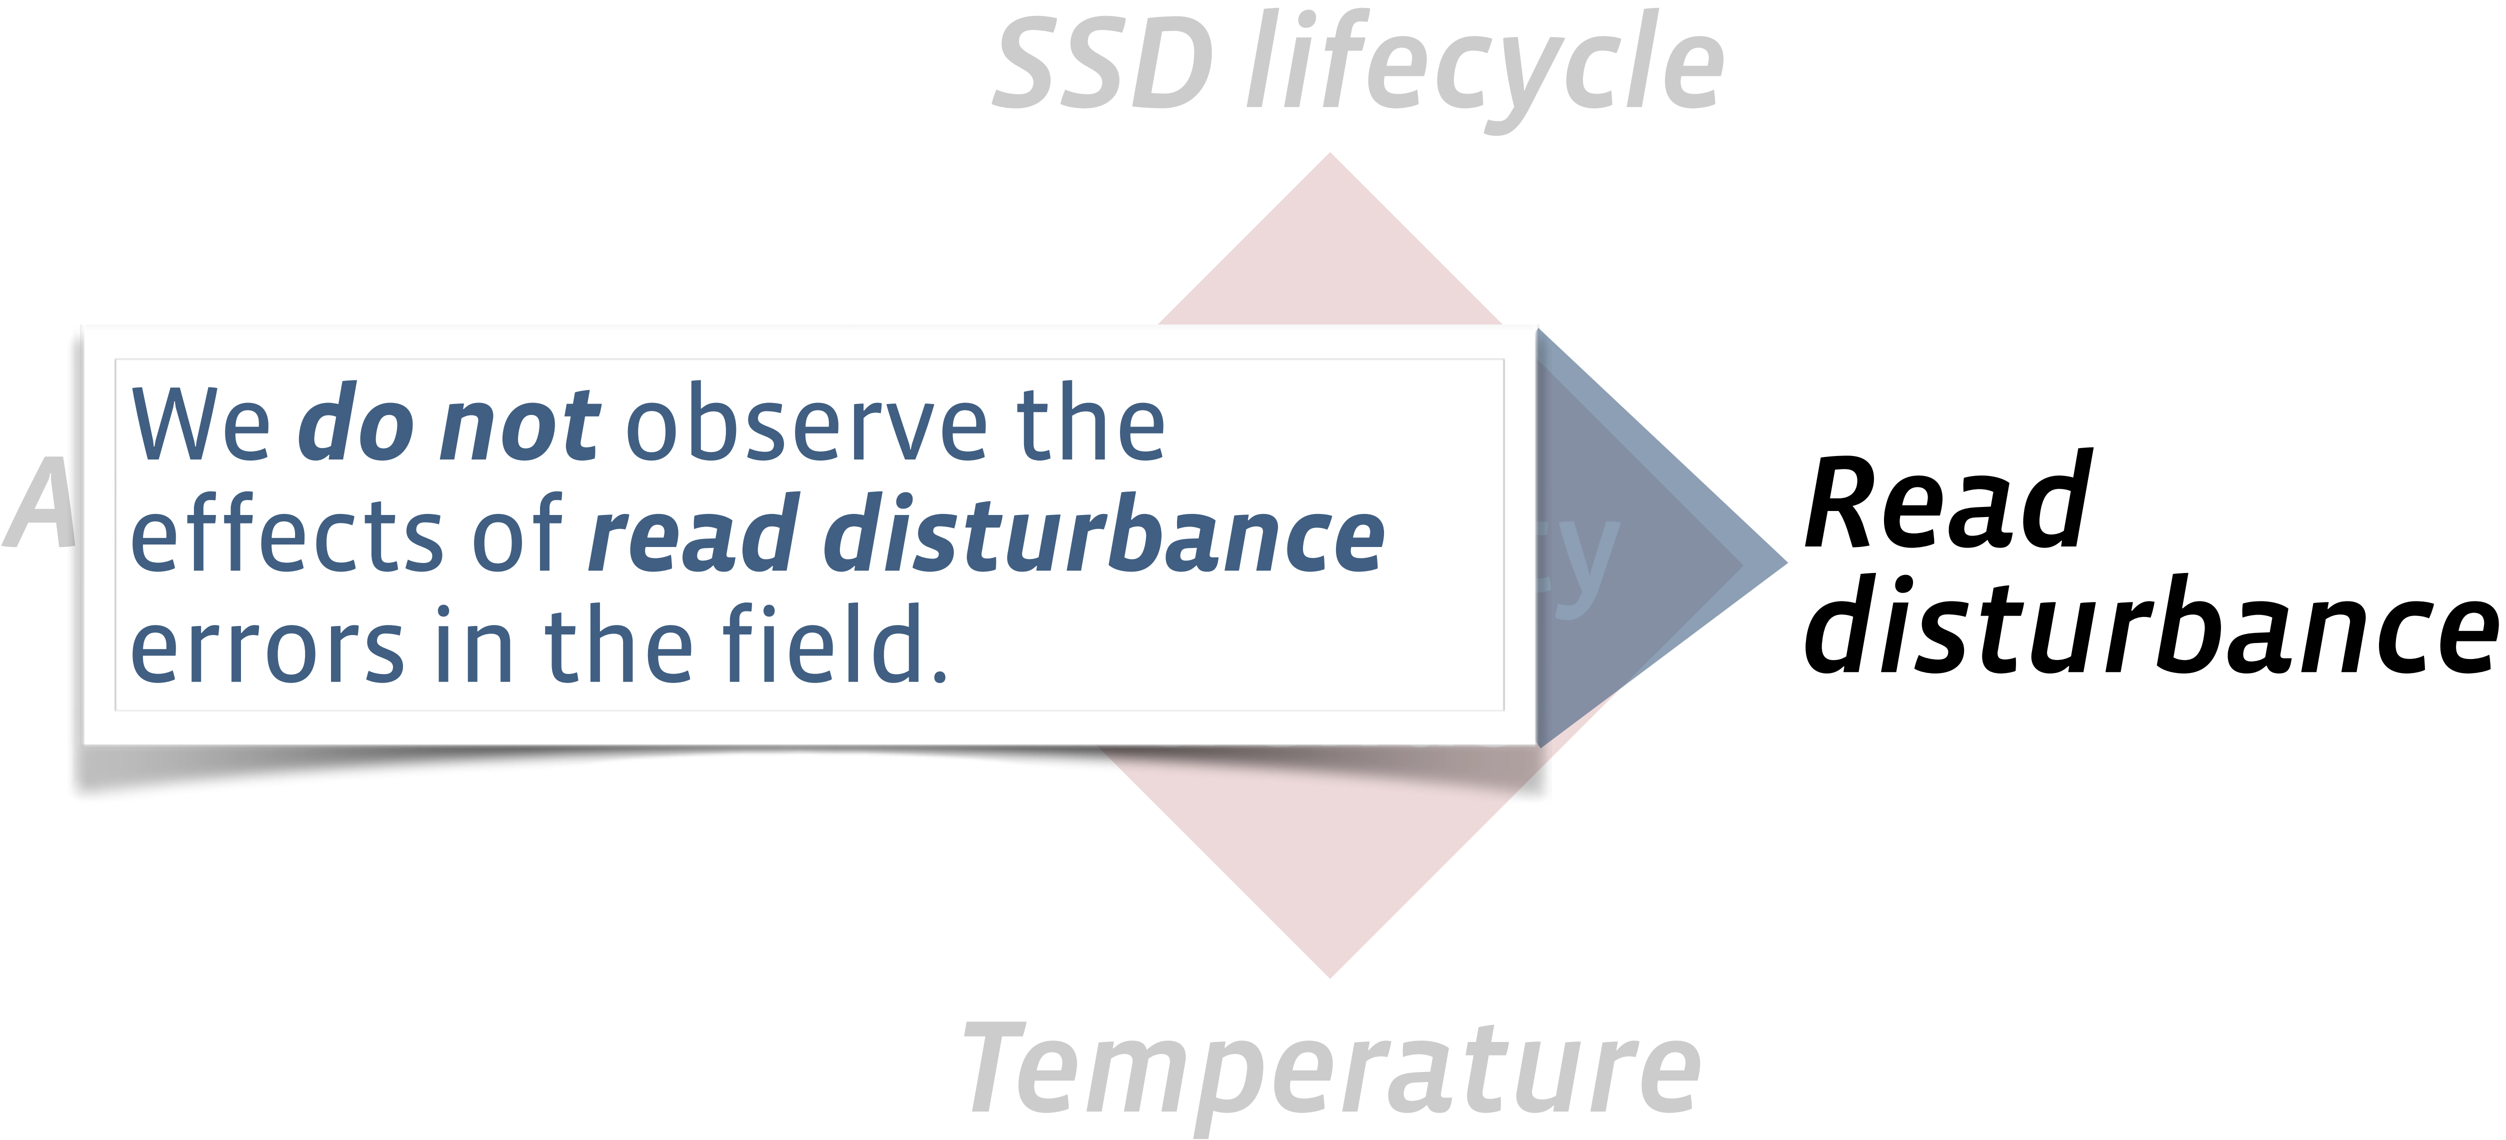

SSD lifecycle

Temperature

Read disturbance

We ***do not*** observe the effects of ***read disturbance*** errors in the field.

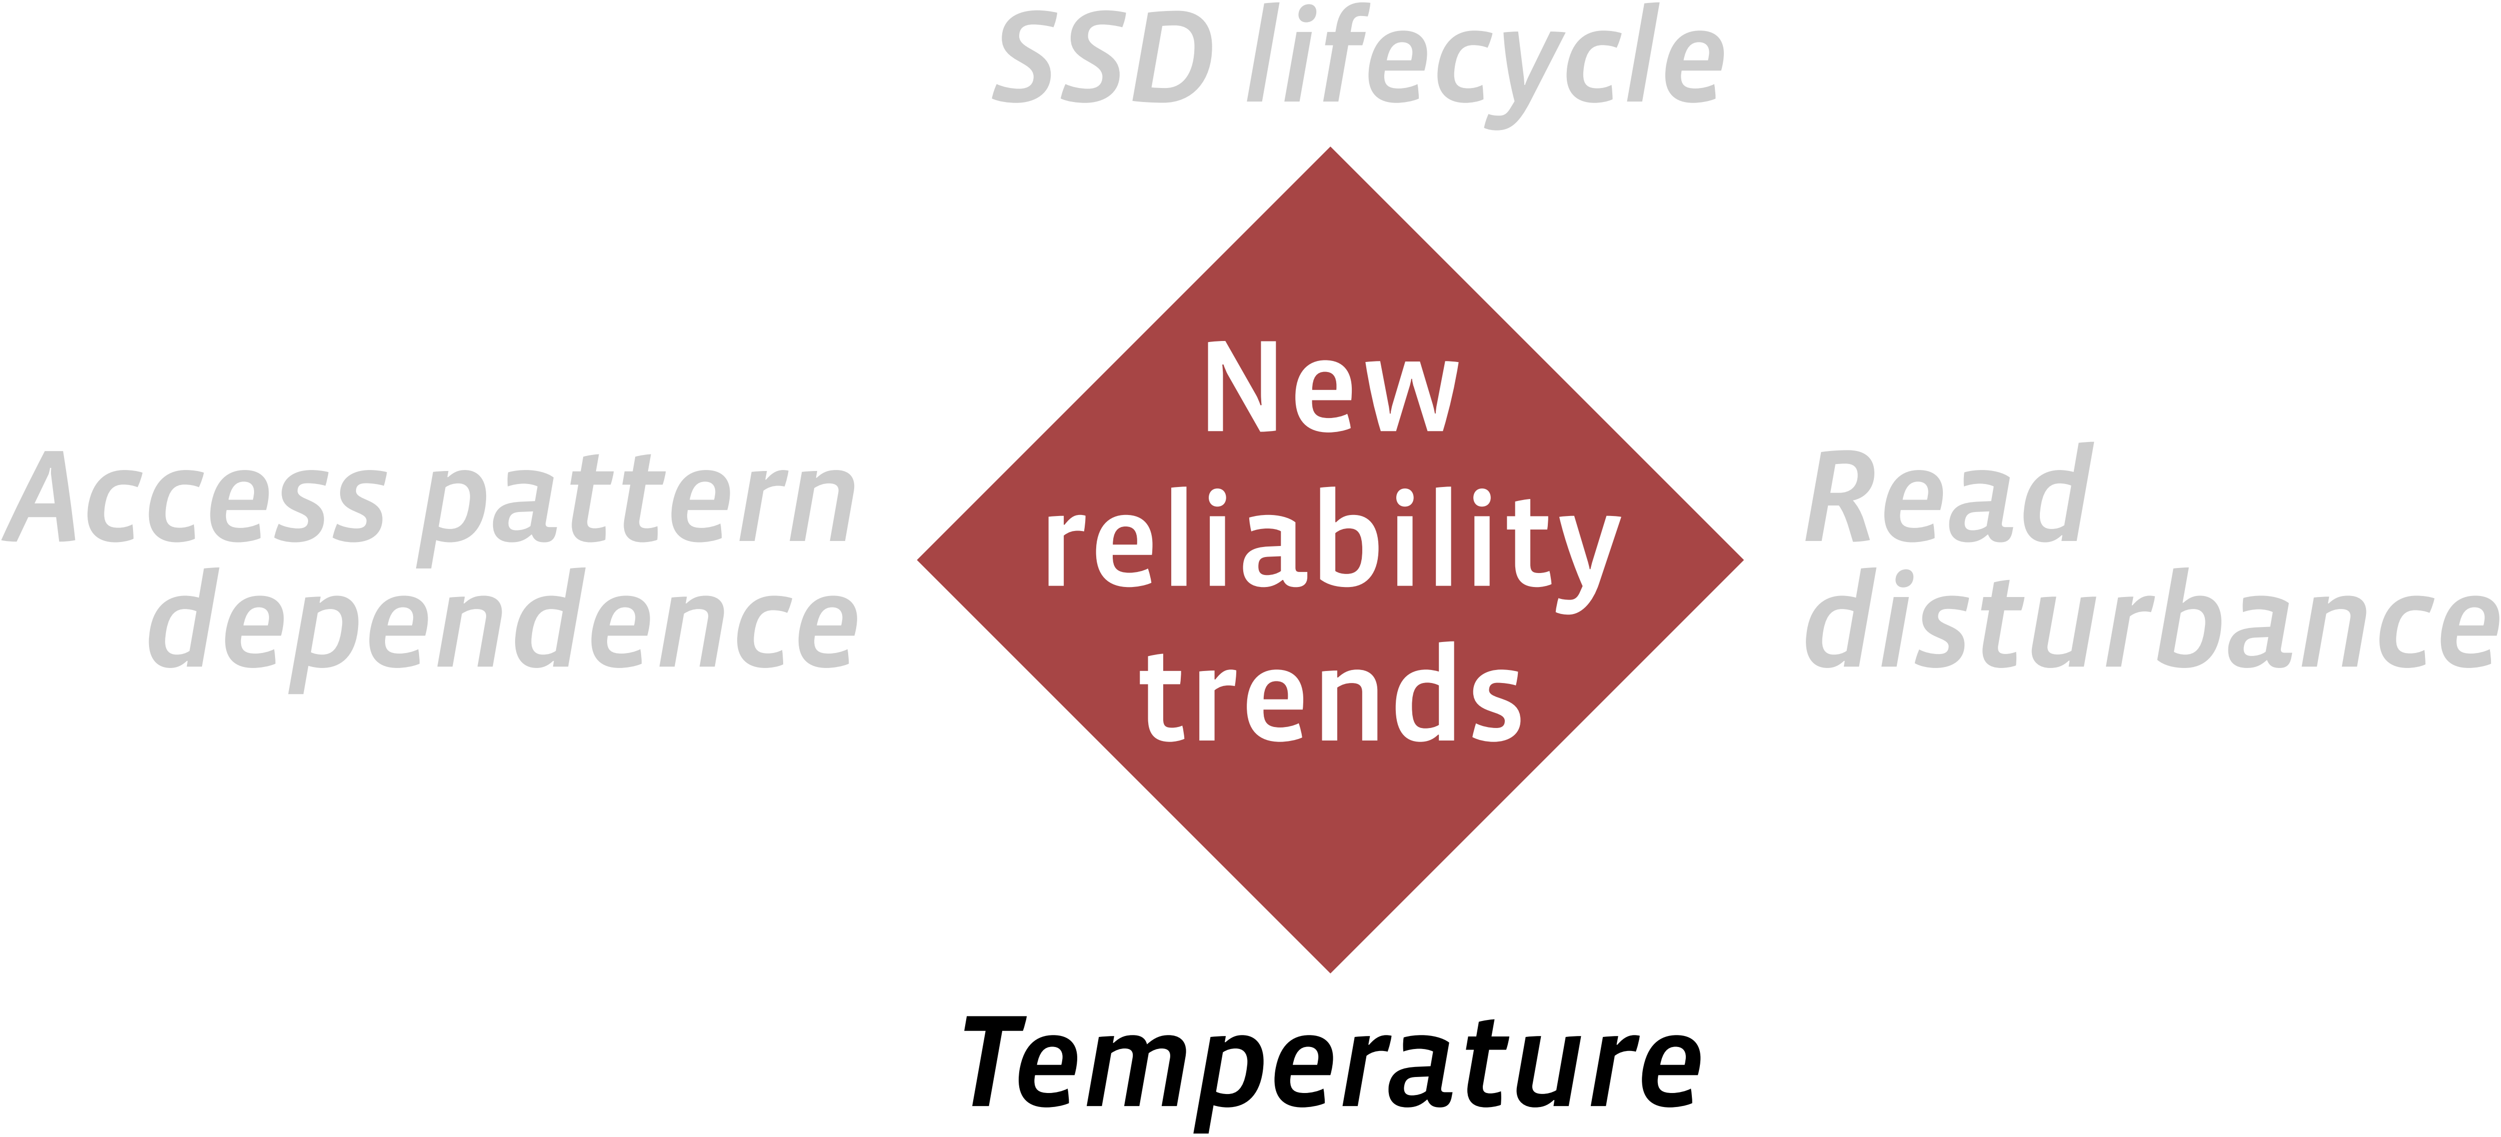

SSD lifecycle

Access pattern dependence

New reliability trends

Read disturbance

Temperature

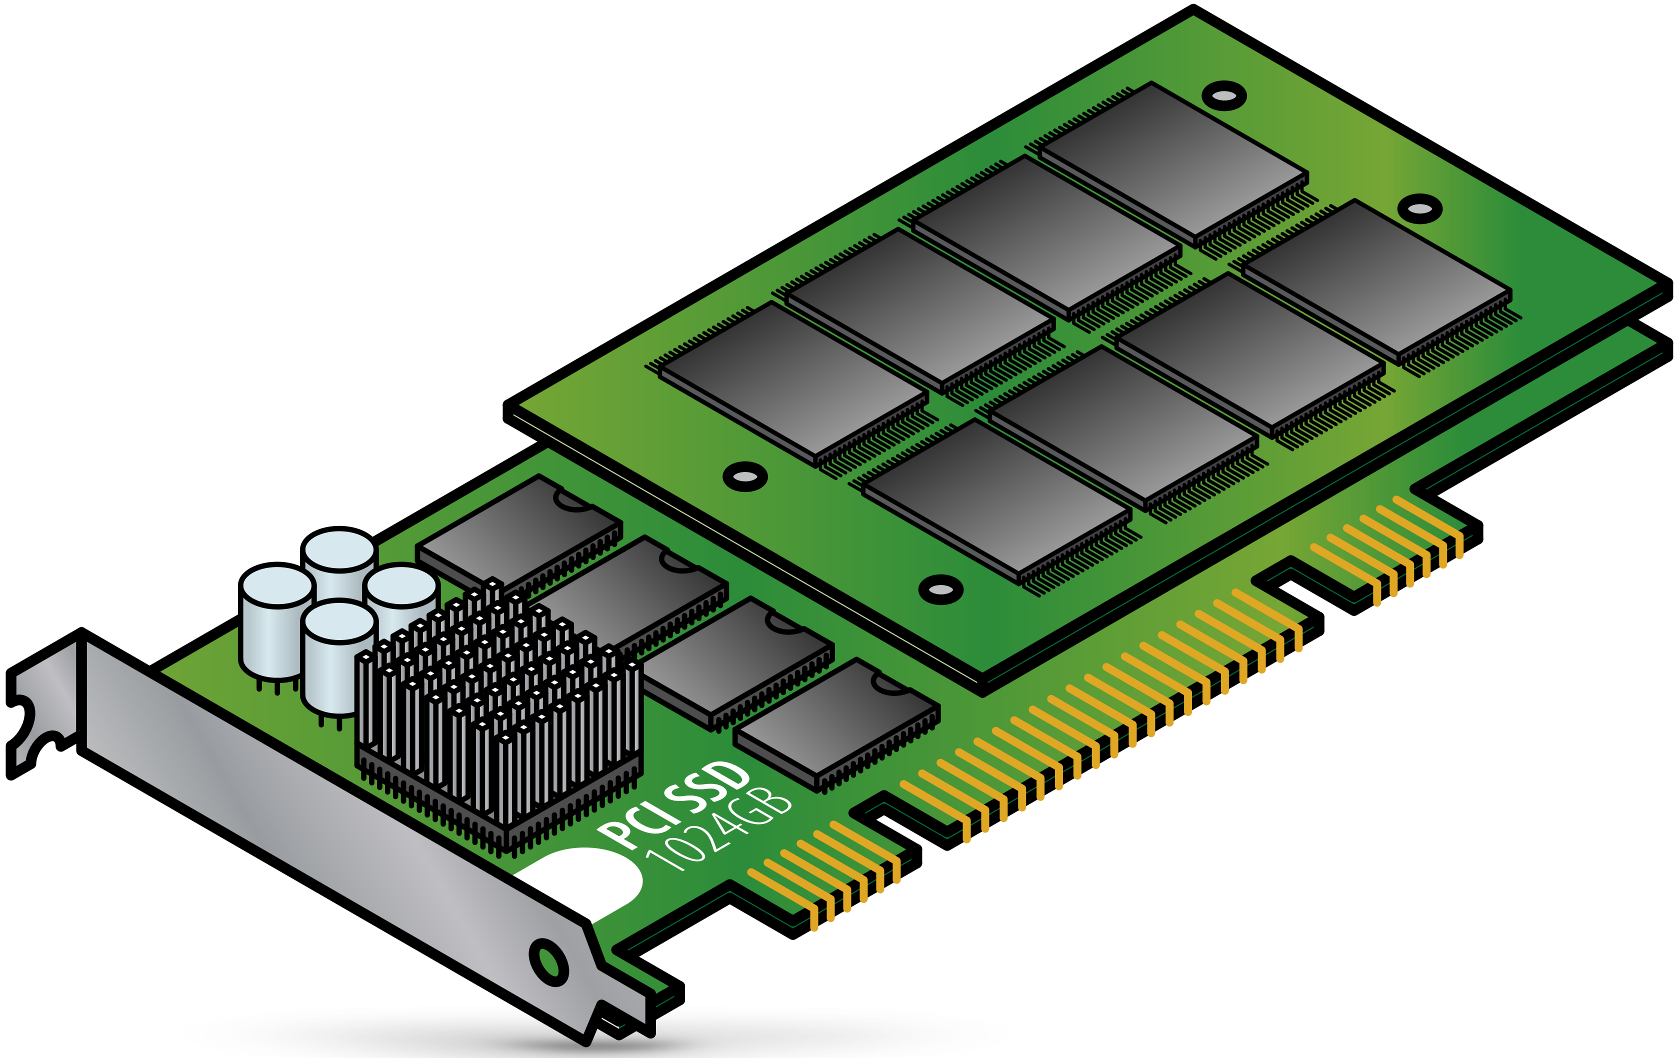
PCI SSD
1024GB

Temperature sensor

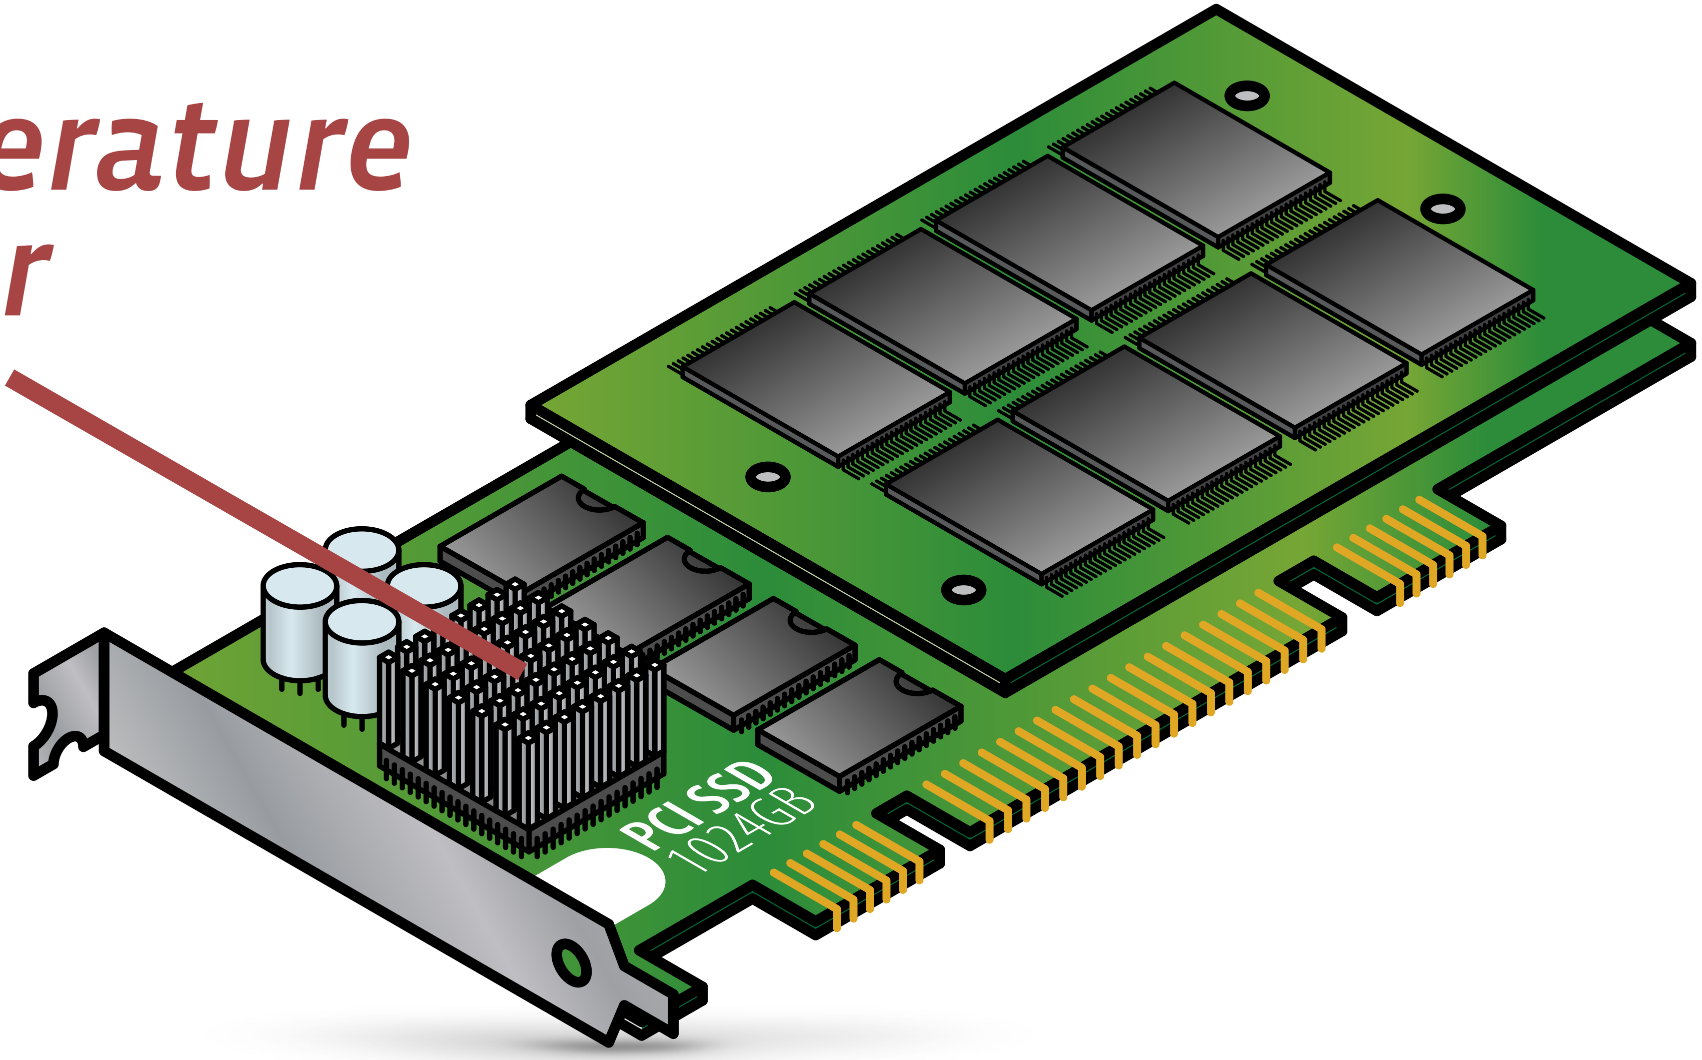

Three Failure Rate Trends with Temperature

- Increasing
  - SSD not throttled

- Decreasing after some temperature
  - SSD could be throttled

- Not sensitive
  - SSD could be throttled

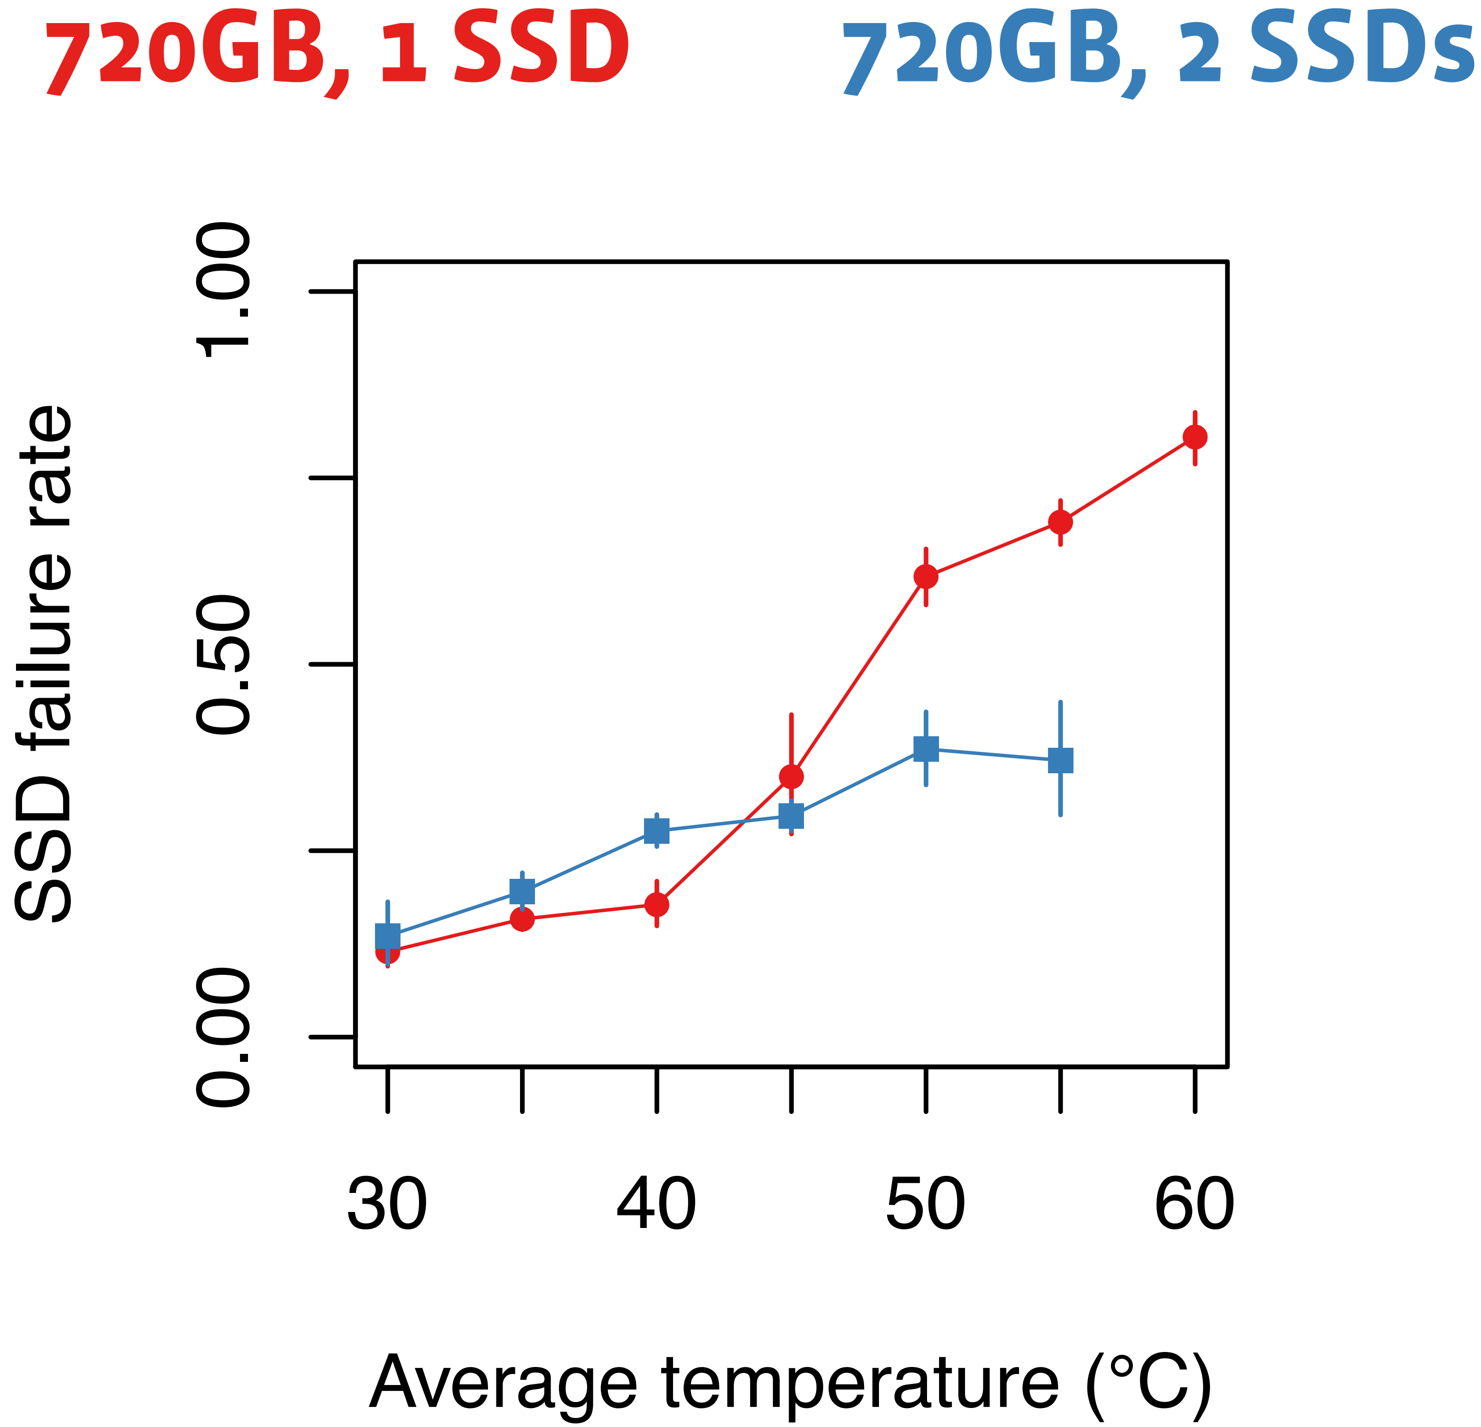
720GB, 1 SSD
720GB, 2 SSDs
SSD failure rate
0.00
0.50
1.00
30
40
50
60
Average temperature (°C)

No Throttling on A & B SSDs

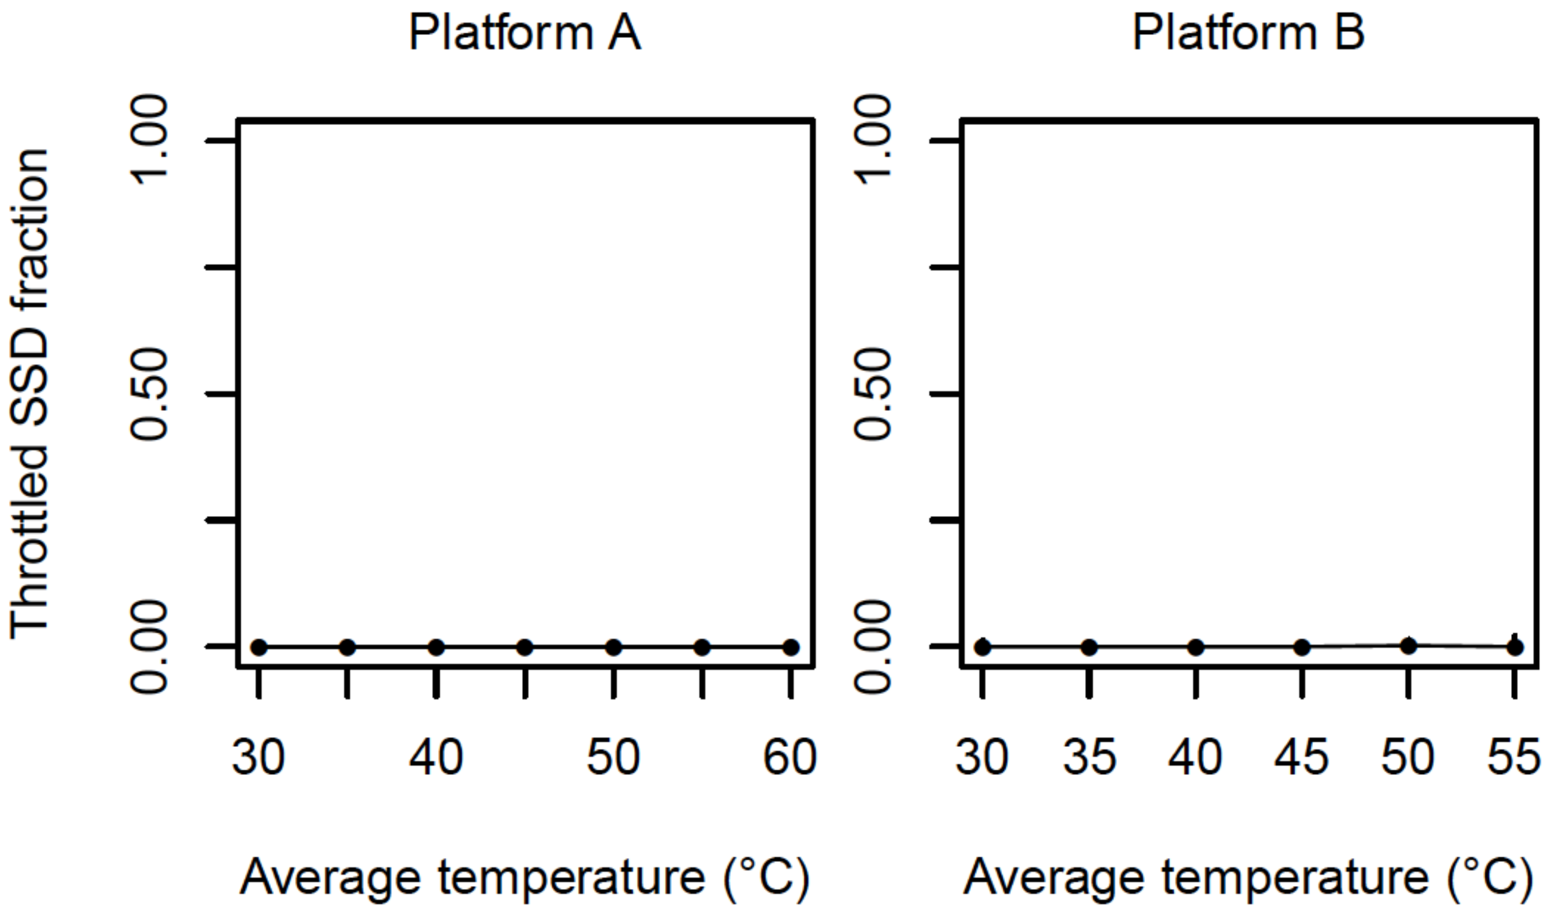

High temperature: may throttle or shut down

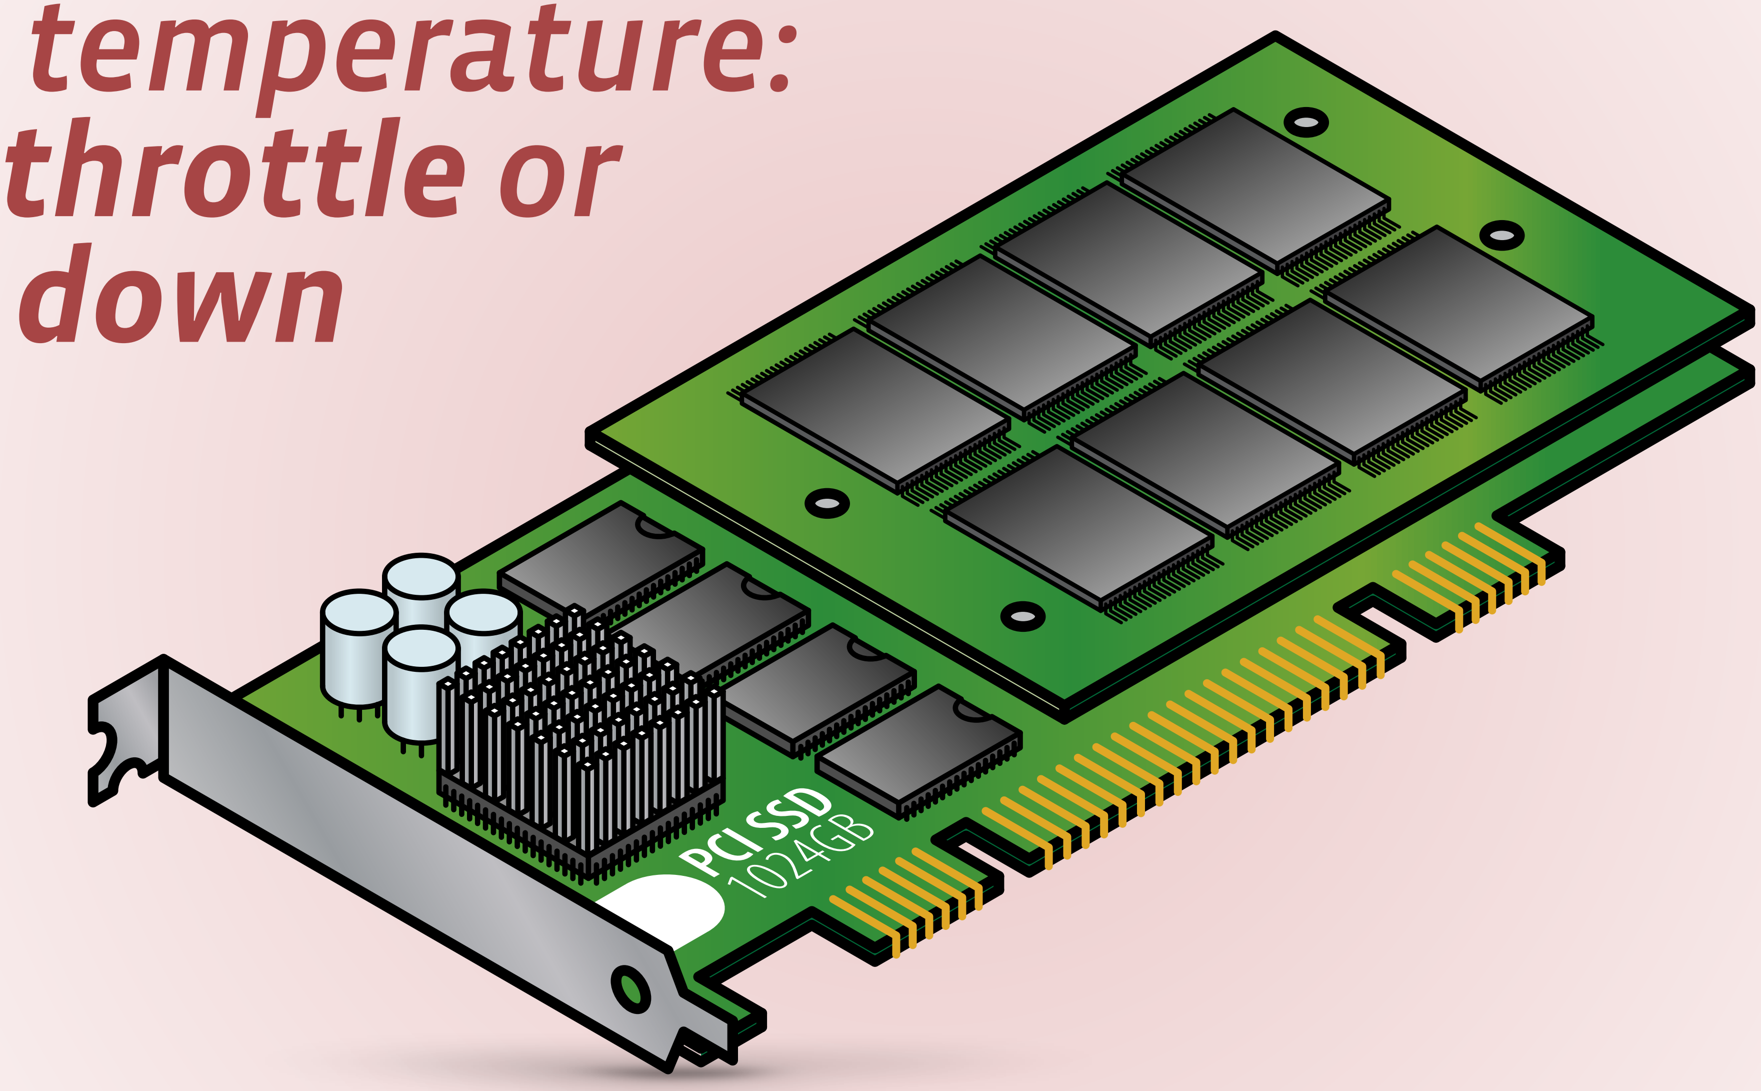

1.2TB, 1 SSD

3.2TB, 1 SSD

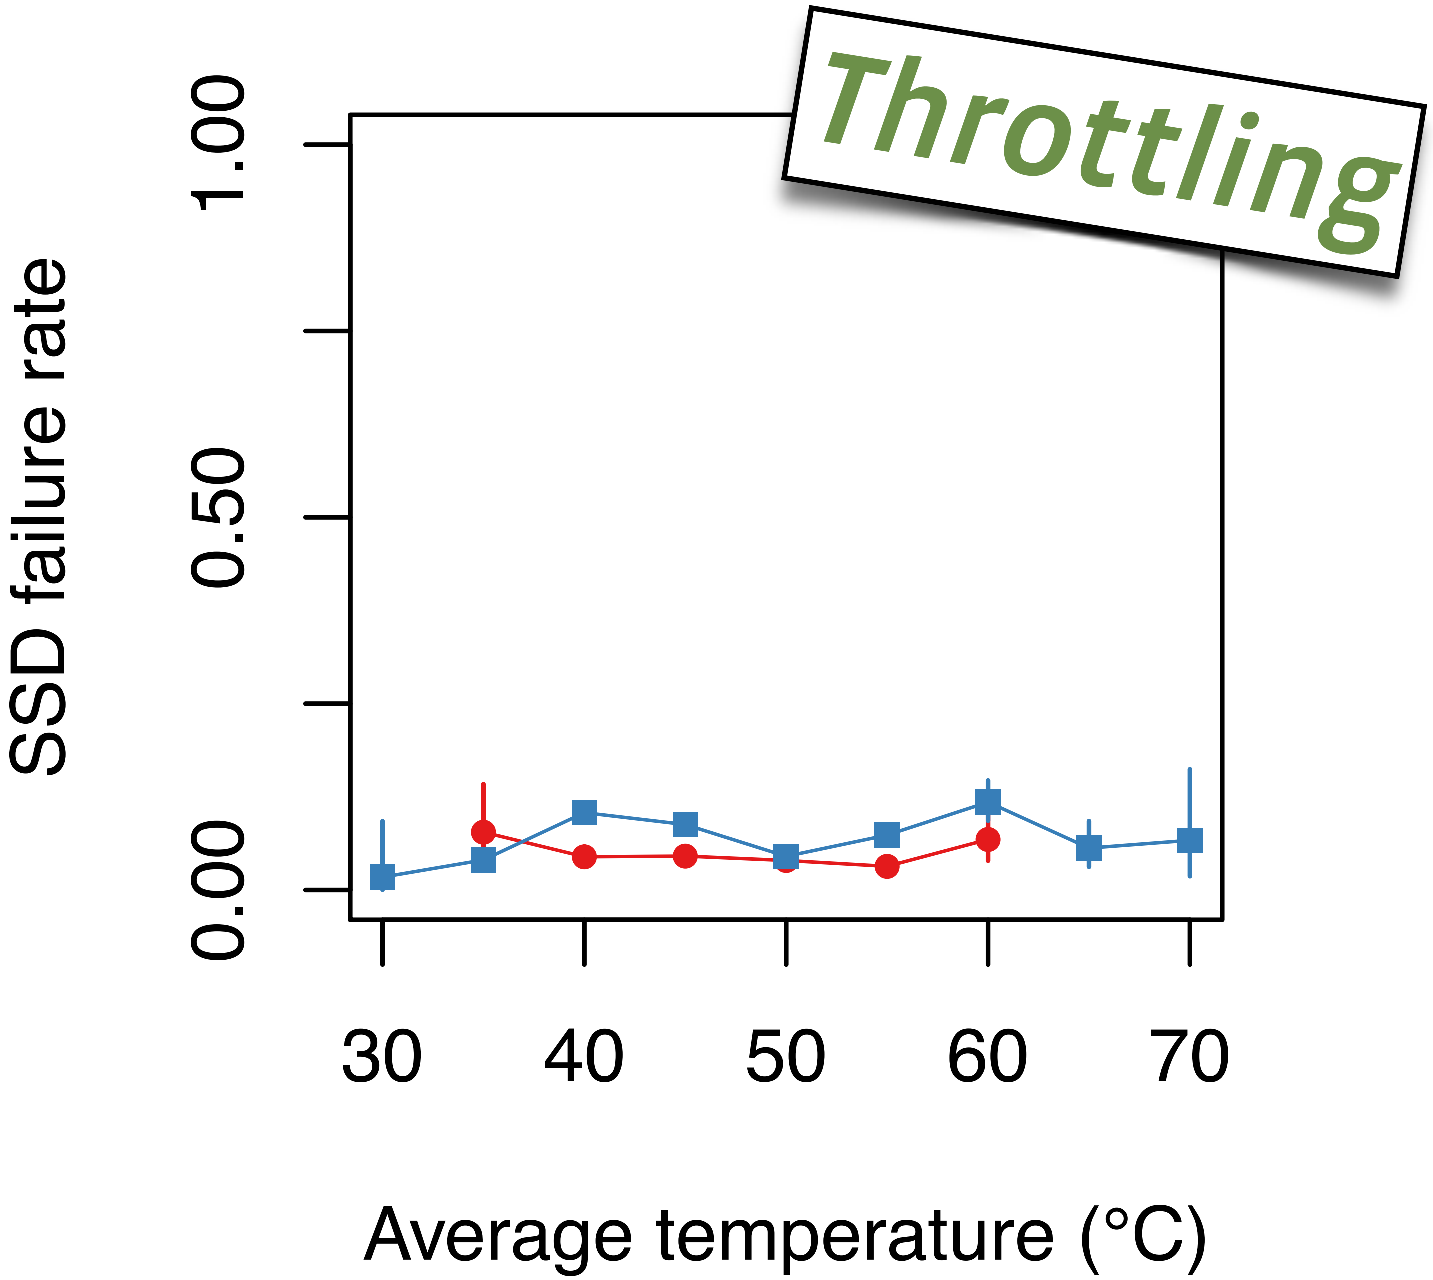

Throttling

Heavy Throttling on C & E SSDs

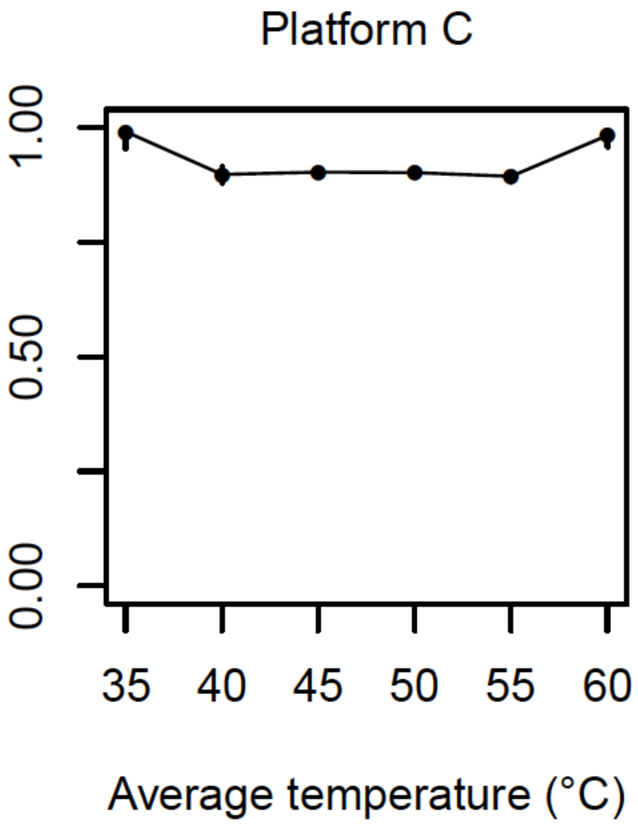

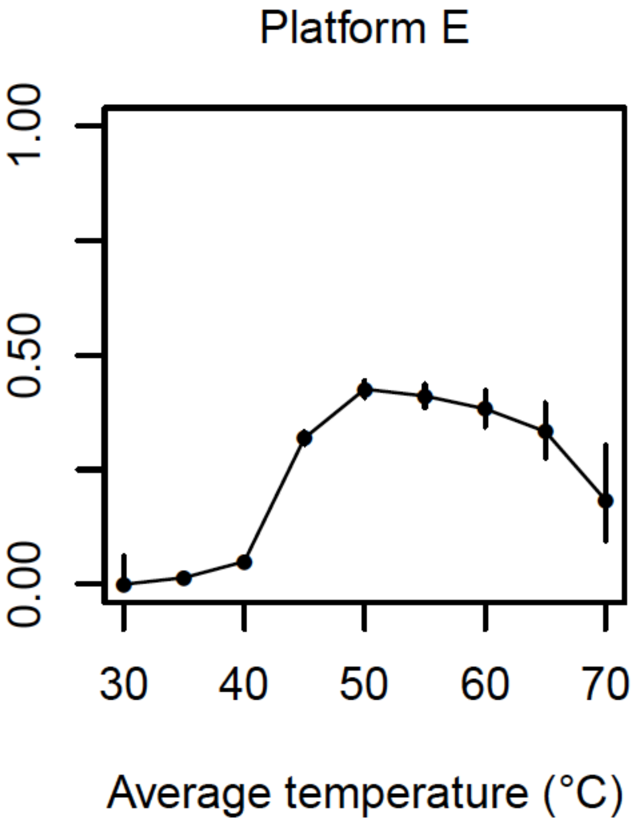

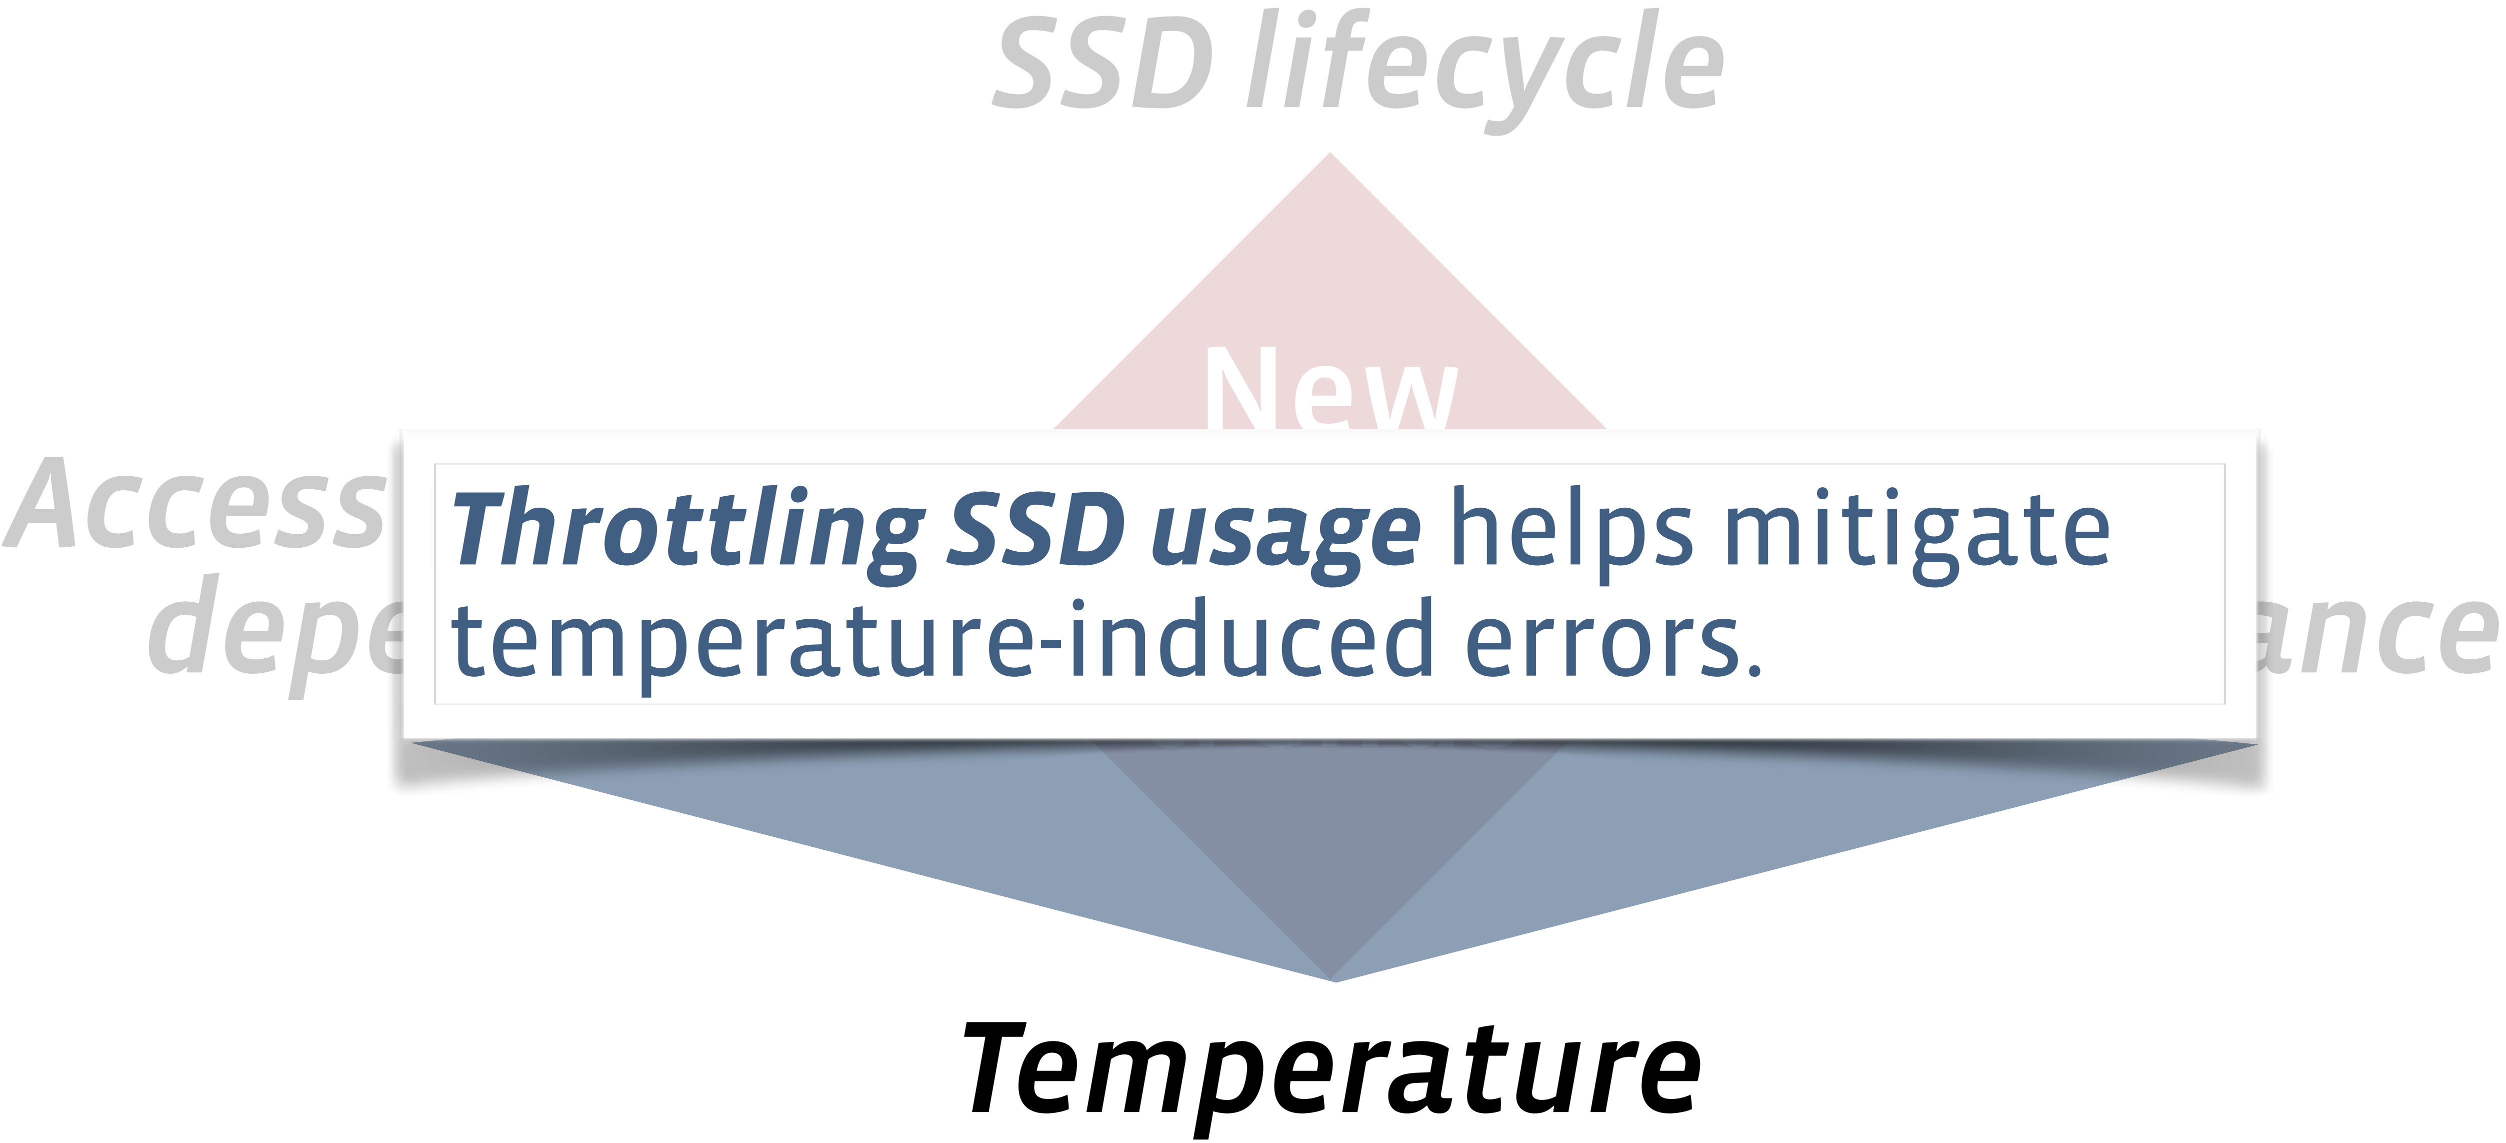

Throttling SSD usage helps mitigate temperature-induced errors.

PCIe Bus Power Consumption

- Trends for Bus Power Consumption vs. Failure Rate
  - Similar to Temperature vs. Failure Rate

- Temperature might be correlated with Bus Power

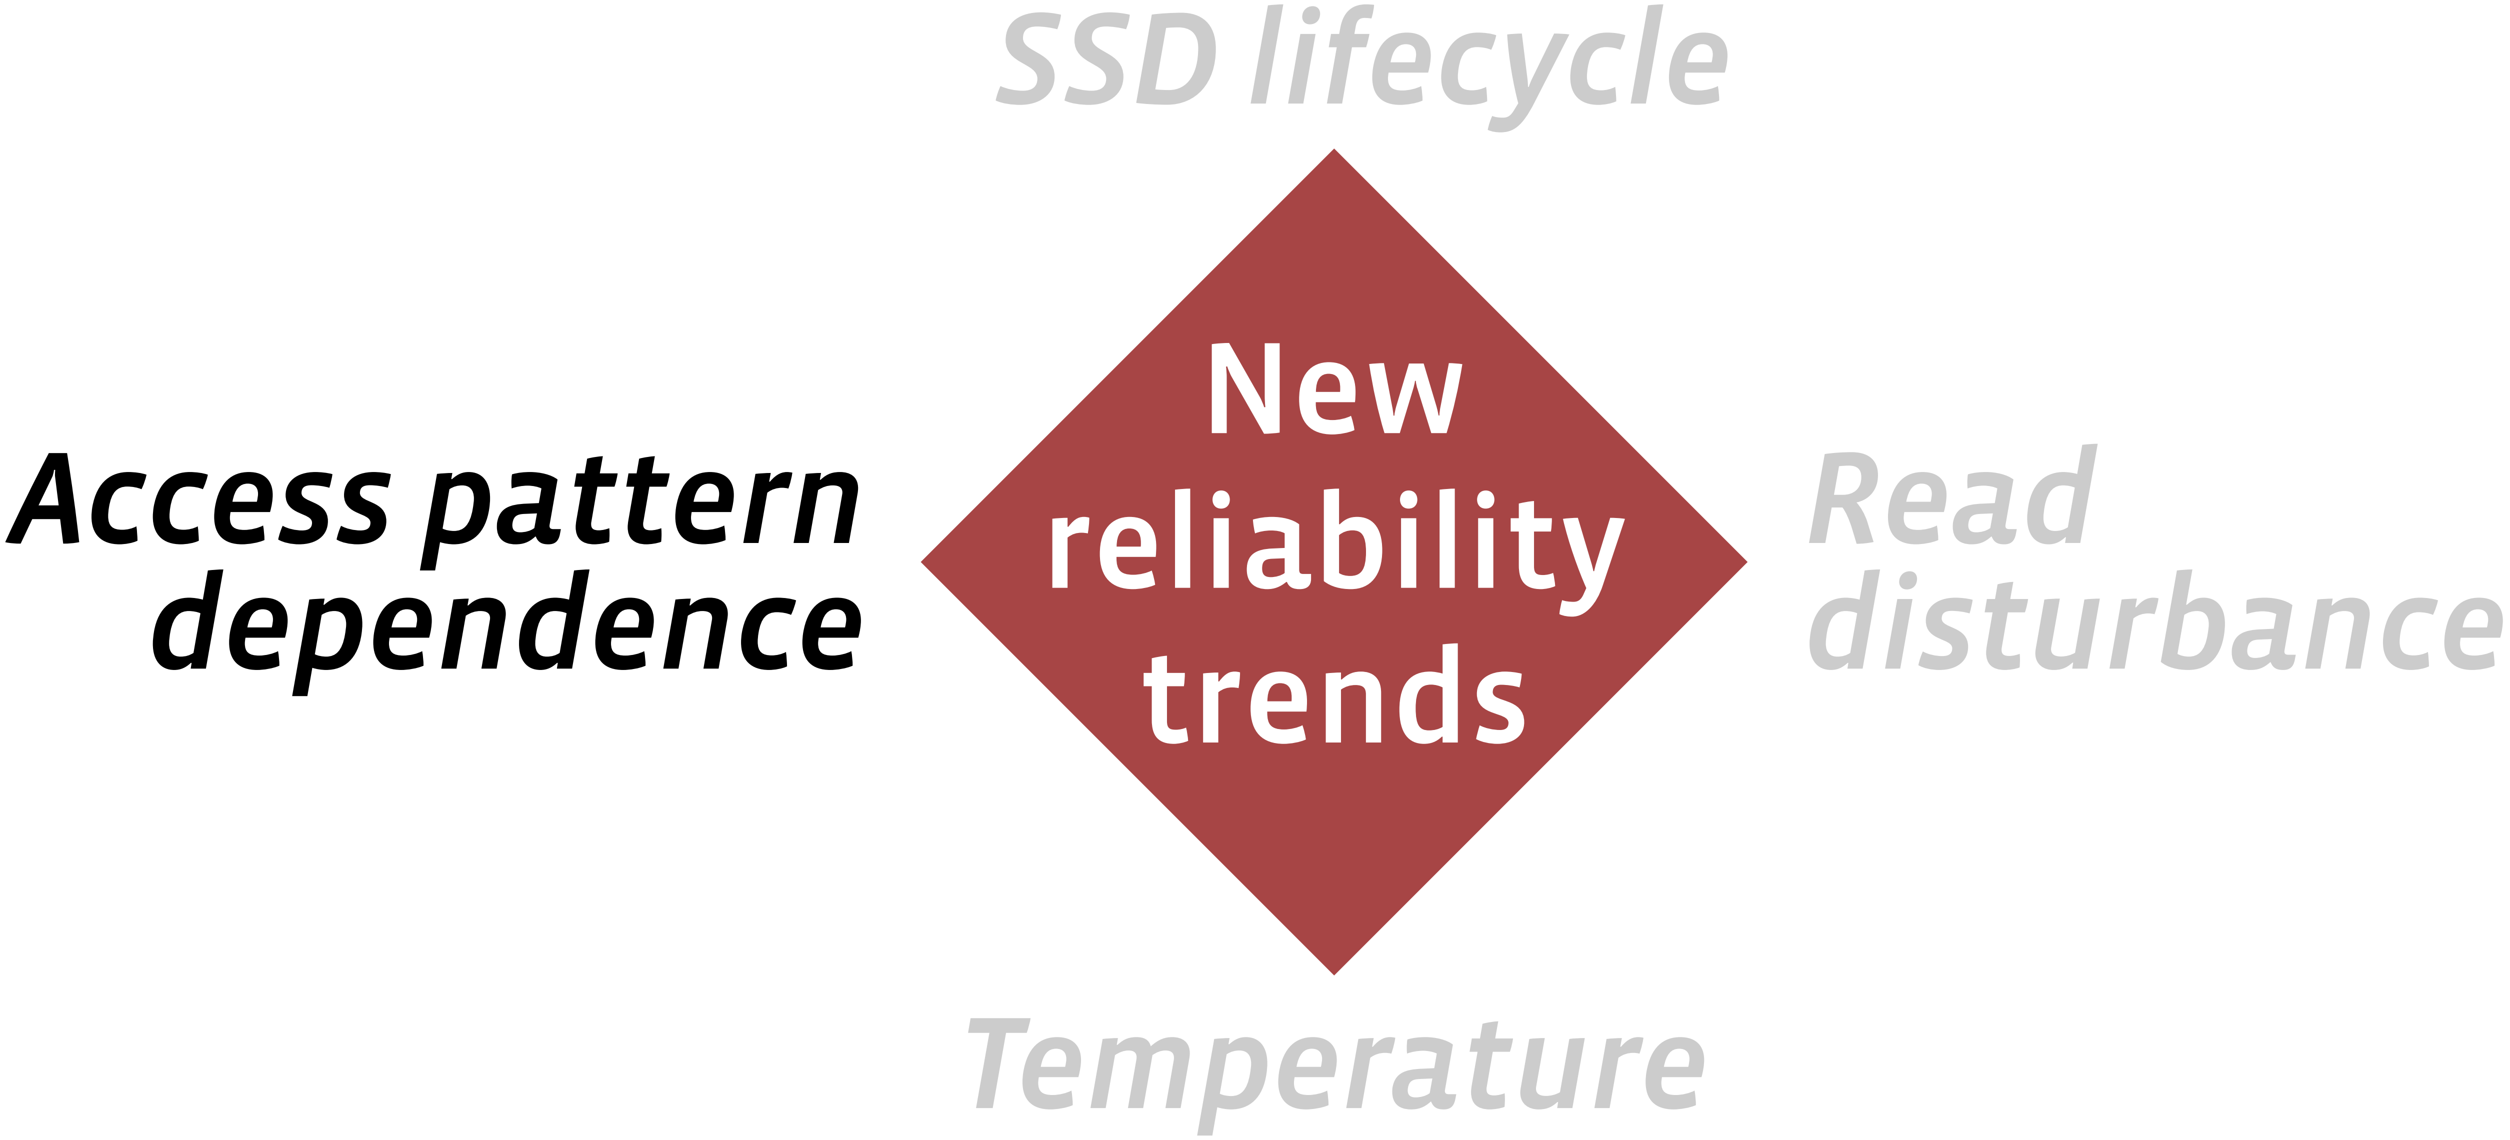

SSD lifecycle

Access pattern dependence

New reliability trends

Read disturbance

Temperature

Access pattern effects

System buffering

- data served from OS caches
- decreases SSD usage

Write amplification

- updates to small amounts of data
- increases erasing and copying

Access pattern effects

System buffering

- data served from OS caches
- decreases SSD usage

Write amplification

- updates to small amounts of data
- increases erasing and copying

PCI SSD
1024GB

OS
PCI SSD
1024GB

Page cache

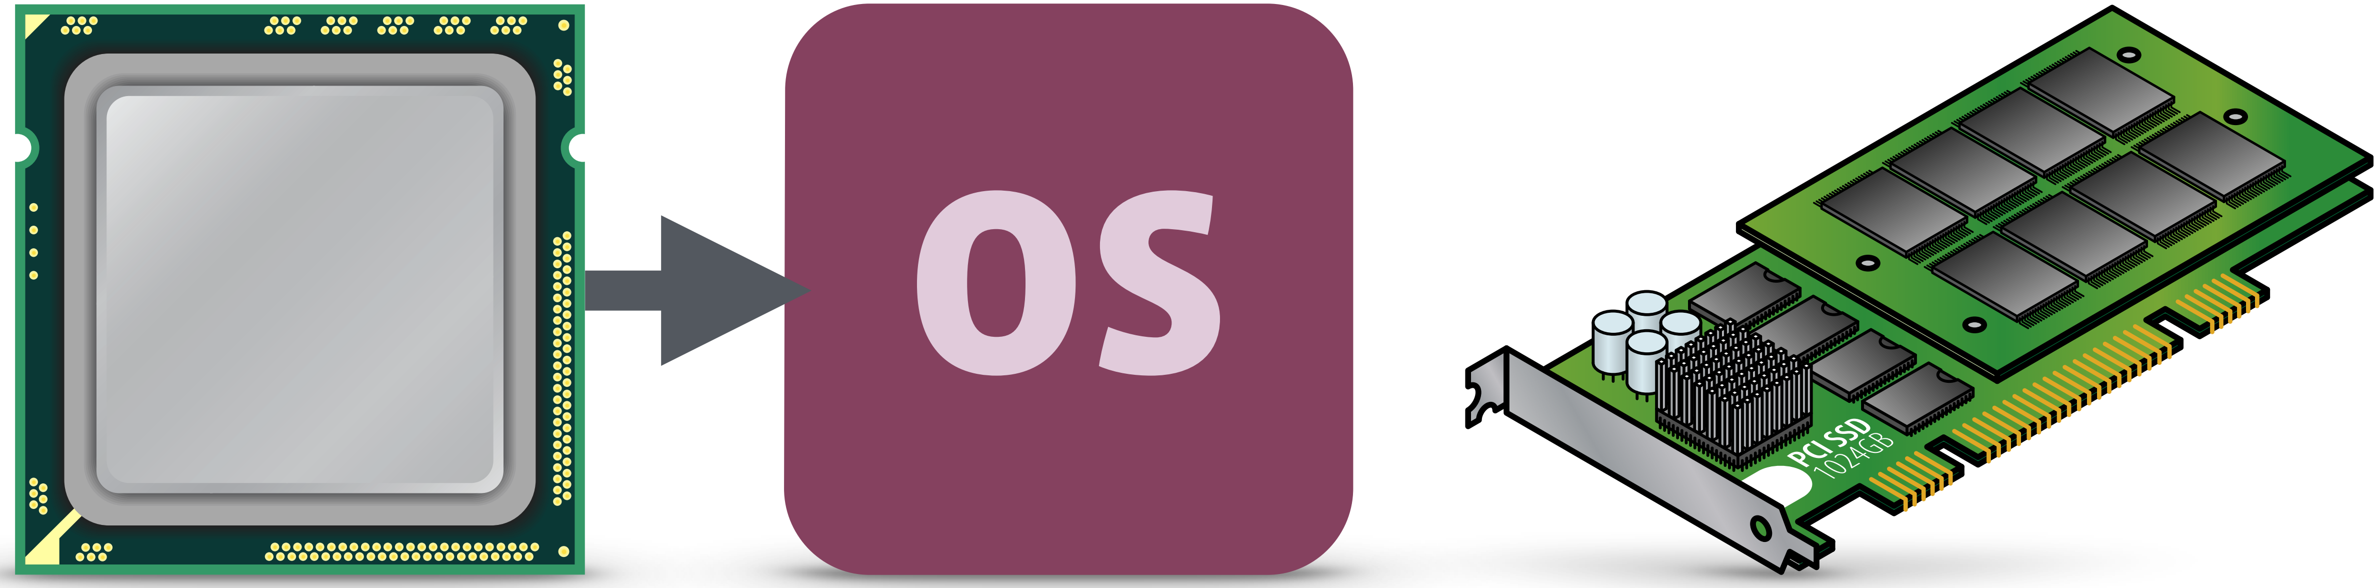

Page cache

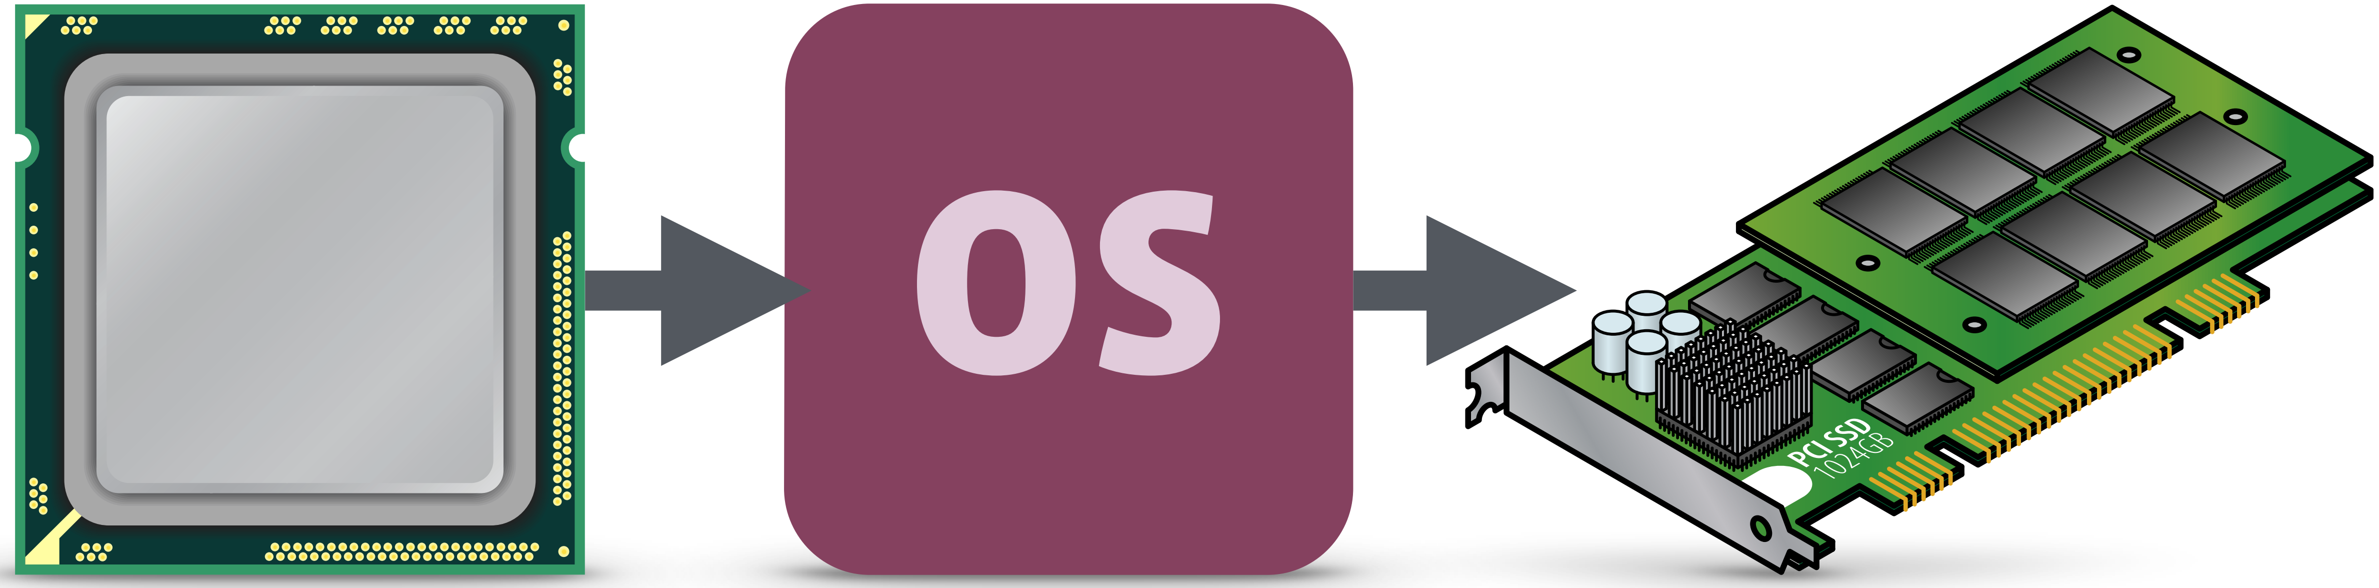

Page cache

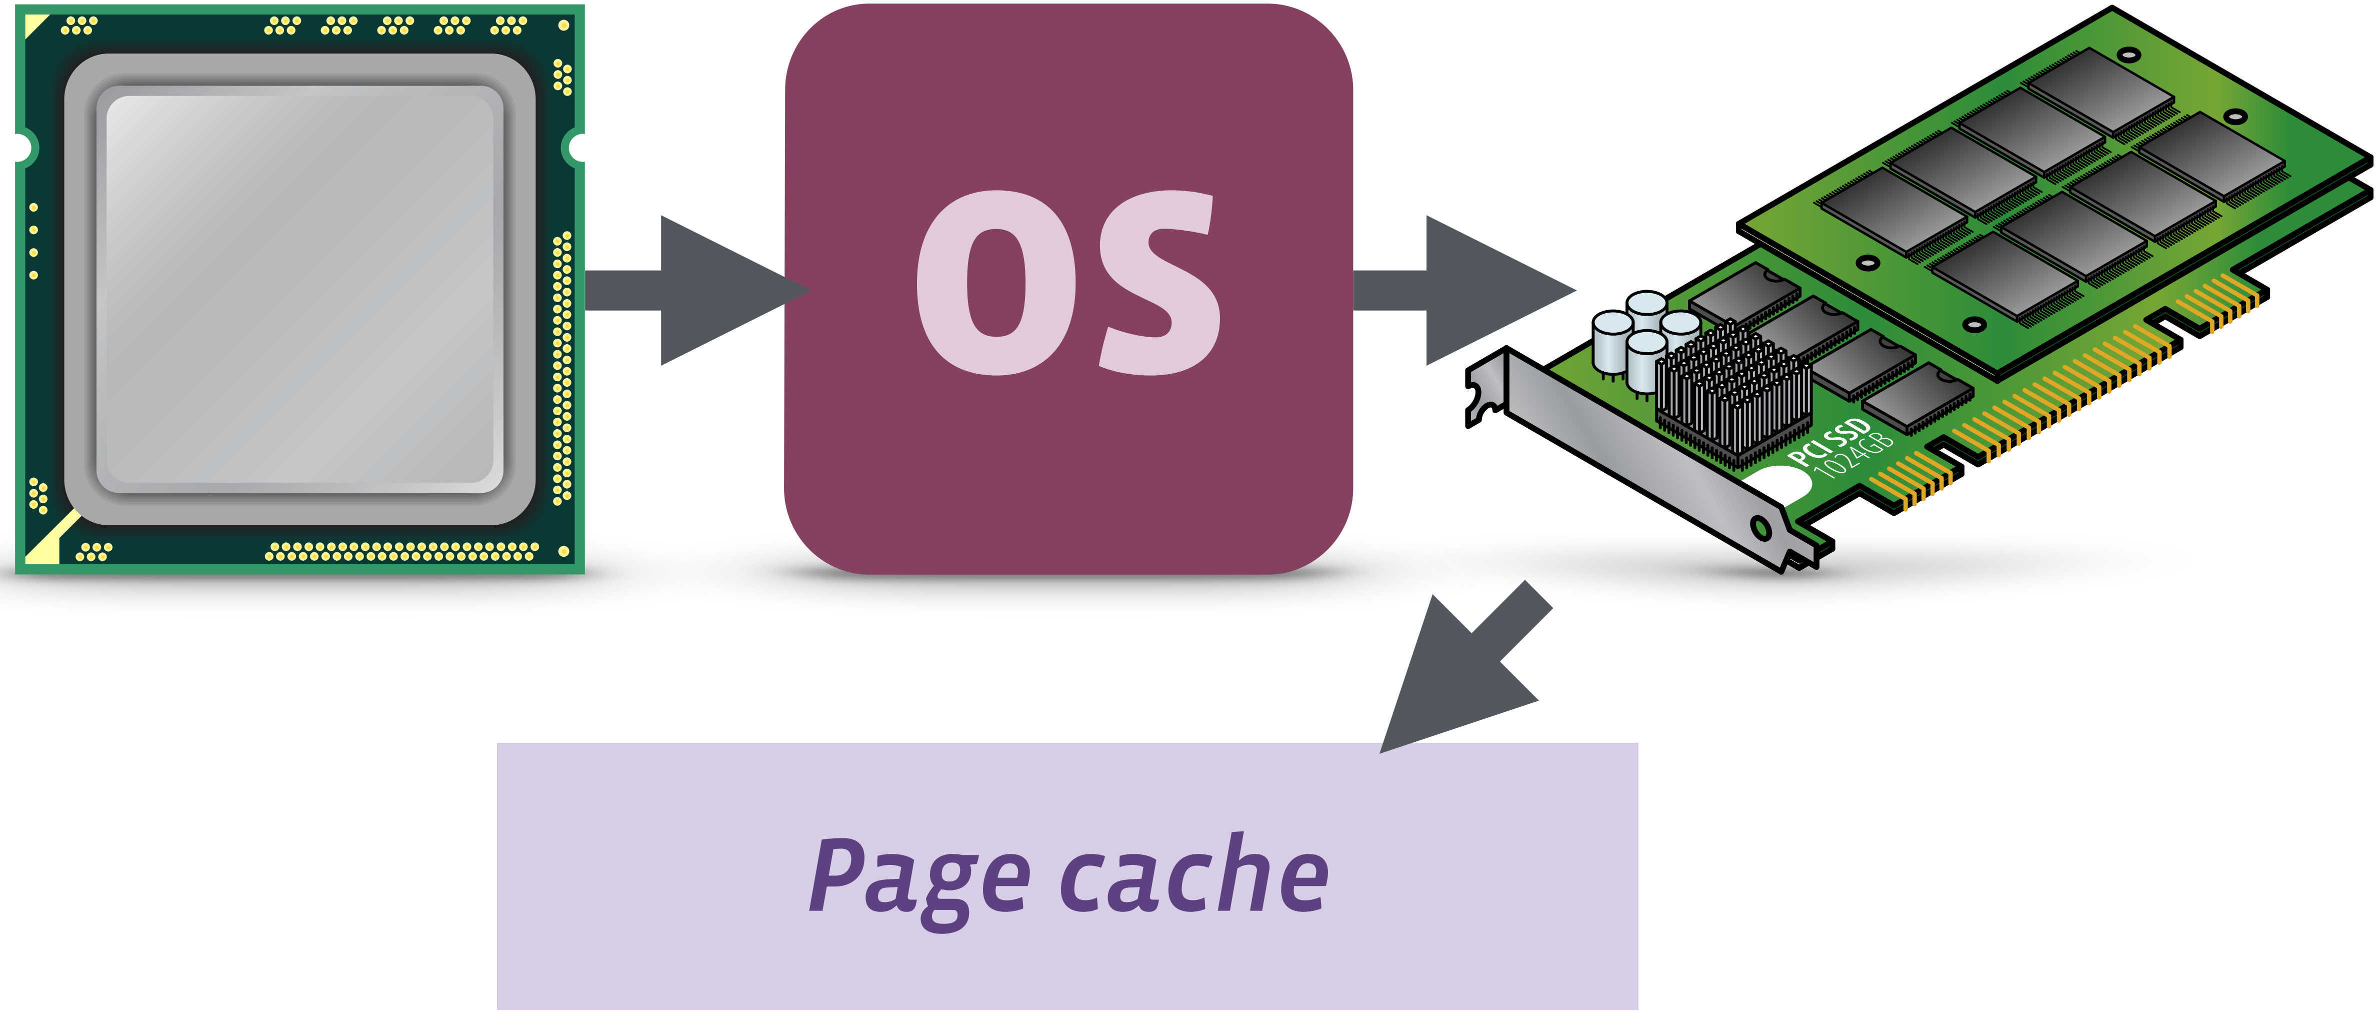

OS

Page cache

Page cache

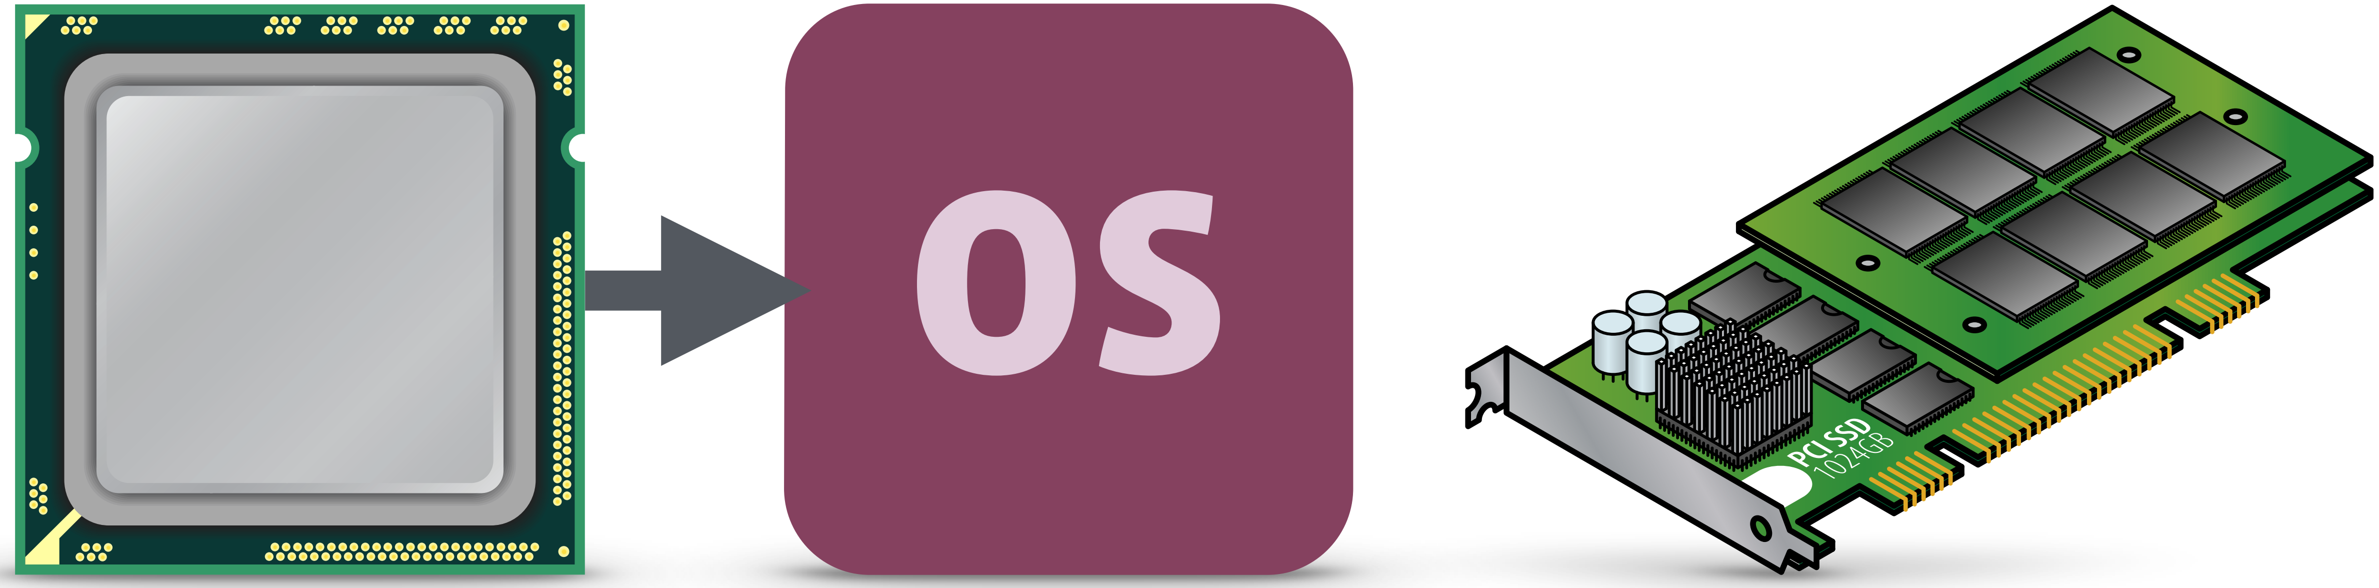

Page cache

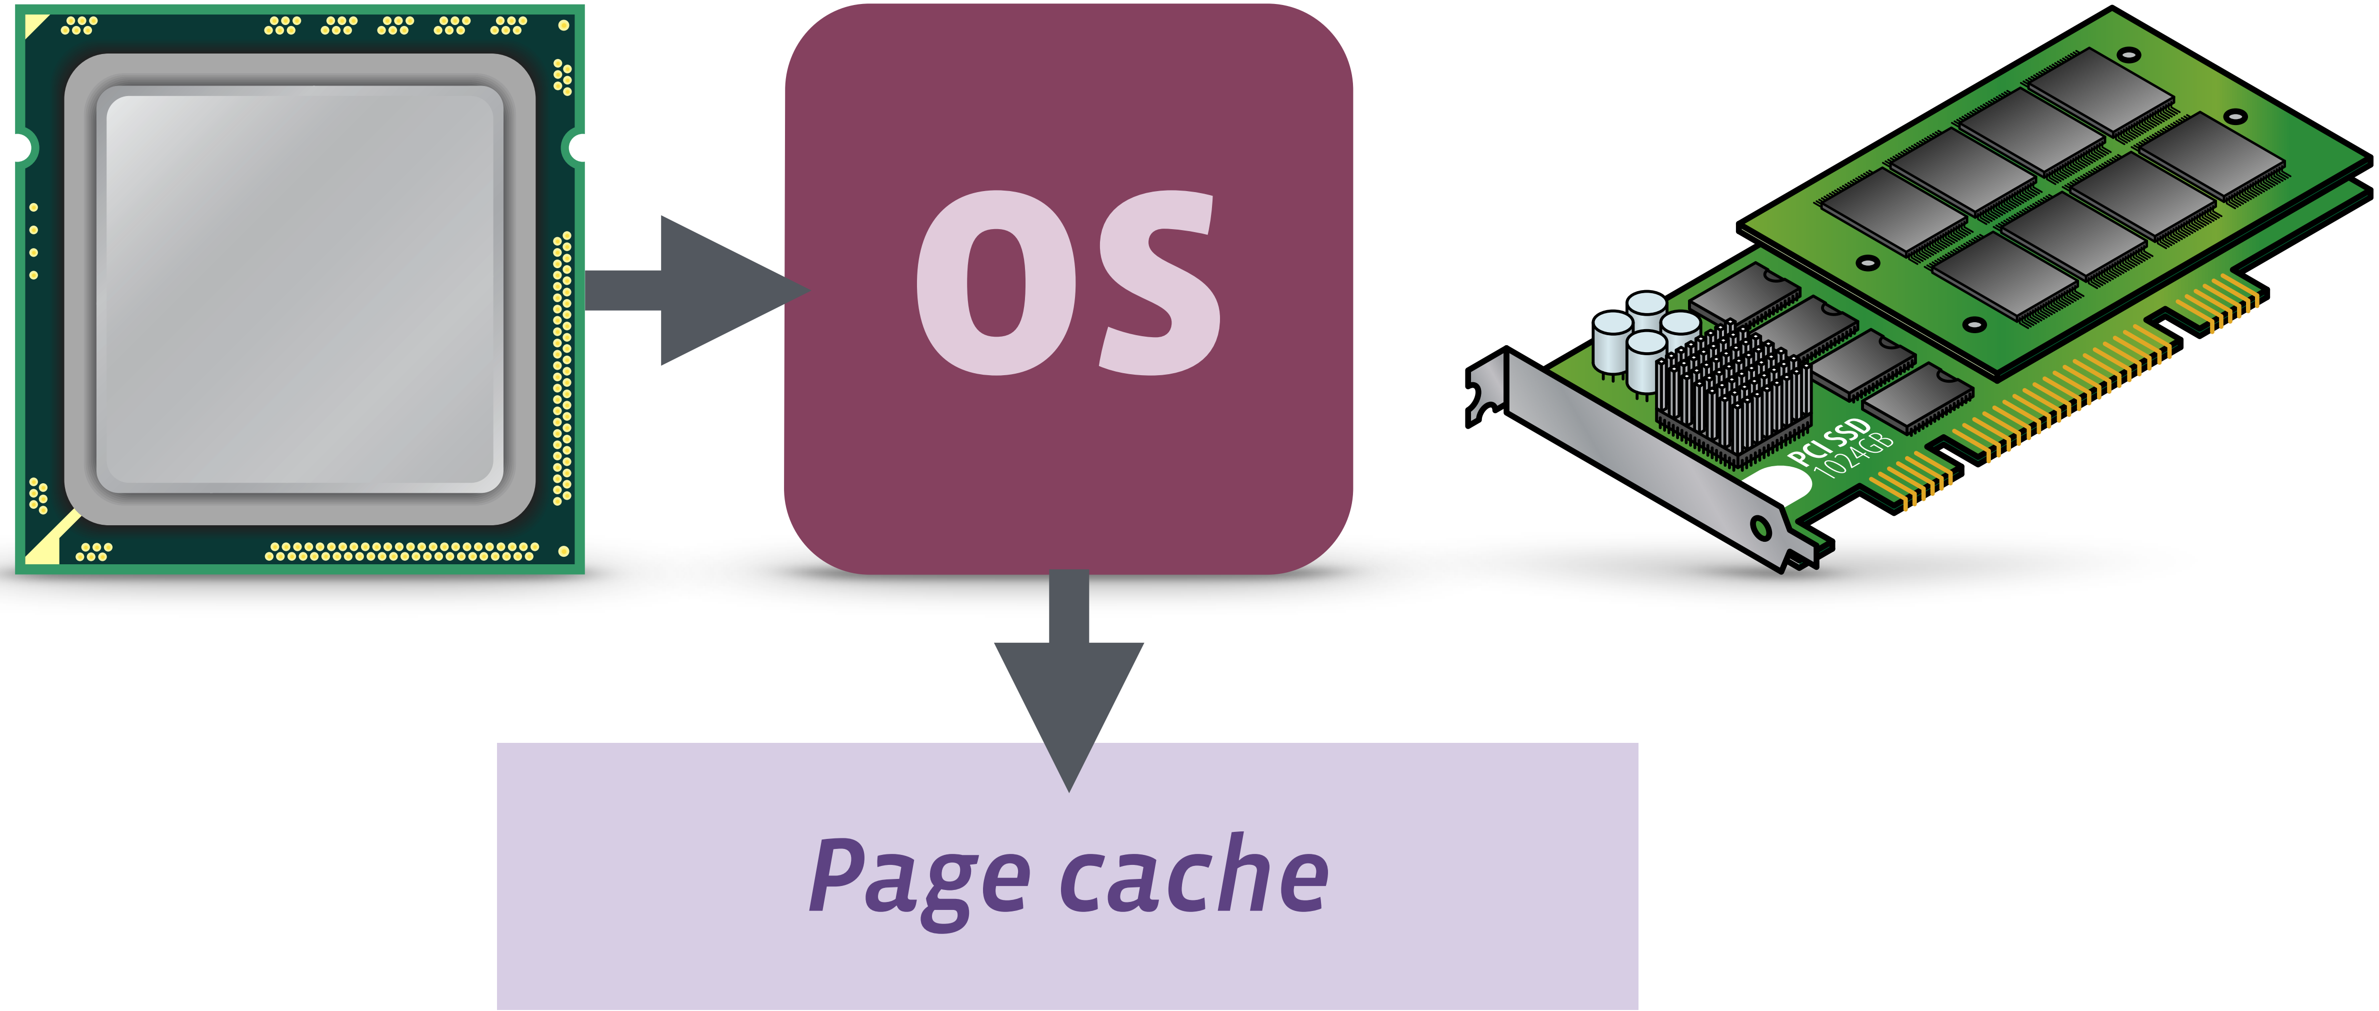

OS

Page cache

System caching reduces the impact of SSD writes

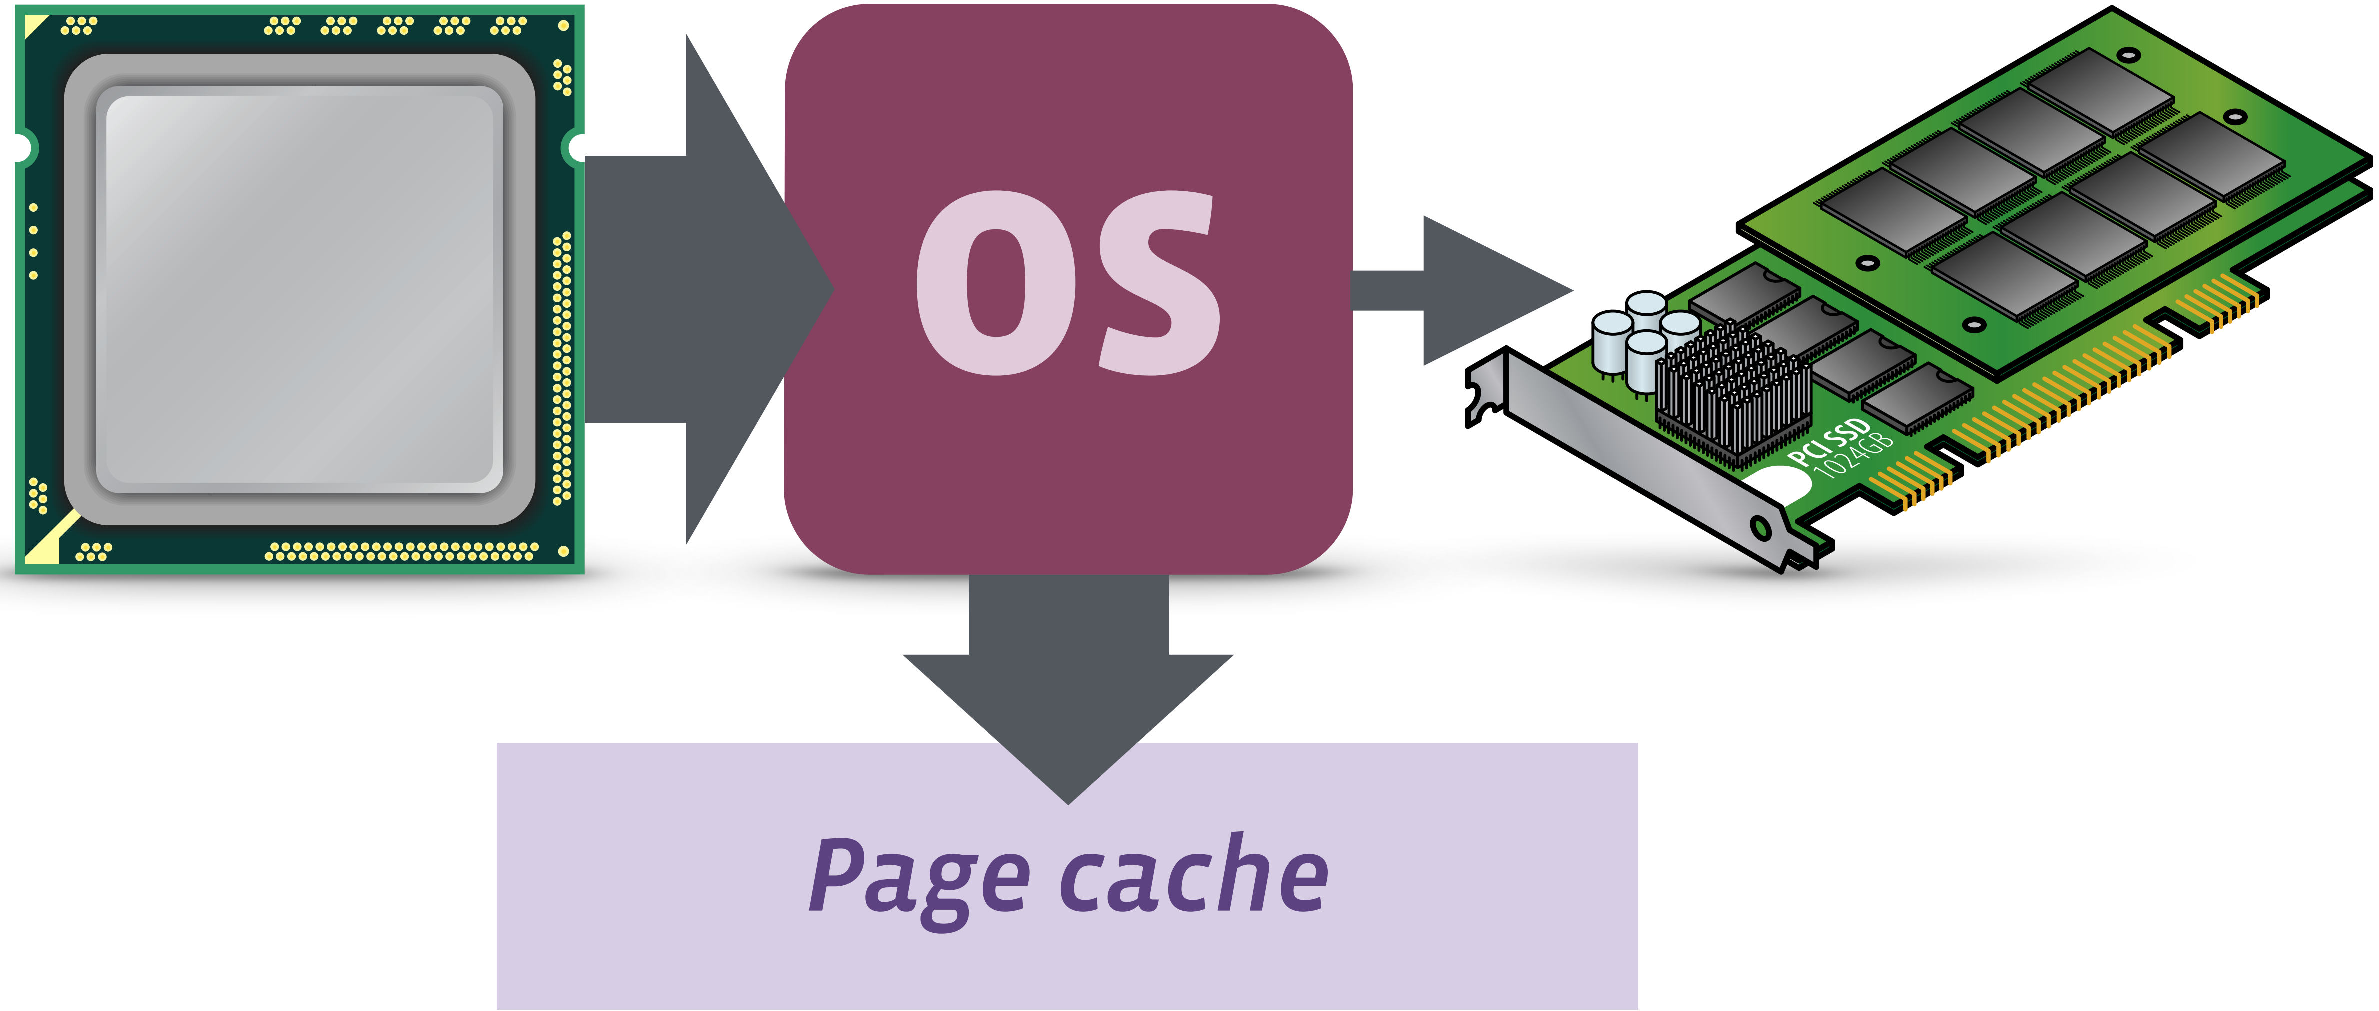

OS

Page cache

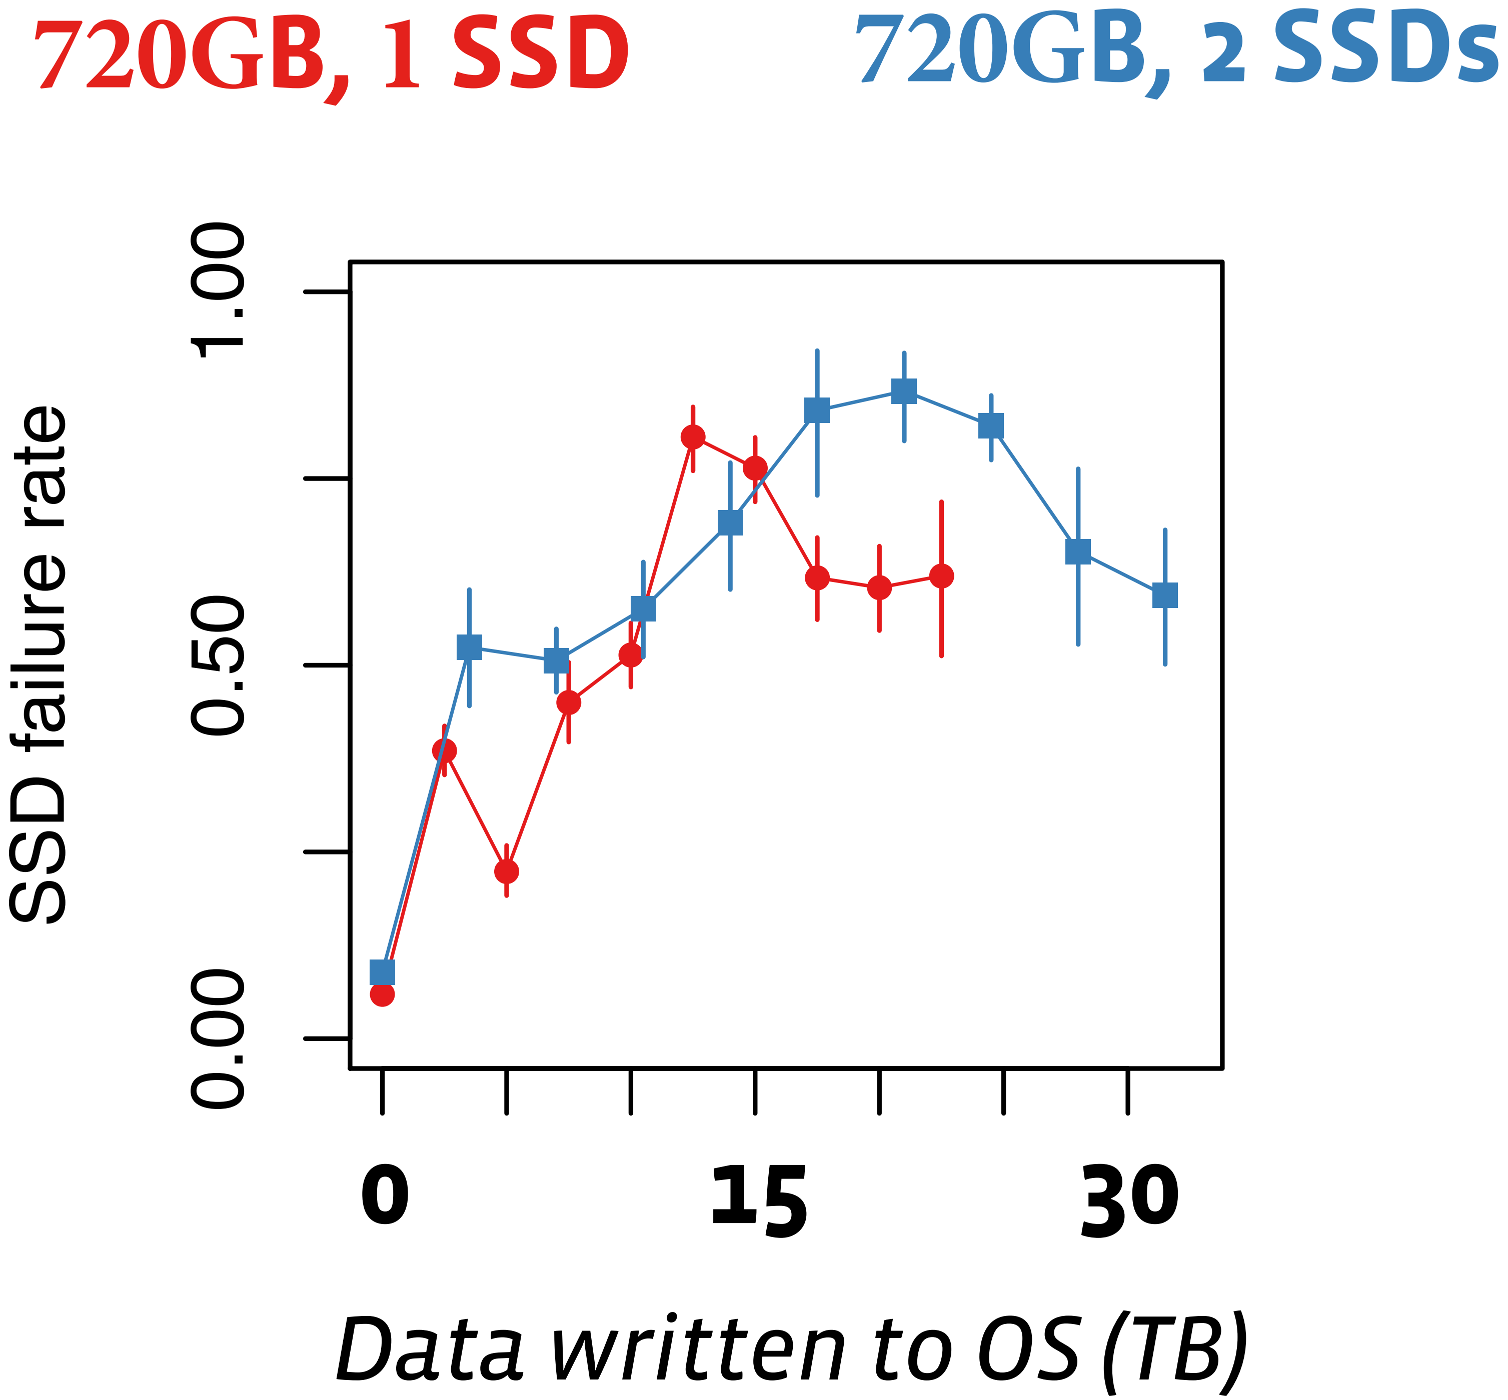
720GB, 1 SSD
720GB, 2 SSDs
SSD failure rate
1.00
0.50
0.00
0
15
30
Data written to OS (TB)

720GB, 2 SSDs

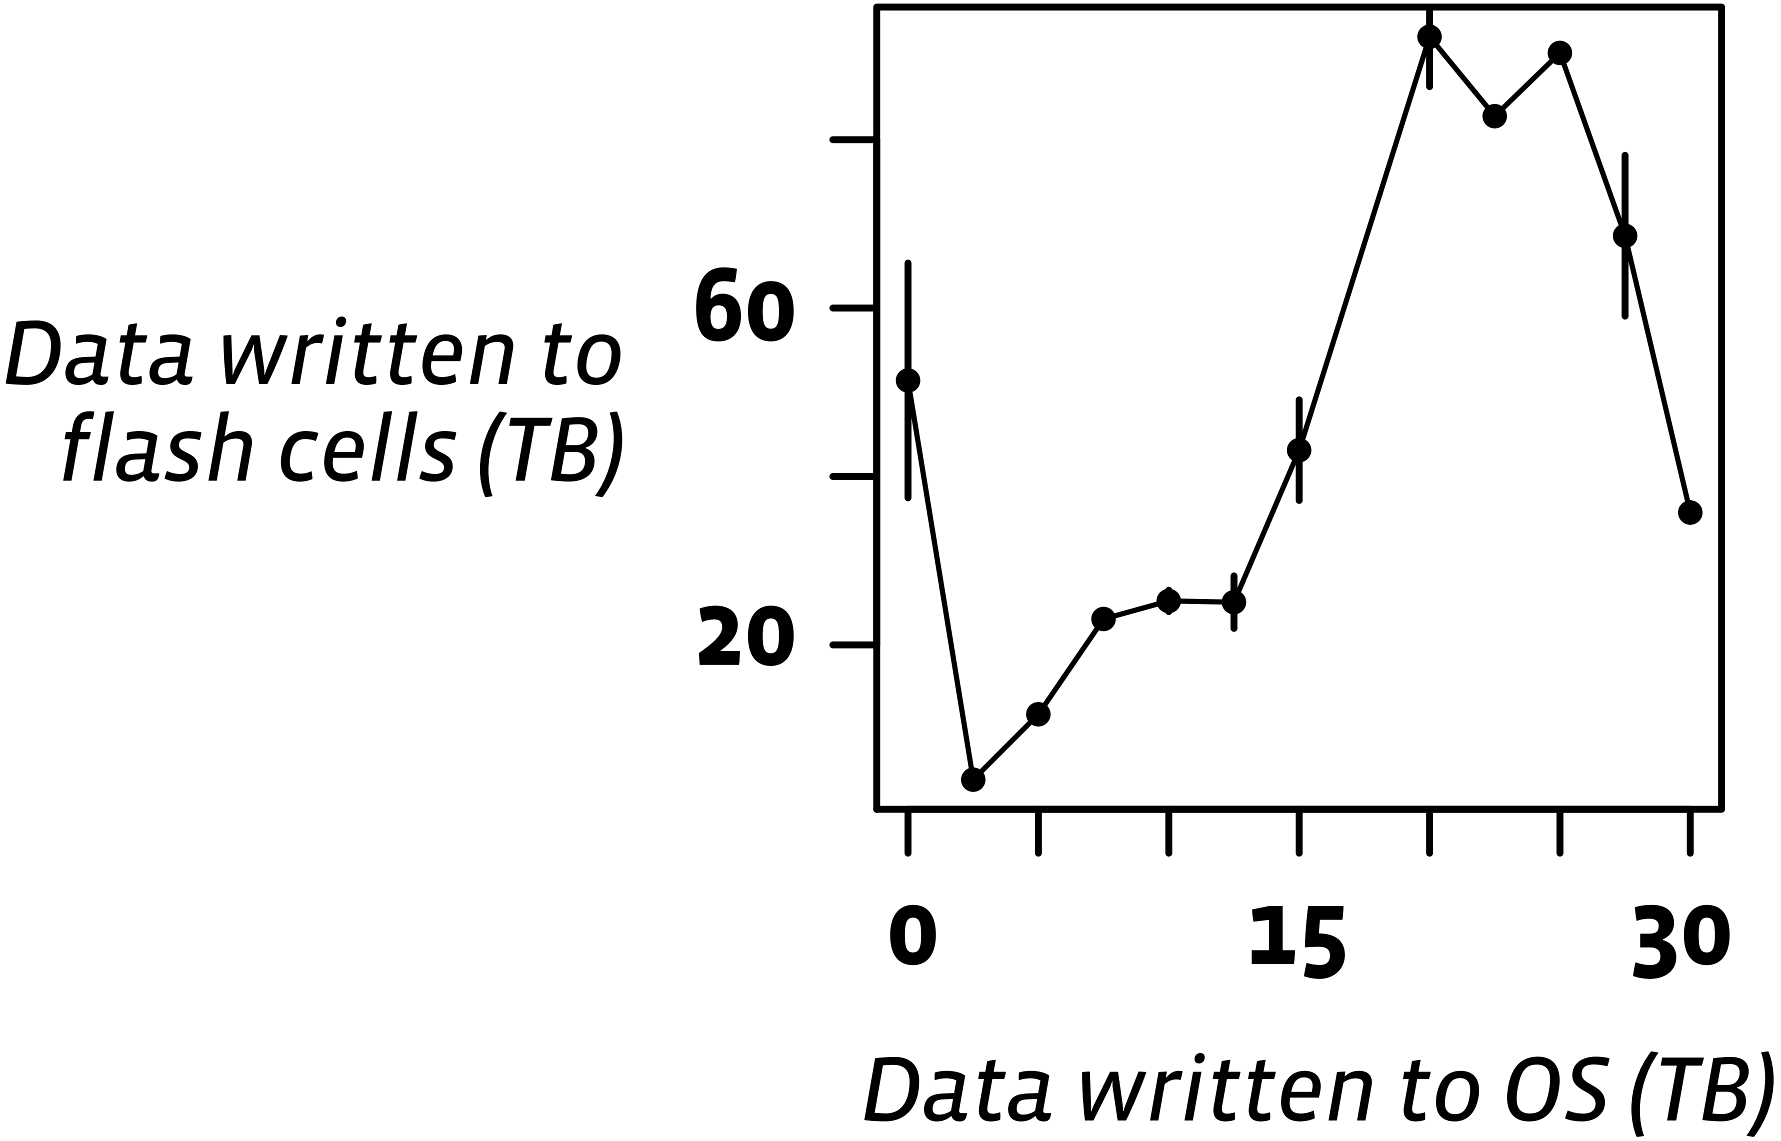
Data written to flash cells (TB)
60
20
0
15
30
Data written to OS (TB)

System-Level Writes vs. Chip-Level Writes

- More data written at the software

  does not imply

- More data written into flash chips

- Due to system level buffering

- More system-level writes can enable more opportunities for coalescing in the system buffers

Access pattern effects

System buffering

- data served from OS caches
- decreases SSD usage

Write amplification

- updates to small amounts of data
- increases erasing and copying

Flash devices use a **translation layer** to locate data

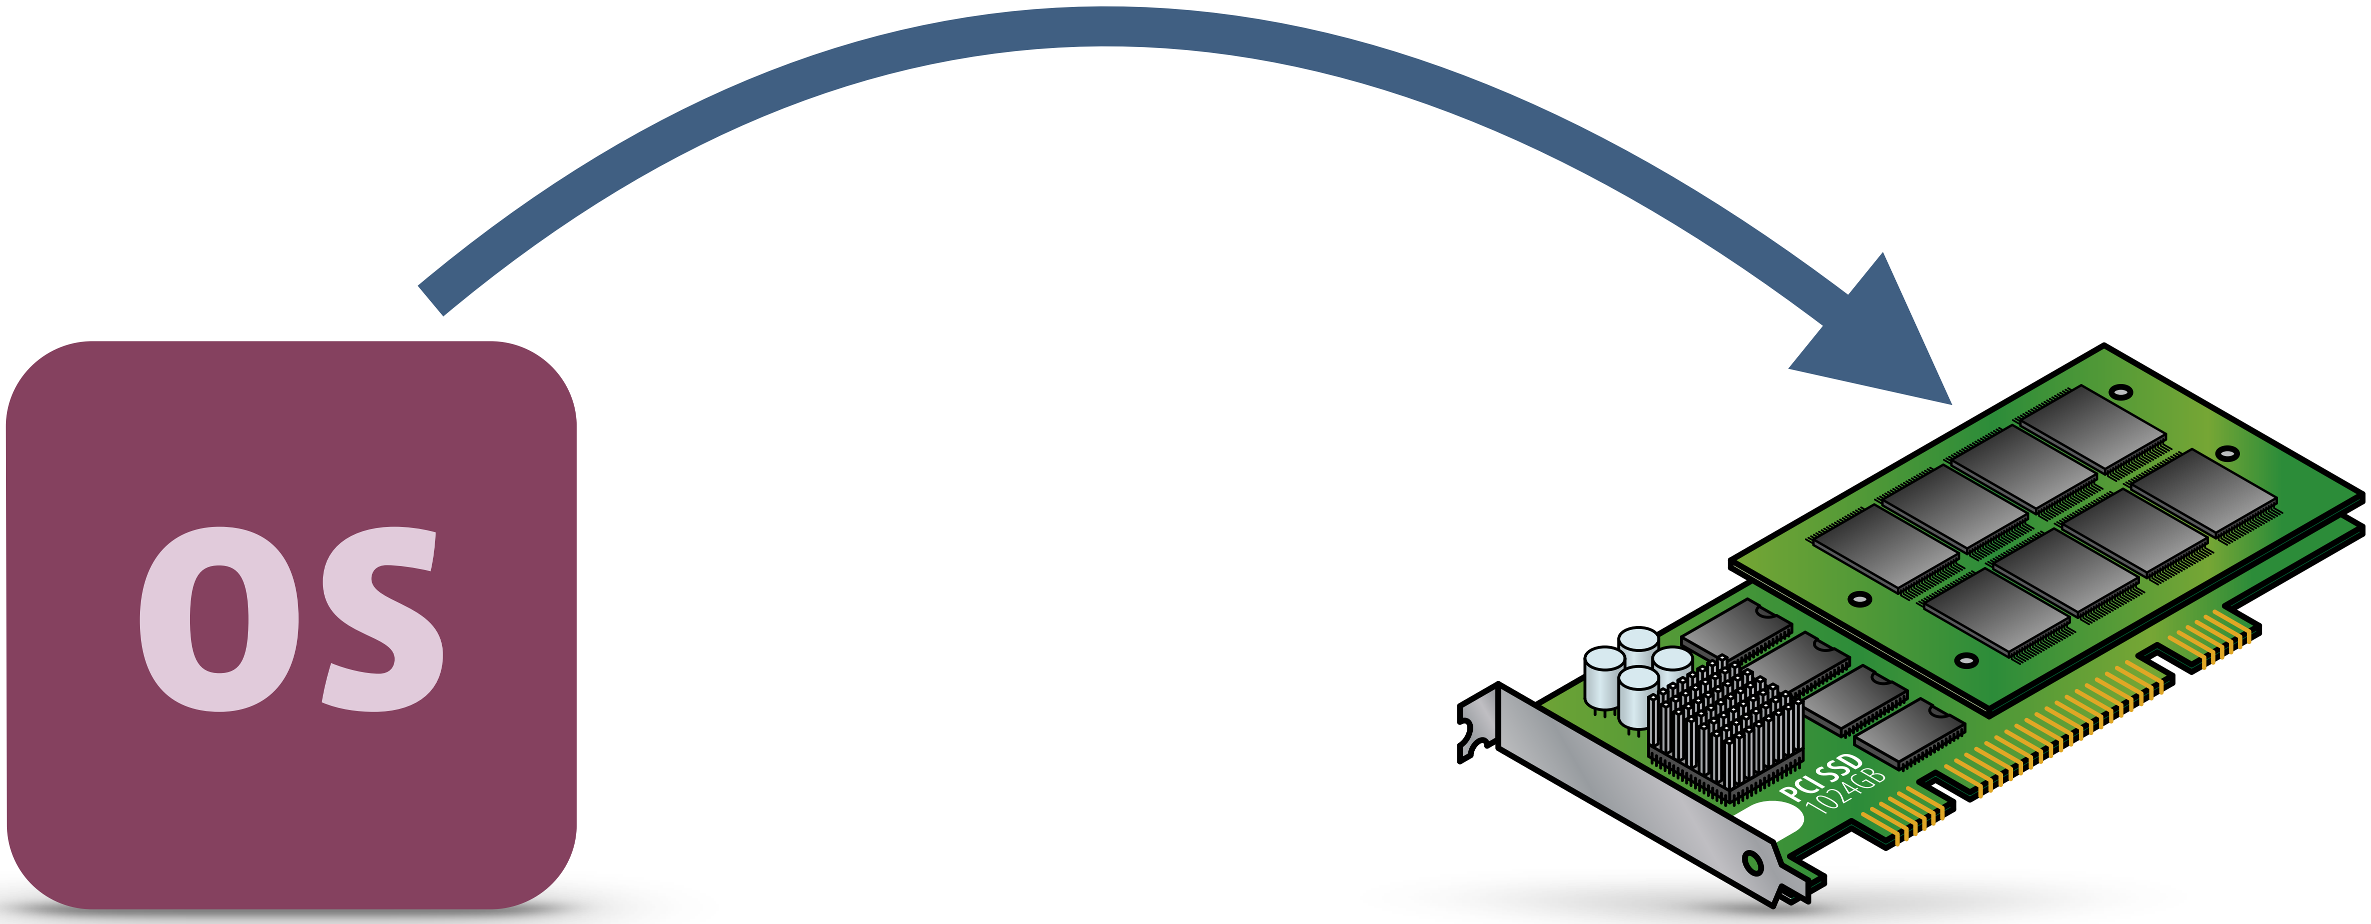

Translation layer

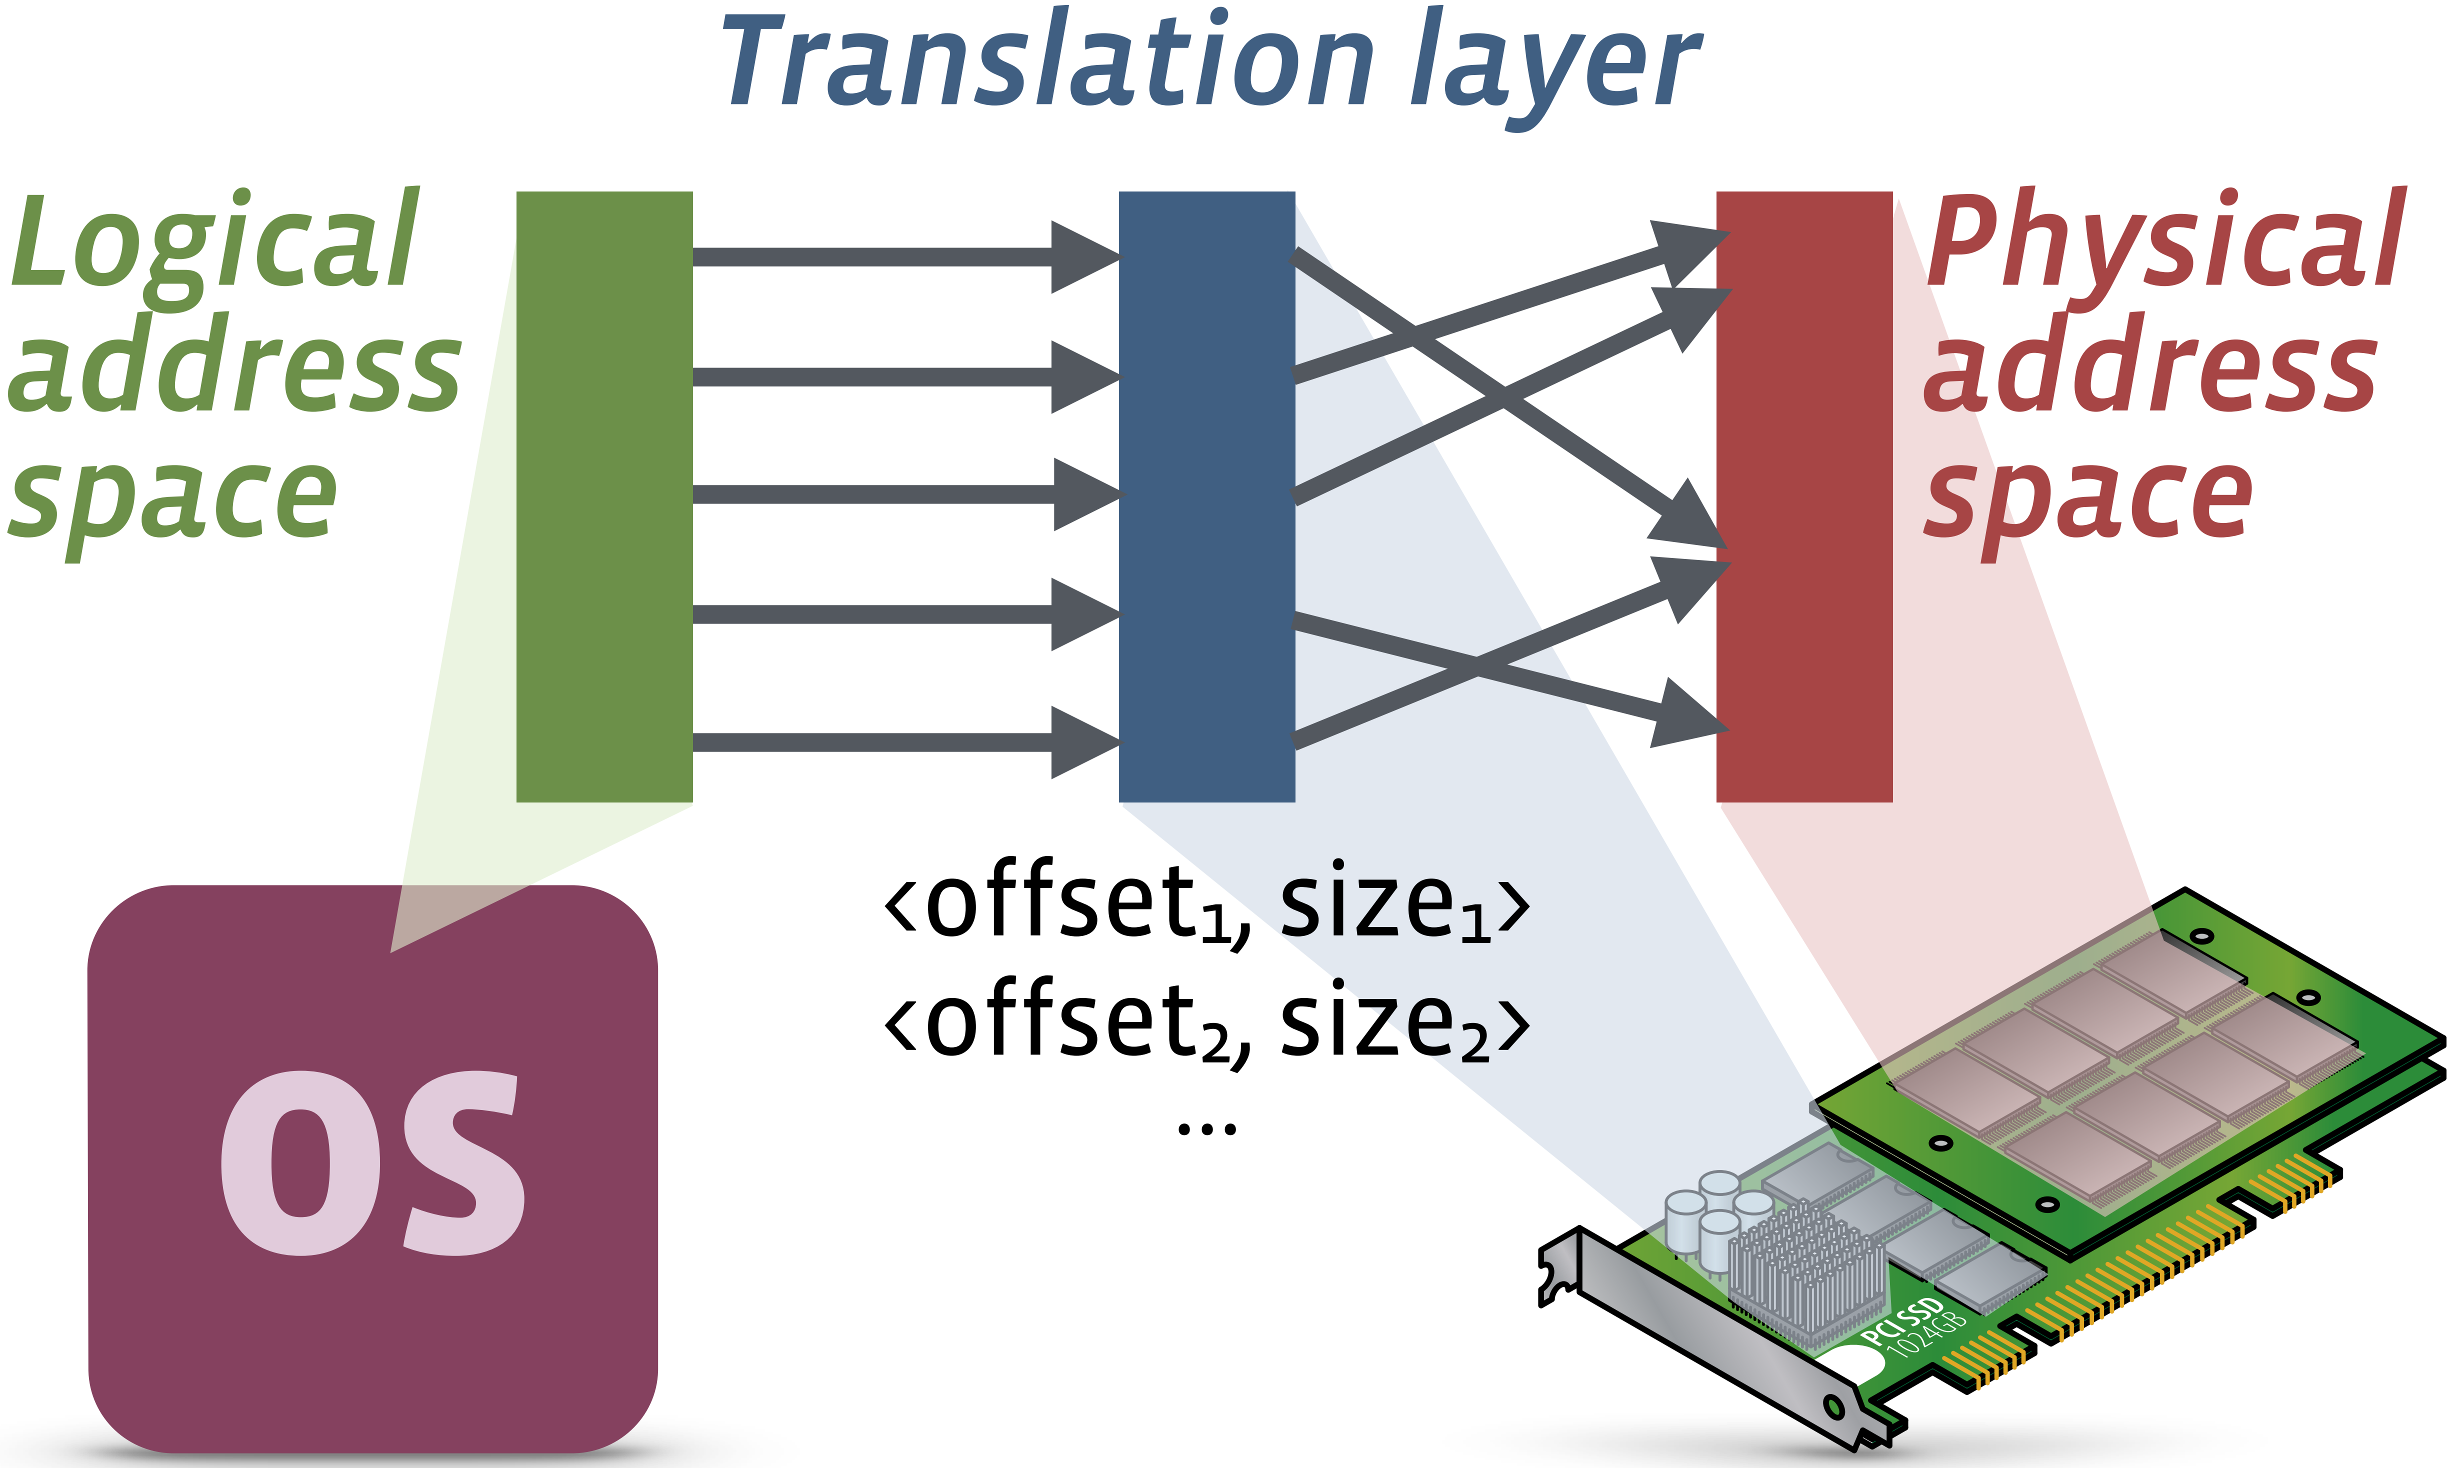

Logical address space

Physical address space

OS

$\langle \text{offset}_1, \text{size}_1 \rangle$

$\langle \text{offset}_2, \text{size}_2 \rangle$

…

Sparse data layout

more ***translation metadata***

potential for ***higher*** **write amplification**

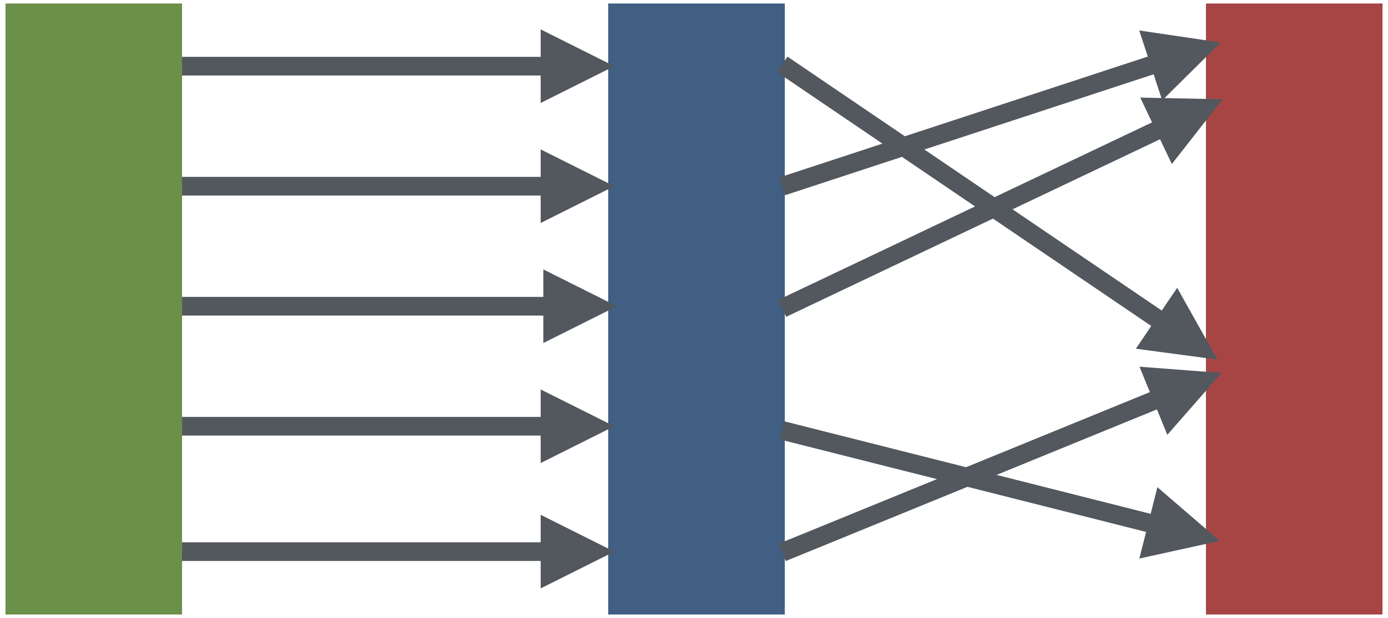

e.g., many small file updates

***Dense* data layout**

less *translation metadata*

potential for ***lower*** **write amplification**

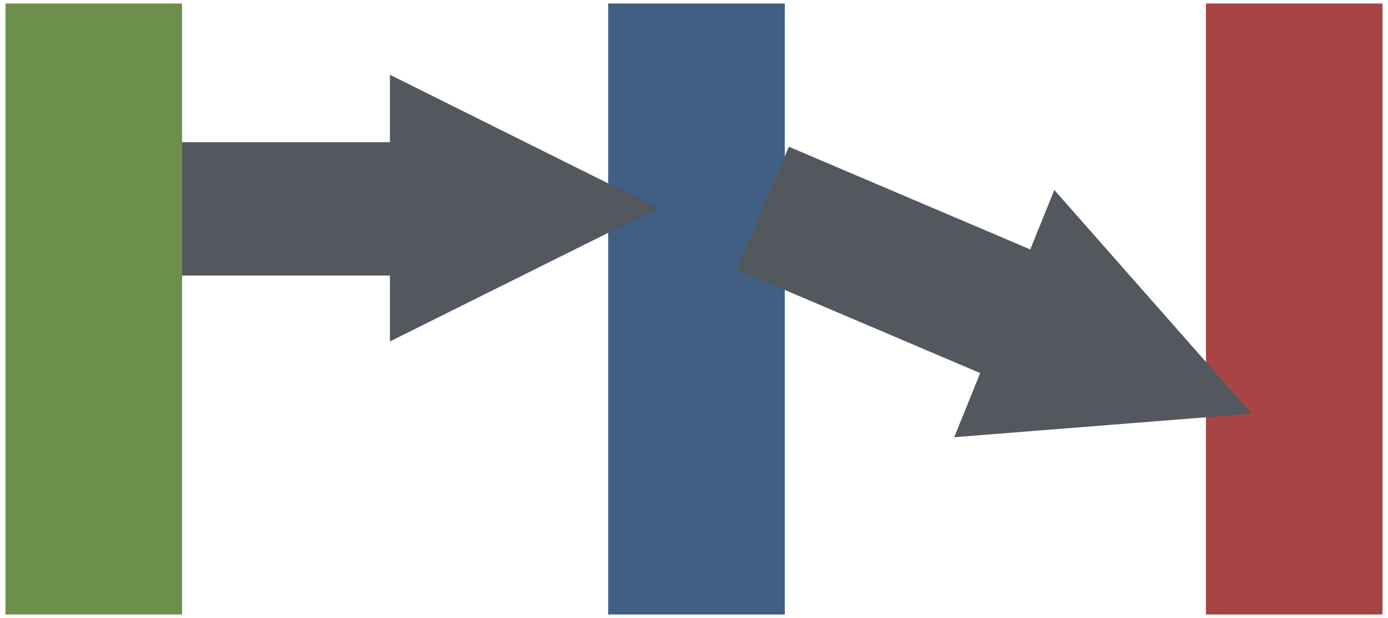

e.g., one huge file update

Use **translation data size**

to examine effects of data layout

(relates to application access patterns)

720GB, 1 SSD

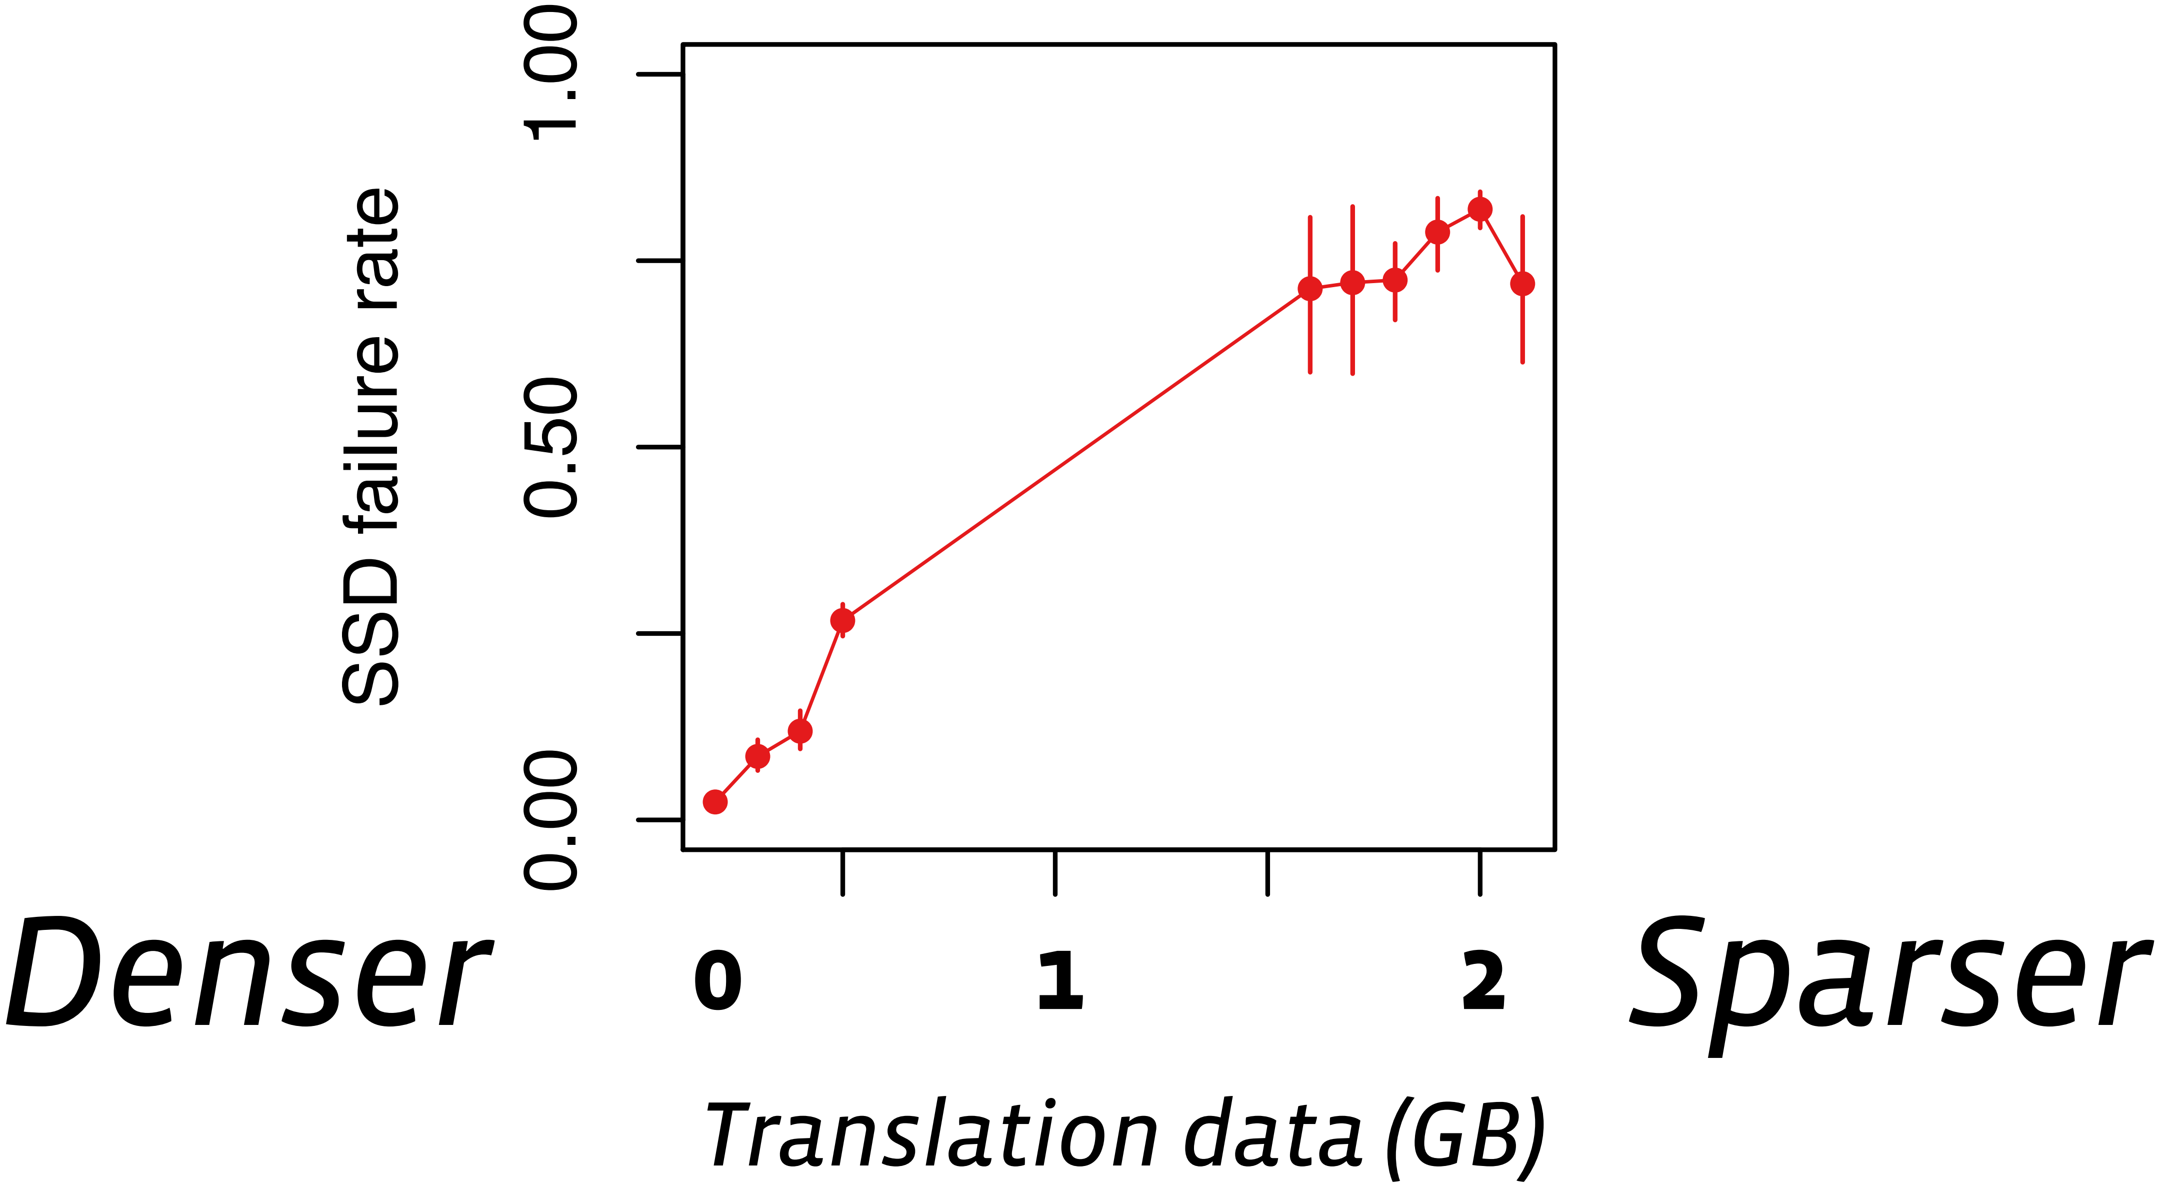
SSD failure rate
0.00
0.50
1.00
Denser
0
1
2
Sparser
Translation data (GB)

Write amplification in the field

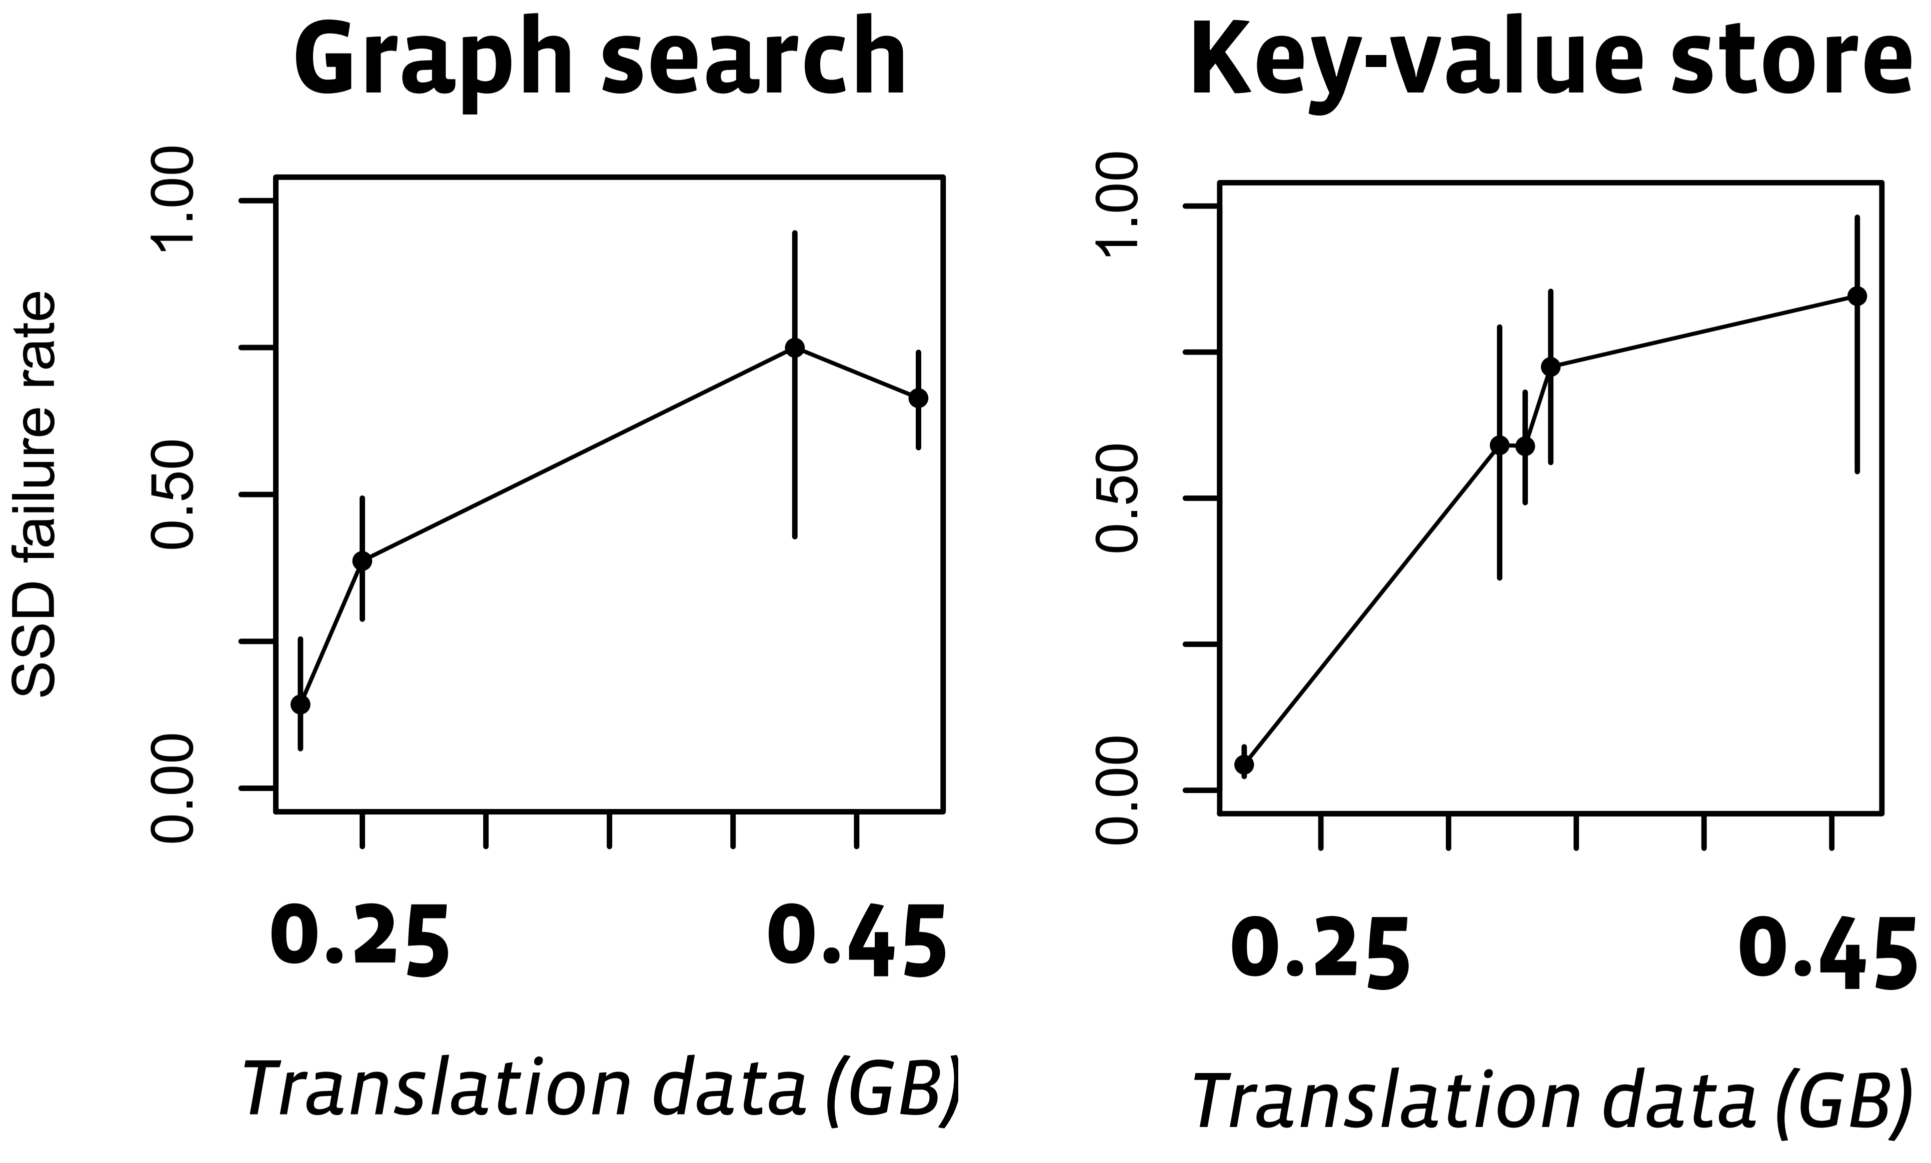

Why Does Sparsity of Data Matter?

- More translation data correlates with higher failure rates

- Sparse data updates, i.e., updates to less contiguous data, lead to more translation data

- Higher failure rates likely due to more frequent erase and copying caused by non-contiguous updates
  - Write amplification

SSD lifecycle

Access pattern dependence

We quantify the effects of the ***page cache*** and ***write amplification*** in the field.

Temperature

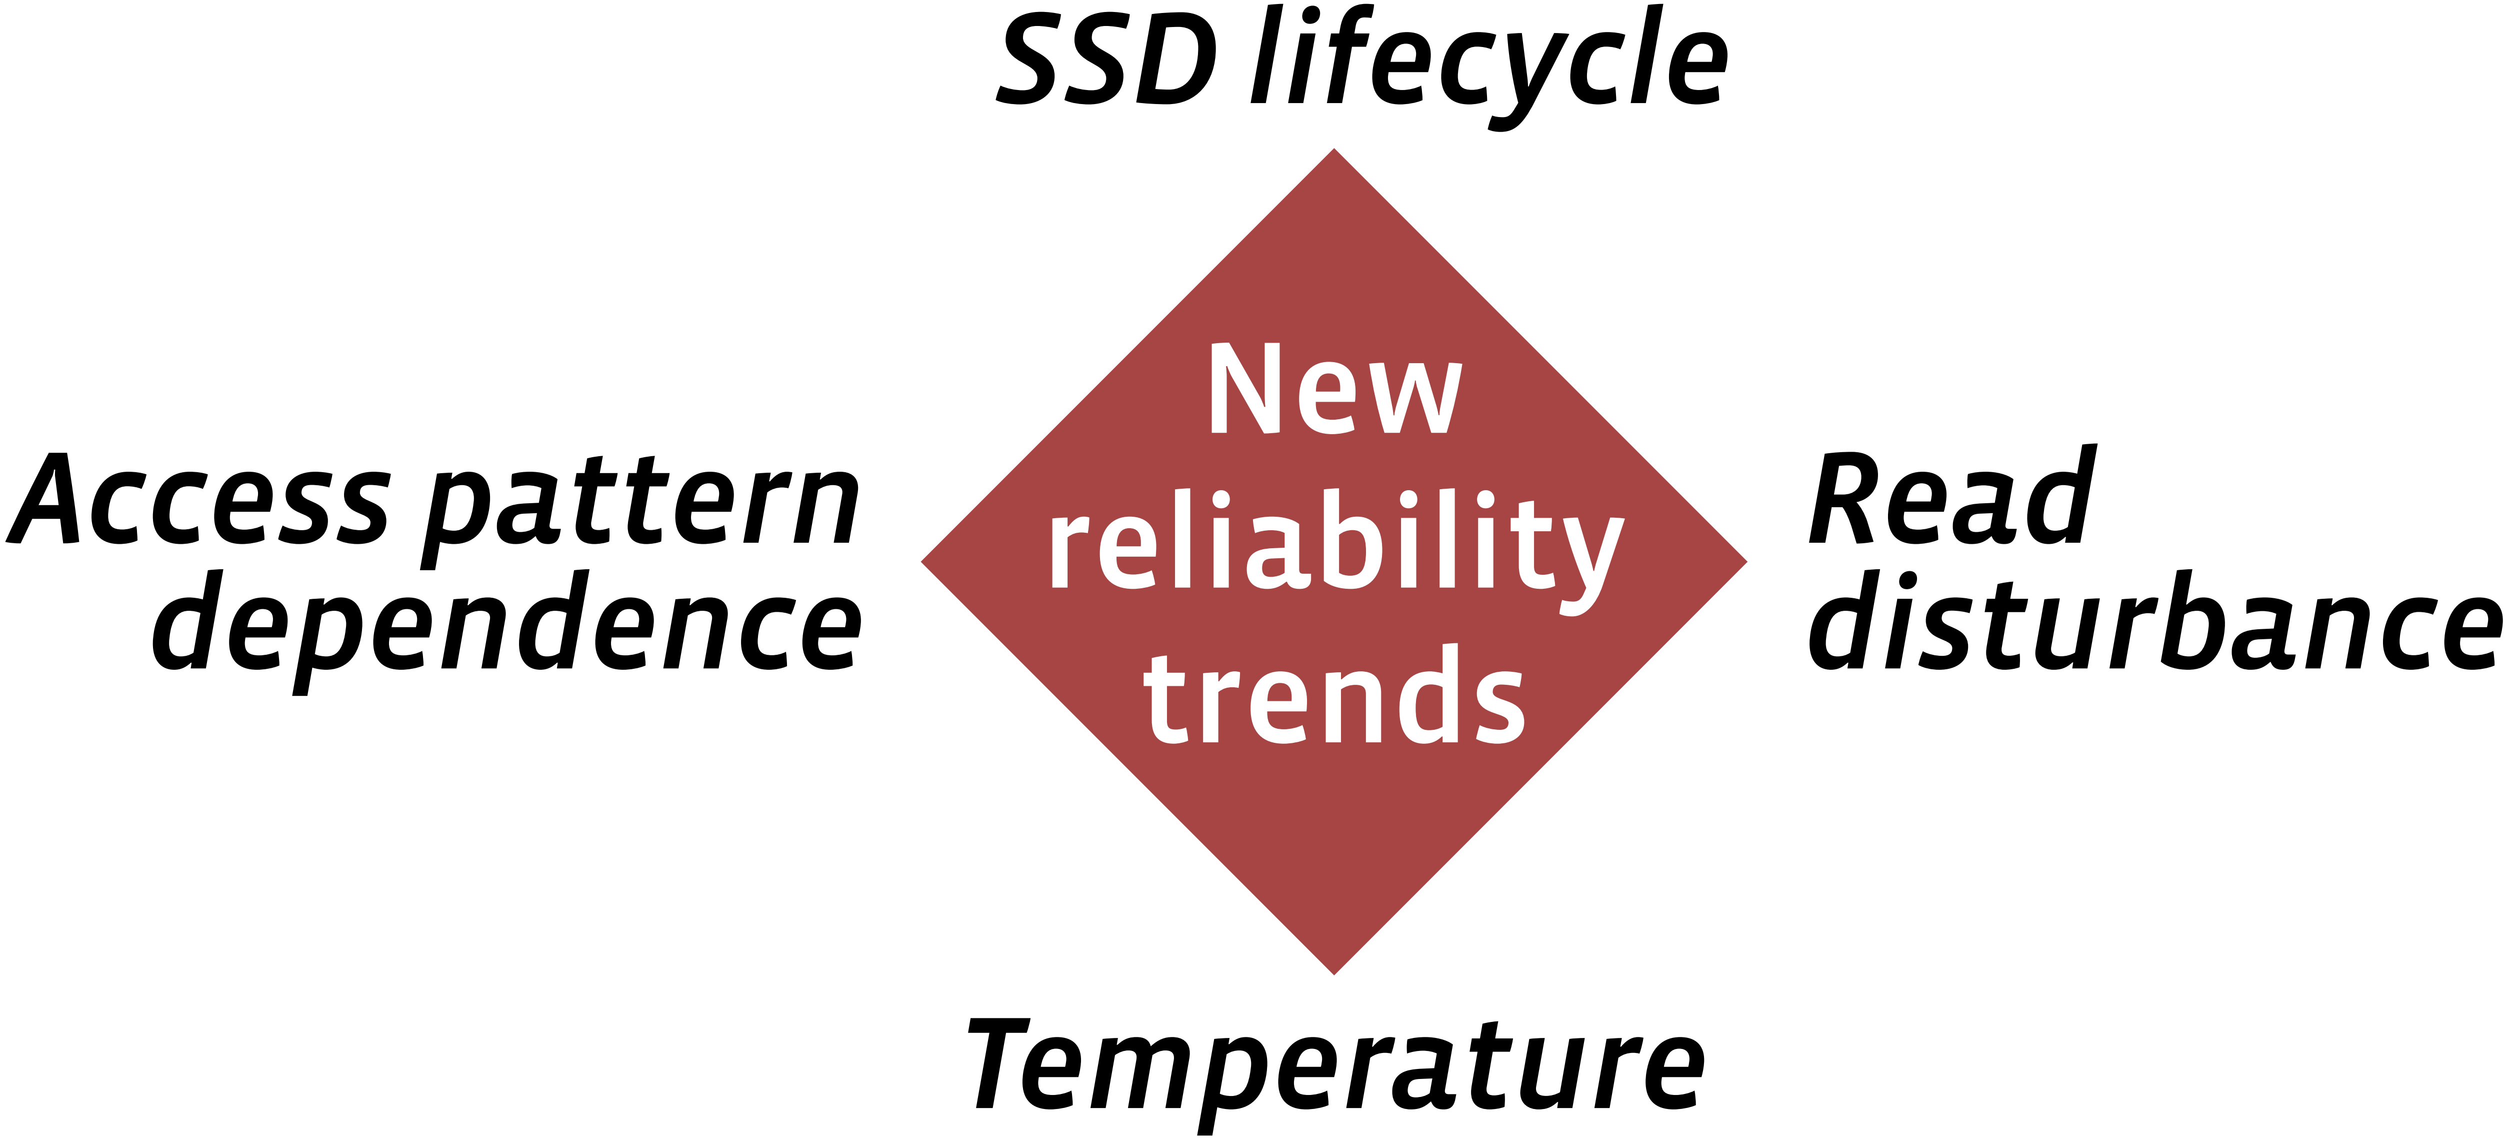

SSD lifecycle

Access pattern dependence

New reliability trends

Read disturbance

Temperature

More results in paper

- *Block erasures and discards*
- *Page copies*
- *Bus power consumption*

Summary

- *Large scale*
- *In the field*

Summary

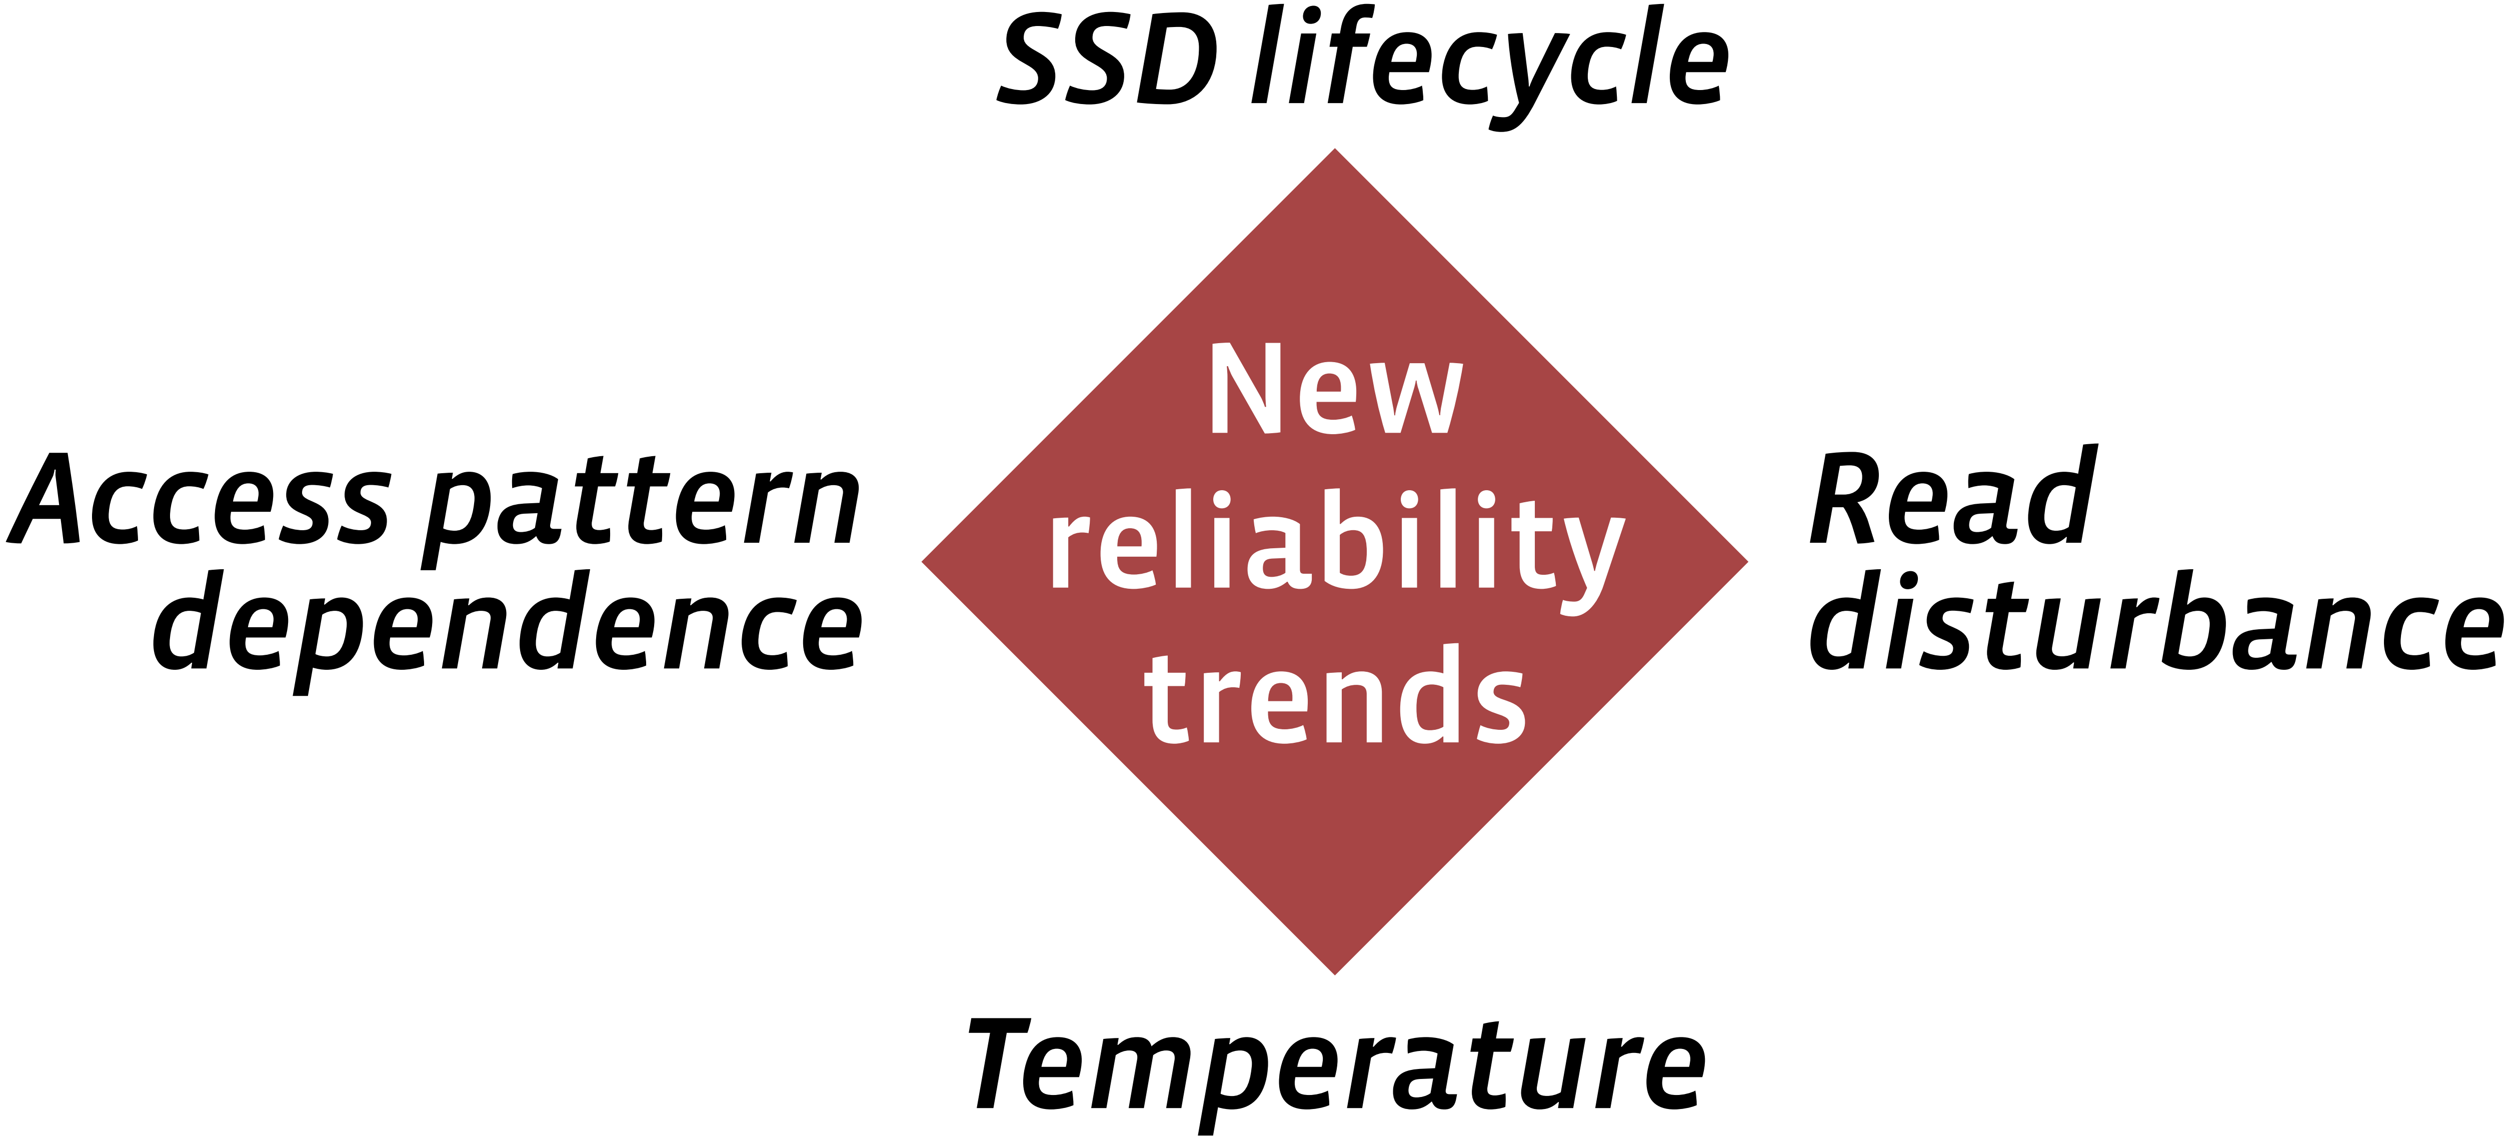

SSD lifecycle

Access pattern dependence

New reliability trends

Read disturbance

Temperature

Summary

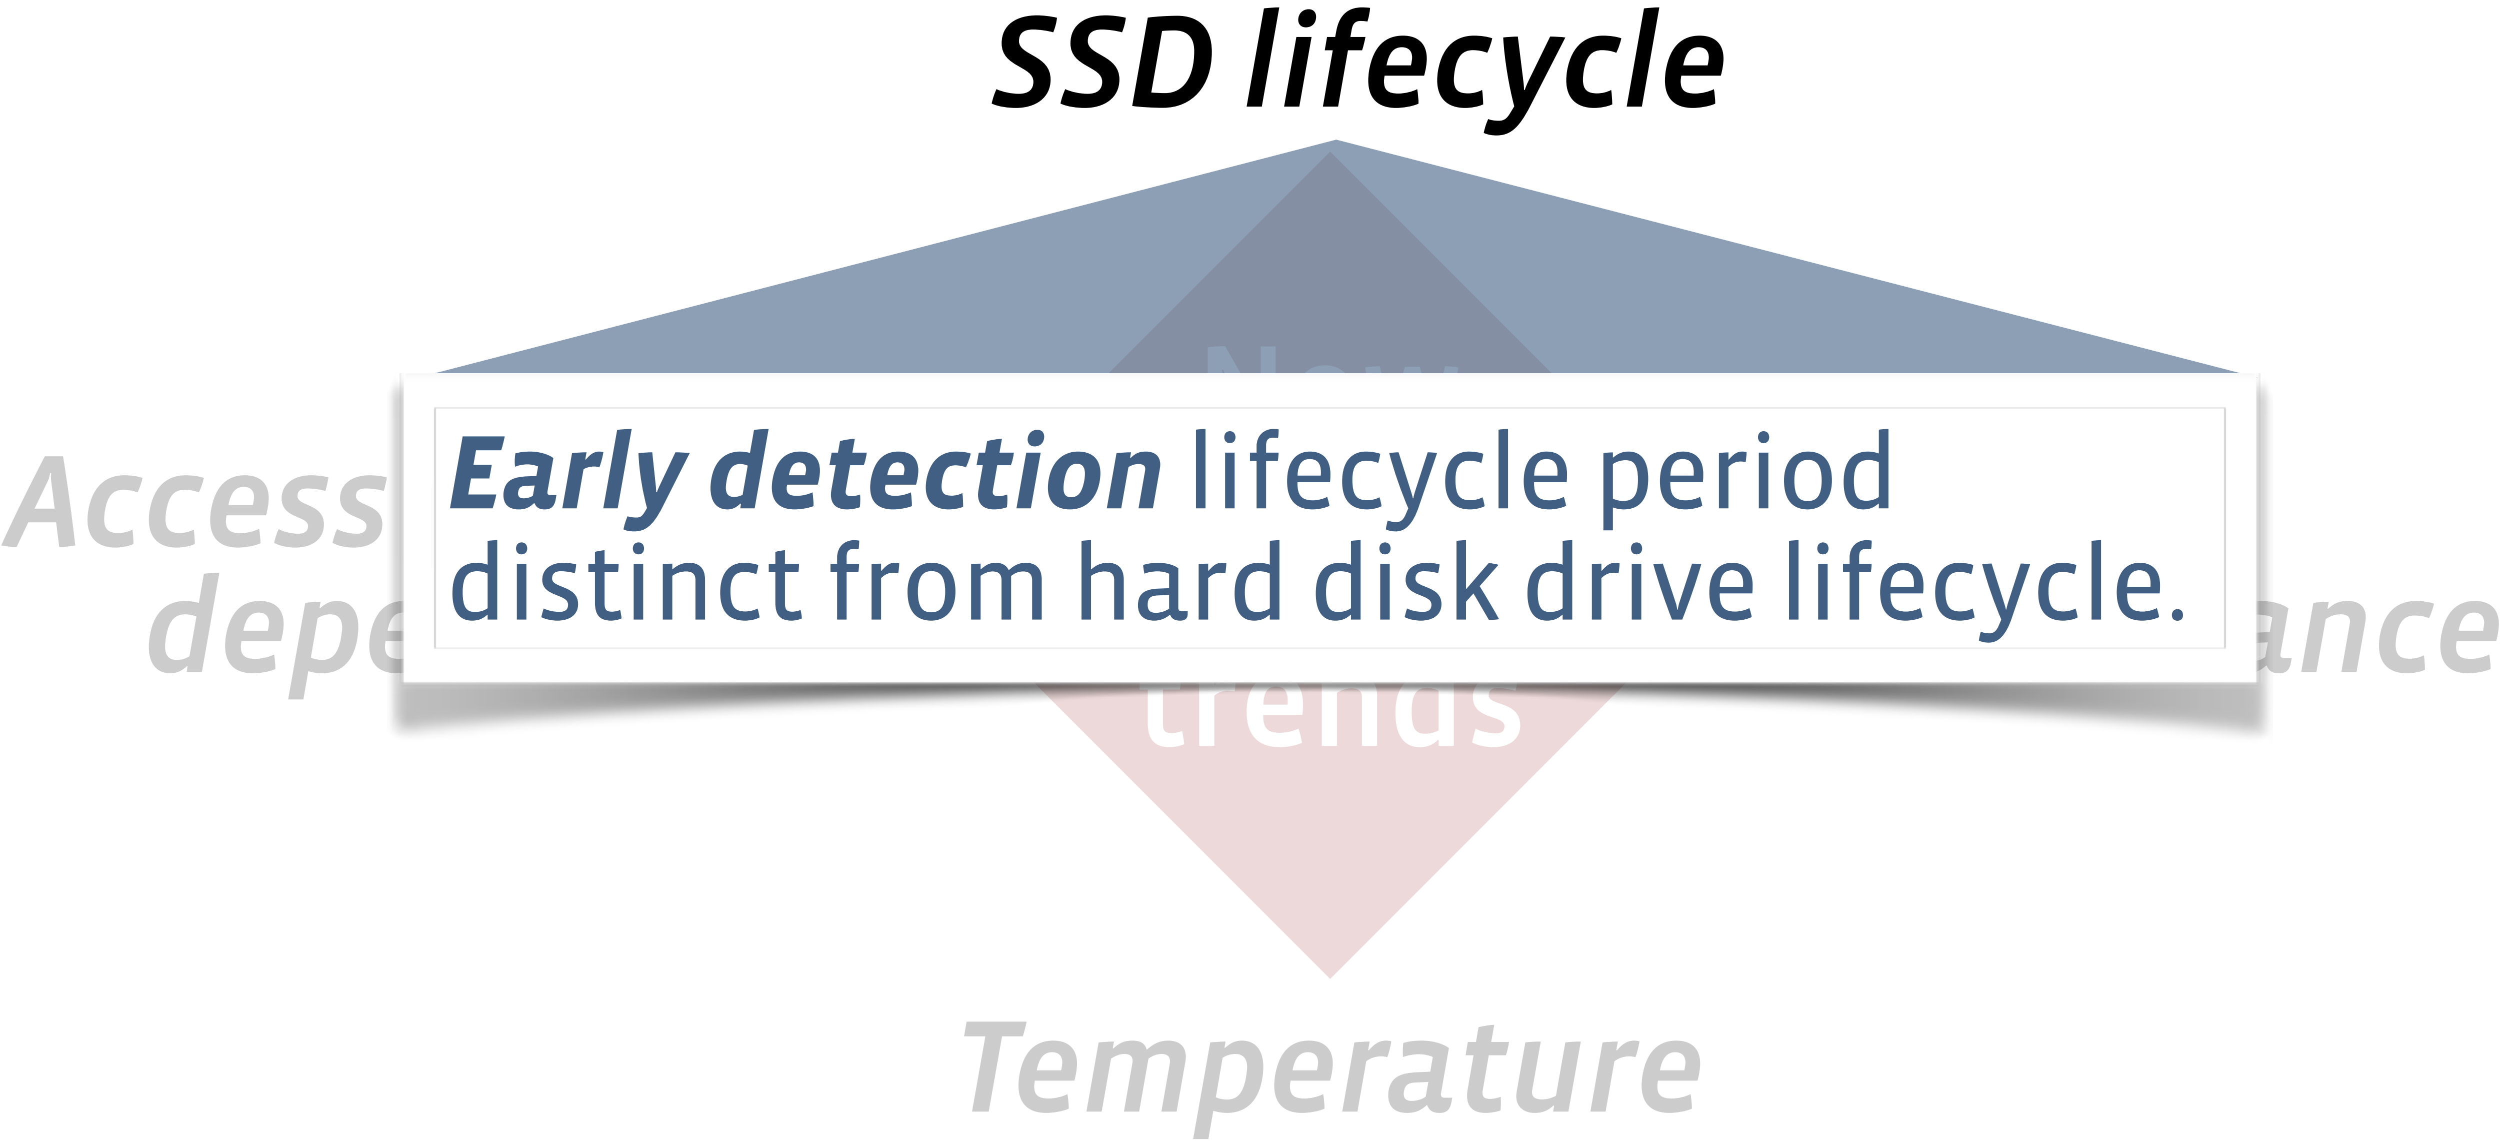

SSD lifecycle

Access depe

New trends

ance

Temperature

Early detection lifecycle period distinct from hard disk drive lifecycle.

Summary

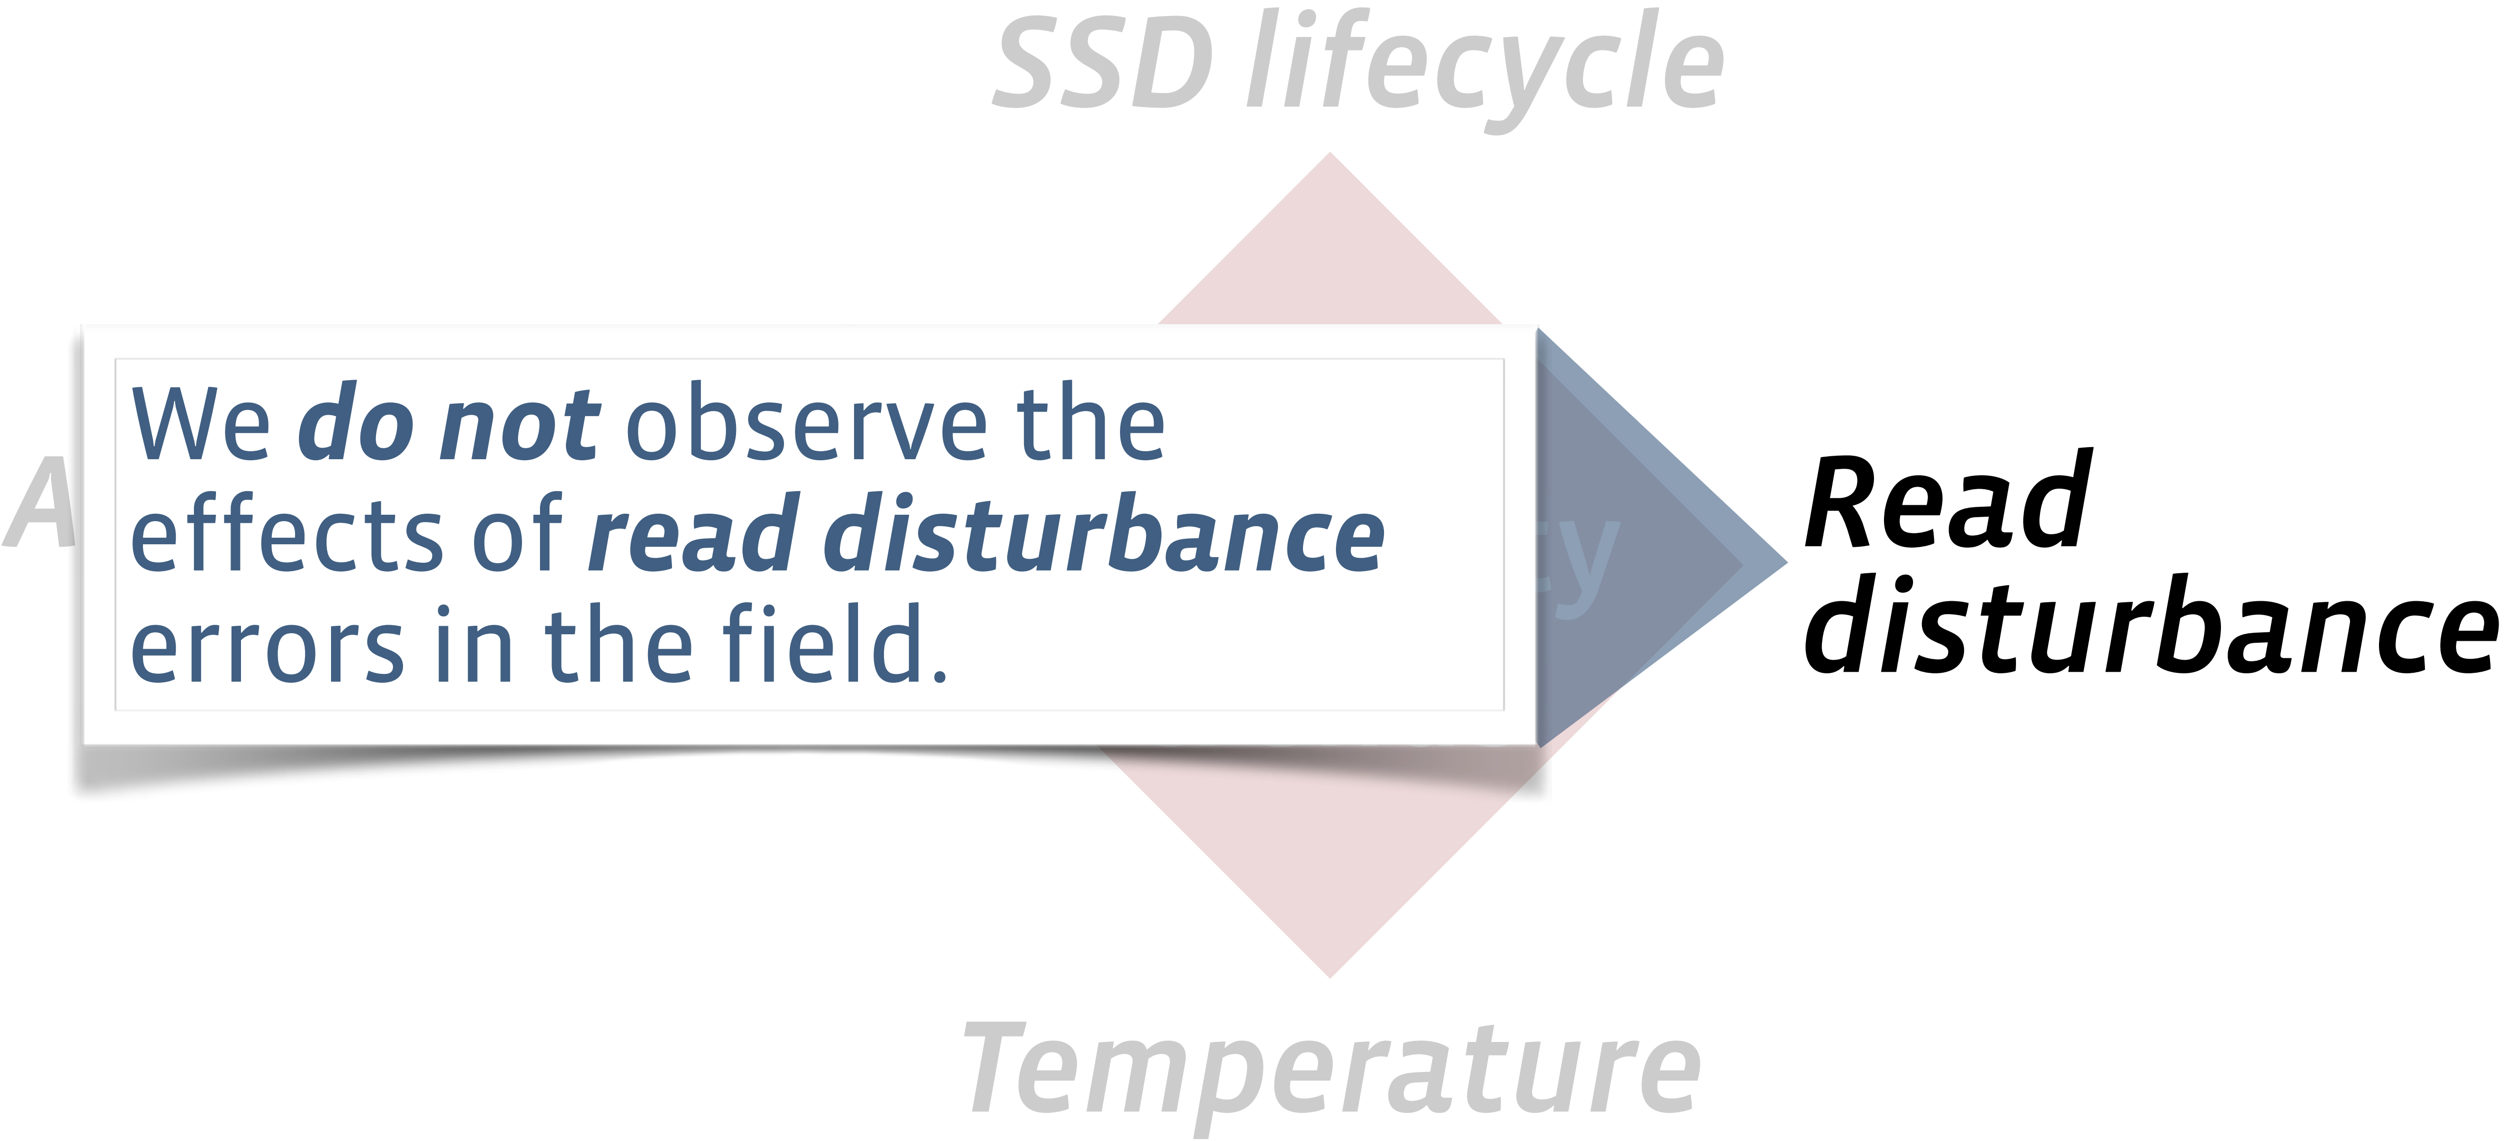

We ***do not*** observe the effects of ***read disturbance*** errors in the field.

Summary

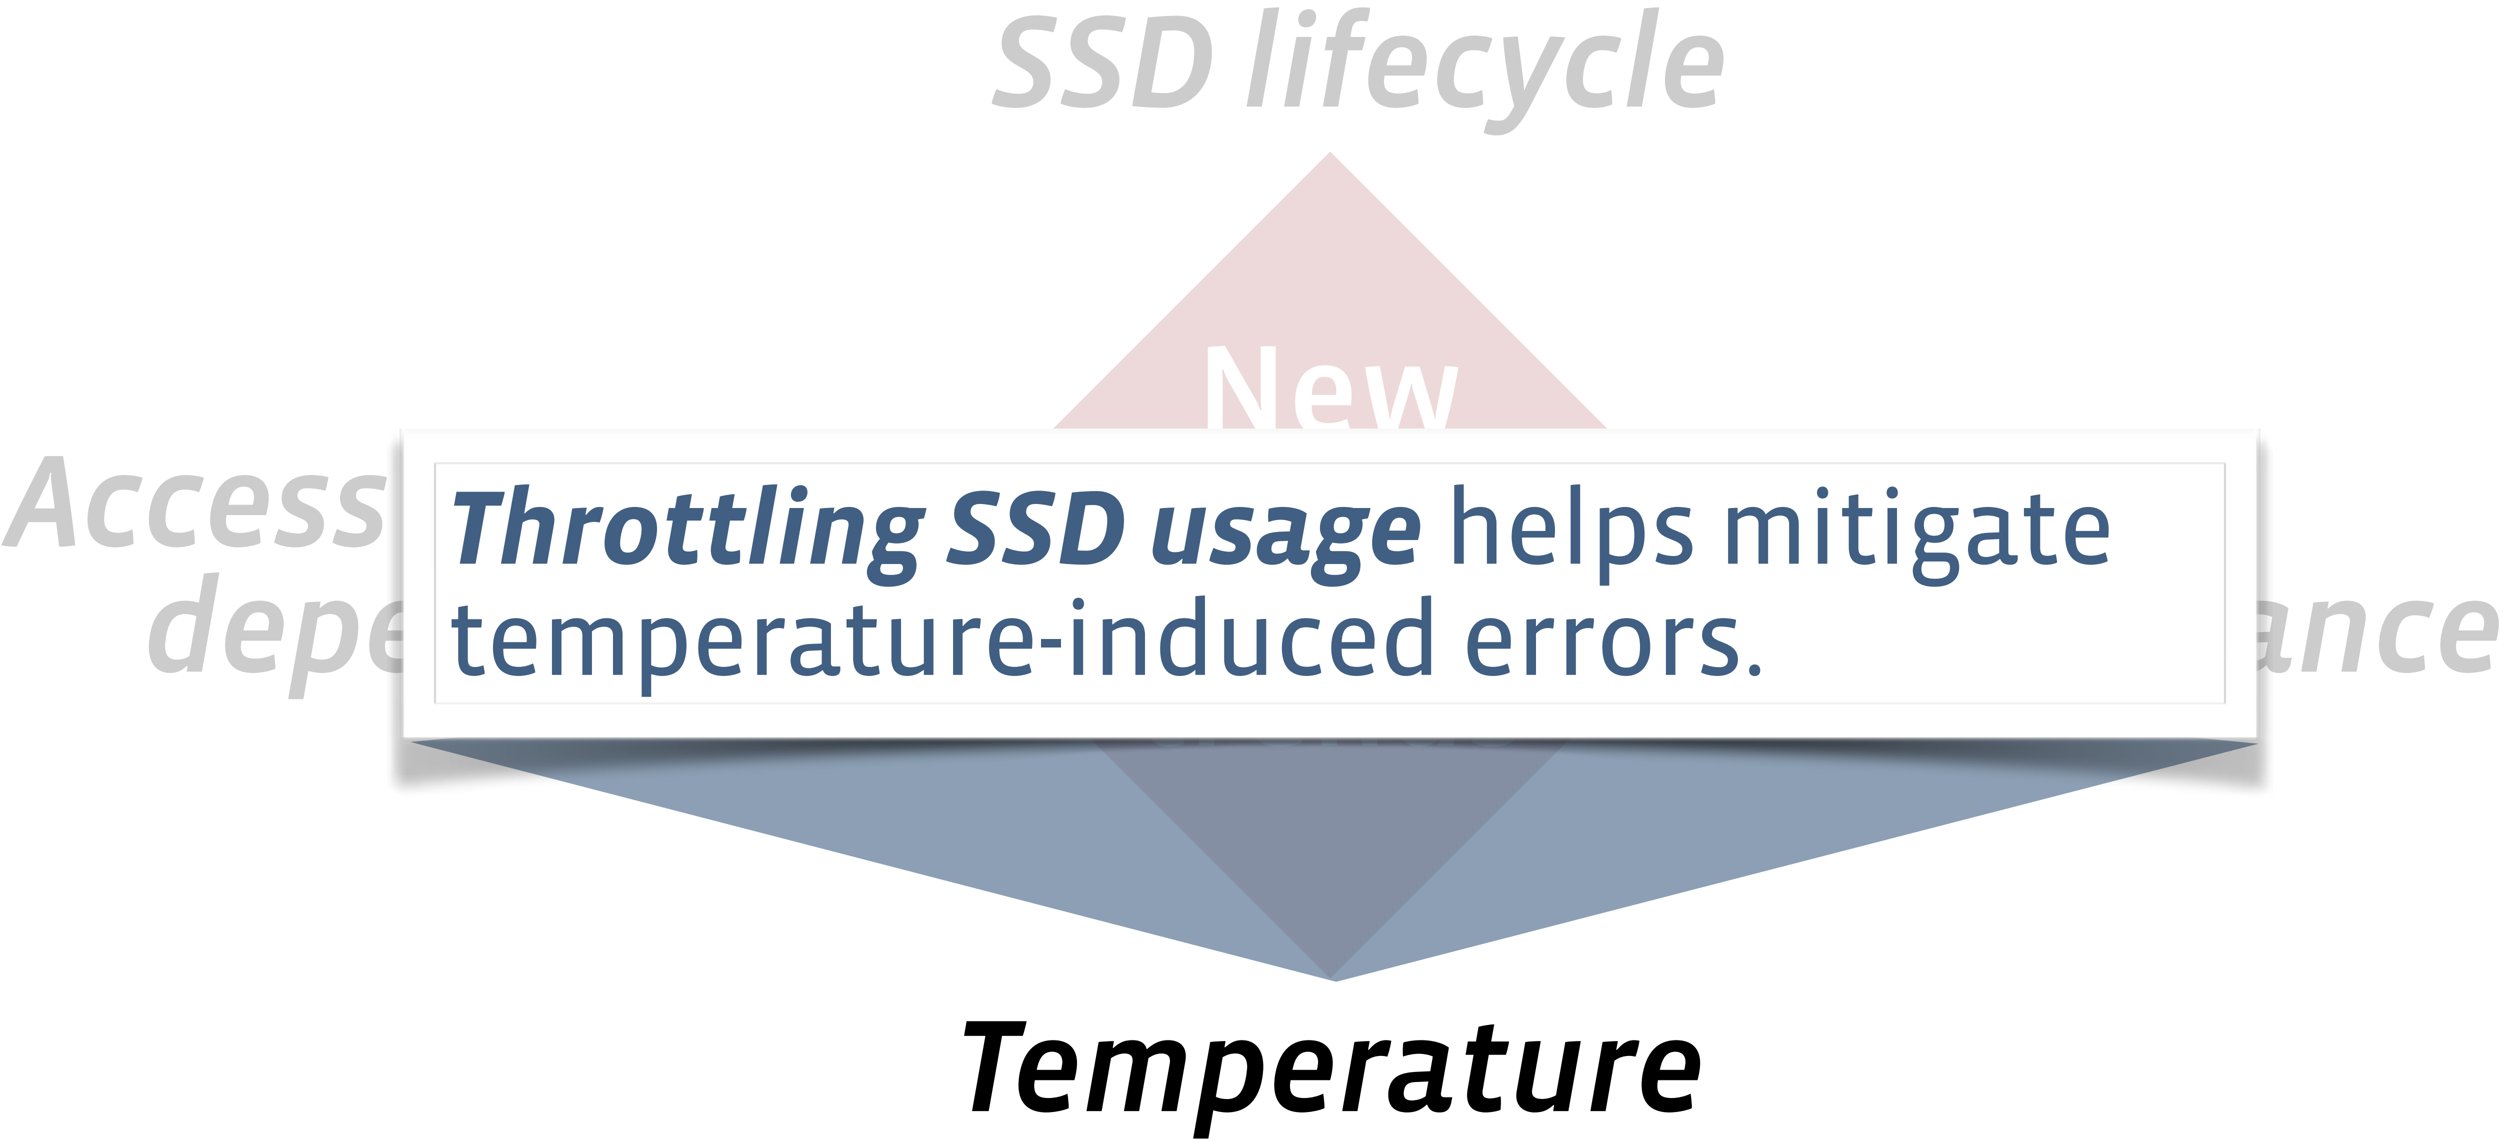

Throttling SSD usage helps mitigate temperature-induced errors.

Summary

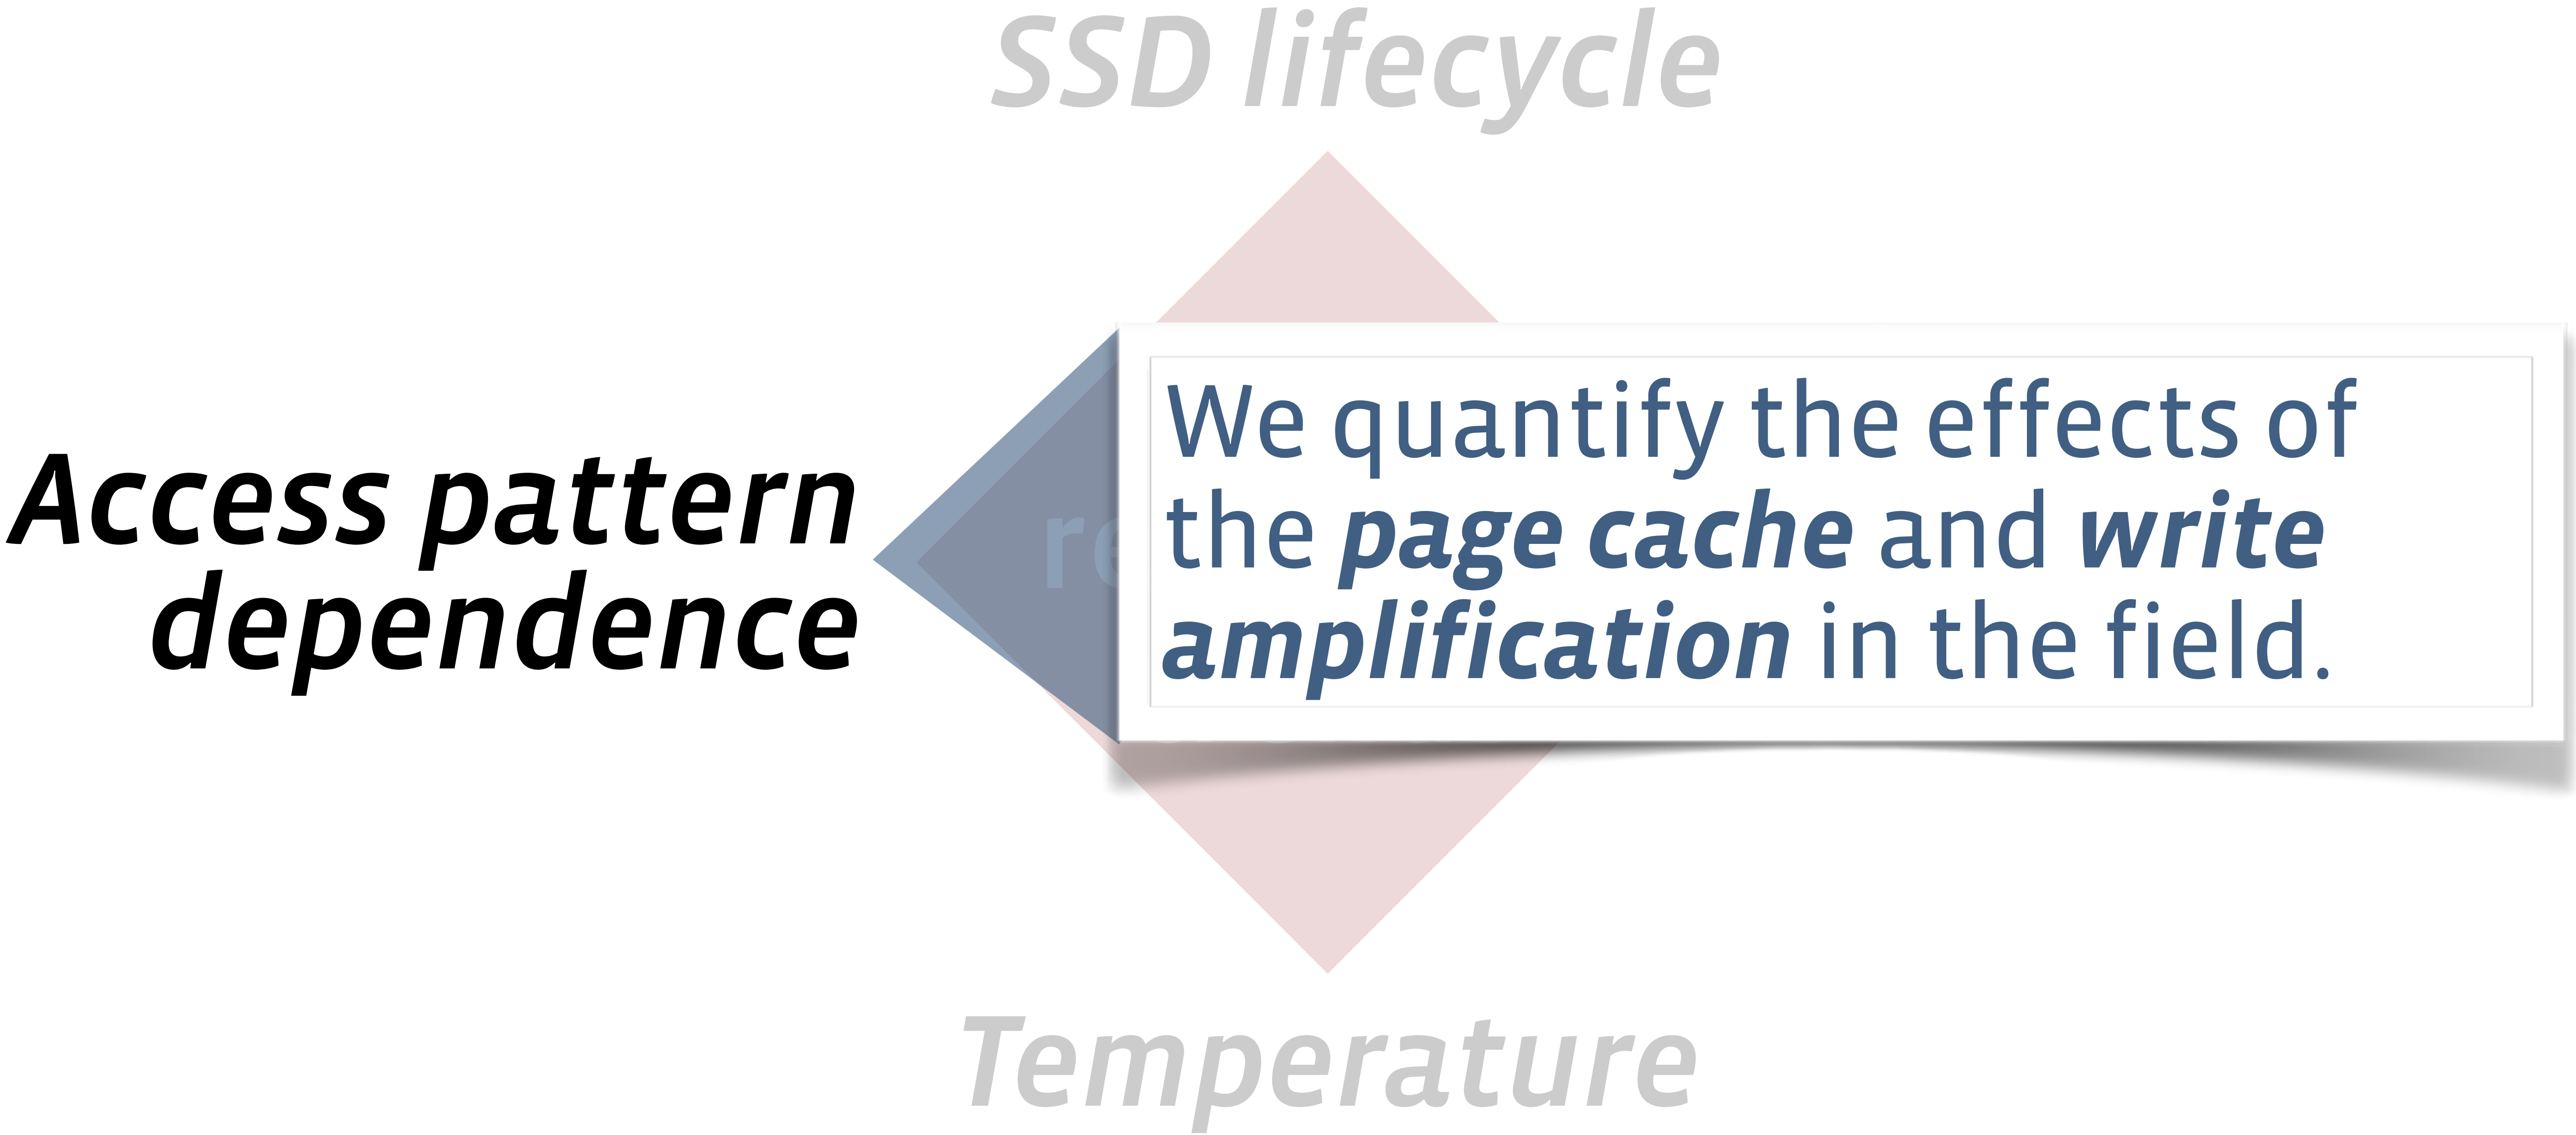

SSD lifecycle

Access pattern dependence

We quantify the effects of the ***page cache*** and ***write amplification*** in the field.

Temperature

A Large-Scale Study of
Flash Memory Errors in the Field

Justin Meza
Qiang Wu
Sanjeev Kumar
Onur Mutlu

facebook
Carnegie Mellon University

Our Remaining FMS 2016 Talks

- At 4:20pm Today
- Practical Threshold Voltage Distribution Modeling
  - Yixin Luo (CMU PhD Student) August 10 @ 4:20pm
  - Forum E-22: Controllers and Flash Technology

- At 5:45pm Today
- "WARM: Improving NAND Flash Memory Lifetime with Write-hotness Aware Retention Management"
  - Saugata Ghose (CMU Researcher) August 10 @ 5:45pm
  - Forum C-22: SSD Concepts (SSDs Track)

Referenced Papers and Talks

- All are available at
  http://users.ece.cmu.edu/~omutlu/projects.htm
  http://users.ece.cmu.edu/~omutlu/talks.htm

- And, many other previous works on
  - NVM & Persistent Memory
  - DRAM
  - Hybrid memories
  - NAND flash memory

Thank you.

Feel free to email me with any questions & feedback

omutlu@ethz.ch
http://users.ece.cmu.edu/~omutlu/

Large-Scale Study of In-the-Field Flash Failures

Onur Mutlu
omutlu@ethz.ch

(joint work with Justin Meza, Qiang Wu, Sanjeev Kumar)

August 10, 2016

Flash Memory Summit 2016, Santa Clara, CA

References to Papers and Talks

Challenges and Opportunities in Memory

- Onur Mutlu,
  "Rethinking Memory System Design"
  Keynote talk at
  2016 ACM SIGPLAN International Symposium on Memory Management (***ISMM***), Santa Barbara, CA, USA, June 2016.
  [Slides (pptx) (pdf)]
  [Abstract]

- Onur Mutlu and Lavanya Subramanian,
  "Research Problems and Opportunities in Memory Systems"
  Invited Article in Supercomputing Frontiers and Innovations (***SUPERFRI***), 2015.

Our FMS Talks and Posters

- Onur Mutlu, *ThyNVM: Software-Transparent Crash Consistency for Persistent Memory*, FMS 2016.
- Onur Mutlu, *Large-Scale Study of In-the-Field Flash Failures*, FMS 2016.
- Yixin Luo, *Practical Threshold Voltage Distribution Modeling*, FMS 2016.
- Saugata Ghose, *Write-hotness Aware Retention Management*, FMS 2016.
- Onur Mutlu, *Read Disturb Errors in MLC NAND Flash Memory*, FMS 2015.
- Yixin Luo, *Data Retention in MLC NAND Flash Memory*, FMS 2015.
- Onur Mutlu, *Error Analysis and Management for MLC NAND Flash Memory*, FMS 2014.

- FMS 2016 posters:
  - WARM: Improving NAND Flash Memory Lifetime with Write-hotness Aware Retention Management
  - Read Disturb Errors in MLC NAND Flash Memory
  - Data Retention in MLC NAND Flash Memory

Our Flash Memory Works (I)

1. Retention noise study and management

1) Yu Cai, Gulay Yalcin, Onur Mutlu, Erich F. Haratsch, Adrian Cristal, Osman Unsal, and Ken Mai, **Flash Correct-and-Refresh: Retention-Aware Error Management for Increased Flash Memory Lifetime**, ICCD 2012.
2) Yu Cai, Yixin Luo, Erich F. Haratsch, Ken Mai, and Onur Mutlu, **Data Retention in MLC NAND Flash Memory: Characterization, Optimization and Recovery**, HPCA 2015.
3) Yixin Luo, Yu Cai, Saugata Ghose, Jongmoo Choi, and Onur Mutlu, **WARM: Improving NAND Flash Memory Lifetime with Write-hotness Aware Retention Management**, MSST 2015.

2. Flash-based SSD prototyping and testing platform

4) Yu Cai, Erich F. Haratsh, Mark McCartney, Ken Mai, **FPGA-based solid-state drive prototyping platform**, FCCM 2011.

Our Flash Memory Works (II)

3. Overall flash error analysis

5) Yu Cai, Erich F. Haratsch, Onur Mutlu, and Ken Mai,
Error Patterns in MLC NAND Flash Memory: Measurement, Characterization, and Analysis, DATE 2012.

6) Yu Cai, Gulay Yalcin, Onur Mutlu, Erich F. Haratsch, Adrian Cristal, Osman Unsal, and Ken Mai,
Error Analysis and Retention-Aware Error Management for NAND Flash Memory, ITJ 2013.

4. Program and erase noise study

7) Yu Cai, Erich F. Haratsch, Onur Mutlu, and Ken Mai,
Threshold Voltage Distribution in MLC NAND Flash Memory: Characterization, Analysis and Modeling, DATE 2013.

Our Flash Memory Works (III)

5. Cell-to-cell interference characterization and tolerance

8) Yu Cai, Onur Mutlu, Erich F. Haratsch, and Ken Mai, **Program Interference in MLC NAND Flash Memory: Characterization, Modeling, and Mitigation**, ICCD 2013.

9) Yu Cai, Gulay Yalcin, Onur Mutlu, Erich F. Haratsch, Osman Unsal, Adrian Cristal, and Ken Mai, **Neighbor-Cell Assisted Error Correction for MLC NAND Flash Memories**, SIGMETRICS 2014.

6. Read disturb noise study

10) Yu Cai, Yixin Luo, Saugata Ghose, Erich F. Haratsch, Ken Mai, and Onur Mutlu, **Read Disturb Errors in MLC NAND Flash Memory: Characterization and Mitigation**, DSN 2015.

Our Flash Memory Works (IV)

7. Flash errors in the field

11) Justin Meza, Qiang Wu, Sanjeev Kumar, and Onur Mutlu, A Large-Scale Study of Flash Memory Errors in the Field, SIGMETRICS 2015.

8. Persistent memory

12) Jinglei Ren, Jishen Zhao, Samira Khan, Jongmoo Choi, Yongwei Wu, and Onur Mutlu, ThyNVM: Enabling Software-Transparent Crash Consistency in Persistent Memory Systems, MICRO 2015.

Phase Change Memory As DRAM Replacement

- Benjamin C. Lee, Engin Ipek, Onur Mutlu, and Doug Burger, **"Architecting Phase Change Memory as a Scalable DRAM Alternative"** *Proceedings of the 36th International Symposium on Computer Architecture (**ISCA**)*, pages 2-13, Austin, TX, June 2009. Slides (pdf)

- Benjamin C. Lee, Ping Zhou, Jun Yang, Youtao Zhang, Bo Zhao, Engin Ipek, Onur Mutlu, and Doug Burger, **"Phase Change Technology and the Future of Main Memory"** *IEEE Micro, Special Issue: Micro's Top Picks from 2009 Computer Architecture Conferences (**MICRO TOP PICKS**)*, Vol. 30, No. 1, pages 60-70, January/February 2010.

STT-MRAM As DRAM Replacement

- Emre Kultursay, Mahmut Kandemir, Anand Sivasubramaniam, and Onur Mutlu, **"Evaluating STT-RAM as an Energy-Efficient Main Memory Alternative"** *Proceedings of the 2013 IEEE International Symposium on Performance Analysis of Systems and Software (**ISPASS**)*, Austin, TX, April 2013. Slides (pptx) (pdf)

Taking Advantage of Persistence in Memory

- Justin Meza, Yixin Luo, Samira Khan, Jishen Zhao, Yuan Xie, and Onur Mutlu,
  "A Case for Efficient Hardware-Software Cooperative Management of Storage and Memory"
  *Proceedings of the 5th Workshop on Energy-Efficient Design (**WEED**)*, Tel-Aviv, Israel, June 2013. Slides (pptx) Slides (pdf)

- Jinglei Ren, Jishen Zhao, Samira Khan, Jongmoo Choi, Yongwei Wu, and Onur Mutlu,
  "ThyNVM: Enabling Software-Transparent Crash Consistency in Persistent Memory Systems"
  *Proceedings of the 48th International Symposium on Microarchitecture (**MICRO**)*, Waikiki, Hawaii, USA, December 2015.
  [Slides (pptx) (pdf)] [Lightning Session Slides (pptx) (pdf)] [
  Poster (pptx) (pdf)]
  [Source Code]

Hybrid DRAM + NVM Systems (I)

- HanBin Yoon, Justin Meza, Rachata Ausavarungnirun, Rachael Harding, and Onur Mutlu,
  "Row Buffer Locality Aware Caching Policies for Hybrid Memories"
  Proceedings of the
  30th IEEE International Conference on Computer Design (***ICCD***), Montreal, Quebec, Canada, September 2012. Slides (pptx) (pdf)
  Best paper award (in Computer Systems and Applications track).

- Justin Meza, Jichuan Chang, HanBin Yoon, Onur Mutlu, and Parthasarathy Ranganathan,
  "Enabling Efficient and Scalable Hybrid Memories Using Fine-Granularity DRAM Cache Management"
  IEEE Computer Architecture Letters (***CAL***), February 2012.

Hybrid DRAM + NVM Systems (II)

- Dongwoo Kang, Seungjae Baek, Jongmoo Choi, Donghee Lee, Sam H. Noh, and Onur Mutlu, **"Amnesic Cache Management for Non-Volatile Memory"** *Proceedings of the 31st International Conference on Massive Storage Systems and Technologies (**MSST**)*, Santa Clara, CA, June 2015. [Slides (pdf)]

NVM Design and Architecture

- HanBin Yoon, Justin Meza, Naveen Muralimanohar, Norman P. Jouppi, and Onur Mutlu,
  "Efficient Data Mapping and Buffering Techniques for Multi-Level Cell Phase-Change Memories"
  *ACM Transactions on Architecture and Code Optimization (**TACO**)*, Vol. 11, No. 4, December 2014. [Slides (ppt) (pdf)]
  Presented at the 10th HiPEAC Conference, Amsterdam, Netherlands, January 2015.
  [Slides (ppt) (pdf)]

- Justin Meza, Jing Li, and Onur Mutlu,
  "Evaluating Row Buffer Locality in Future Non-Volatile Main Memories"
  SAFARI Technical Report, TR-SAFARI-2012-002, Carnegie Mellon University, December 2012.

Referenced Papers and Talks

- All are available at
  http://users.ece.cmu.edu/~omutlu/projects.htm
  http://users.ece.cmu.edu/~omutlu/talks.htm

- And, many other previous works on NAND flash memory errors and management

Related Videos and Course Materials

- Undergraduate Computer Architecture Course Lecture Videos (2013, 2014, 2015)

- Undergraduate Computer Architecture Course Materials (2013, 2014, 2015)

- Graduate Computer Architecture Lecture Videos (2013, 2015)

- Parallel Computer Architecture Course Materials (Lecture Videos)

- Memory Systems Short Course Materials (Lecture Video on Main Memory and DRAM Basics)

Additional Slides

Backup slides

System characteristics

SSD capacity	PCIe	Average age (years)	SSDs per server	Average written (TB)	Average read (TB)
720GB	v1, x4	2.4	1	27.2	23.8
			2	48.5	45.1
1.2TB	v2, x4	1.6	1	37.8	43.4
			2	18.9	30.6
3.2TB	v2, x4	0.5	1	23.9	51.1
			2	14.8	18.2

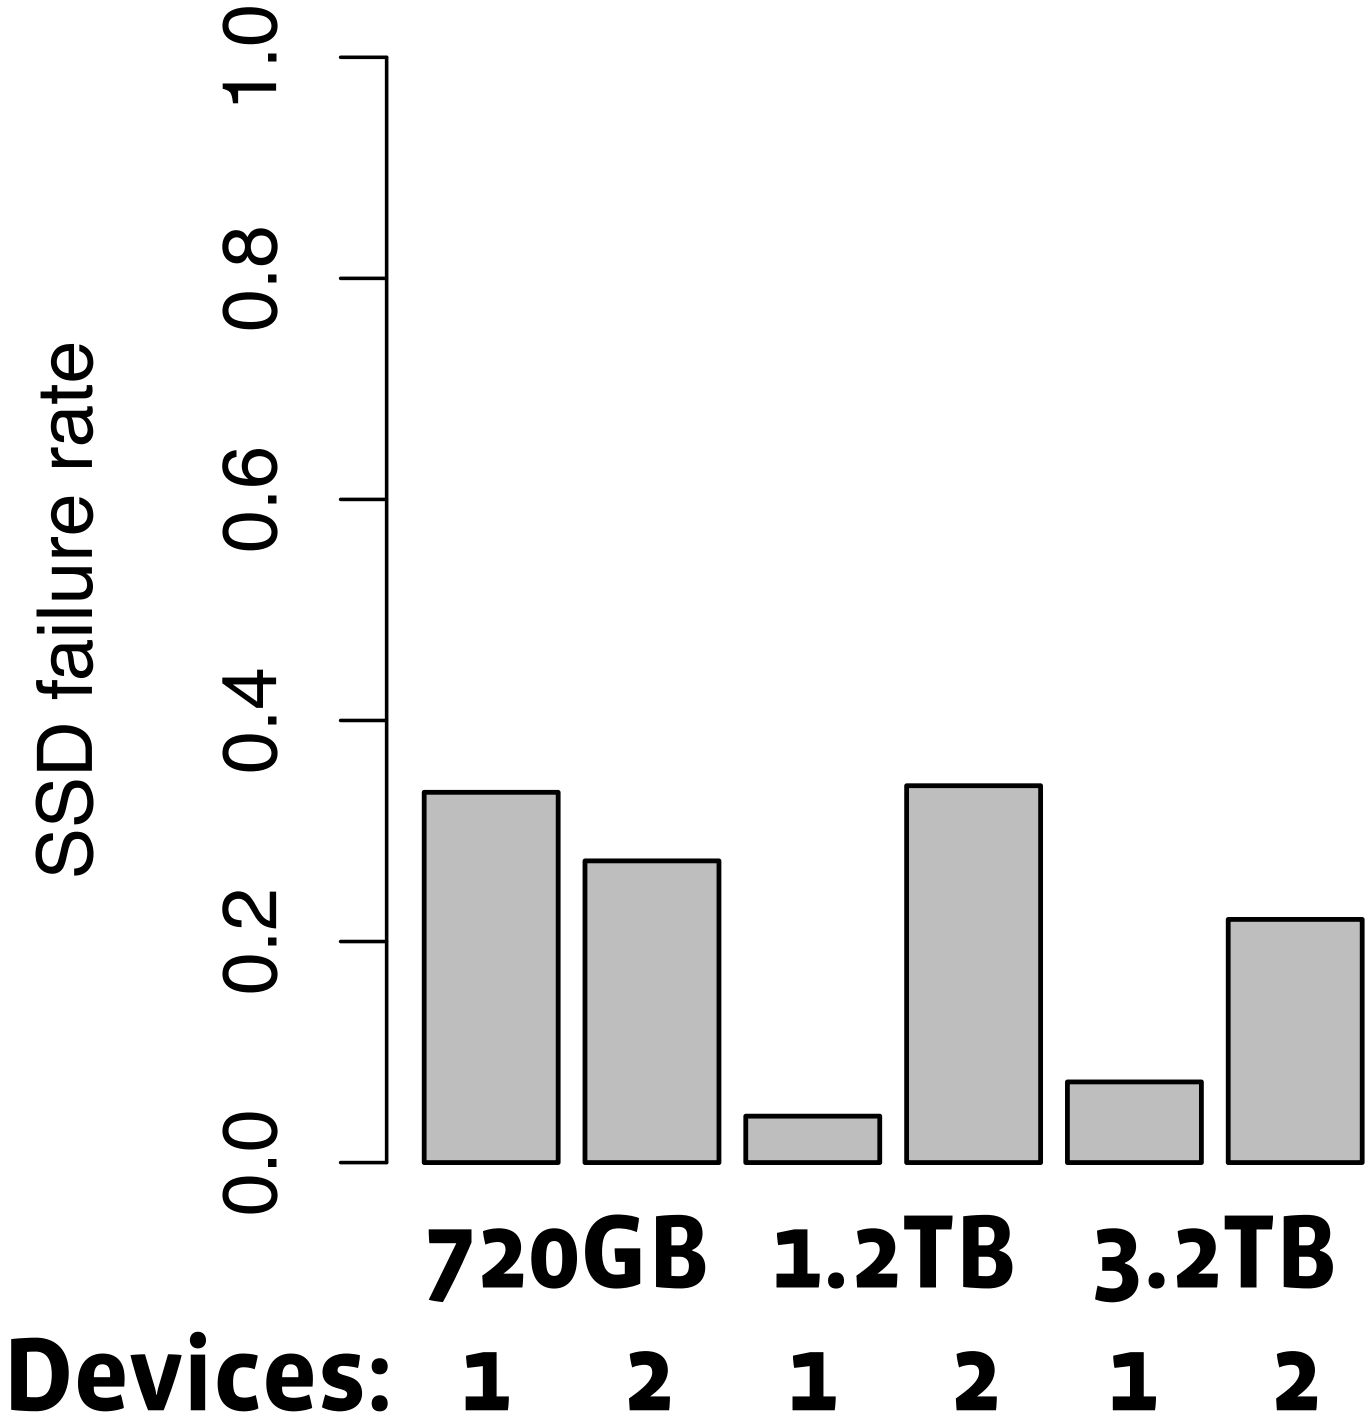
SSD failure rate
0.0
0.2
0.4
0.6
0.8
1.0
720GB
1.2TB
3.2TB
Devices:
1
2
1
2
1
2

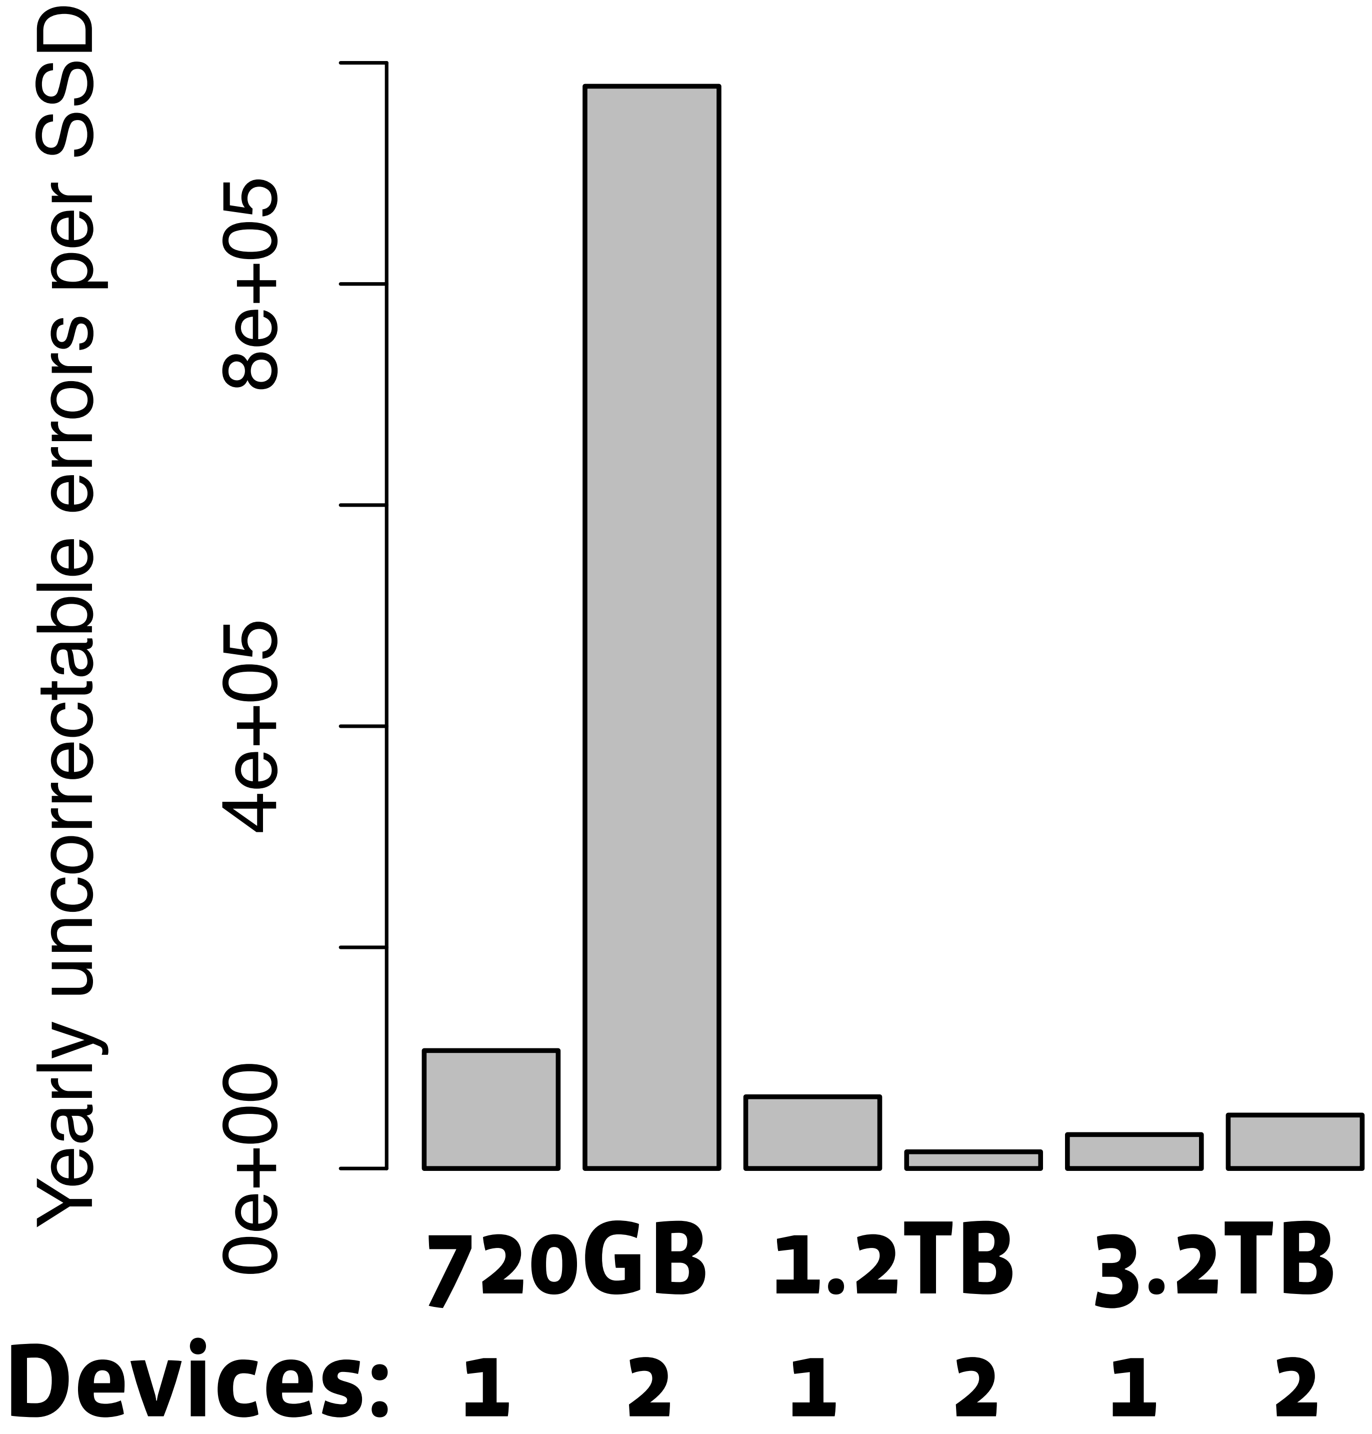
Yearly uncorrectable errors per SSD
0e+00
4e+05
8e+05
720GB
1.2TB
3.2TB
Devices:
1
2
1
2
1
2

Channels operate in parallel

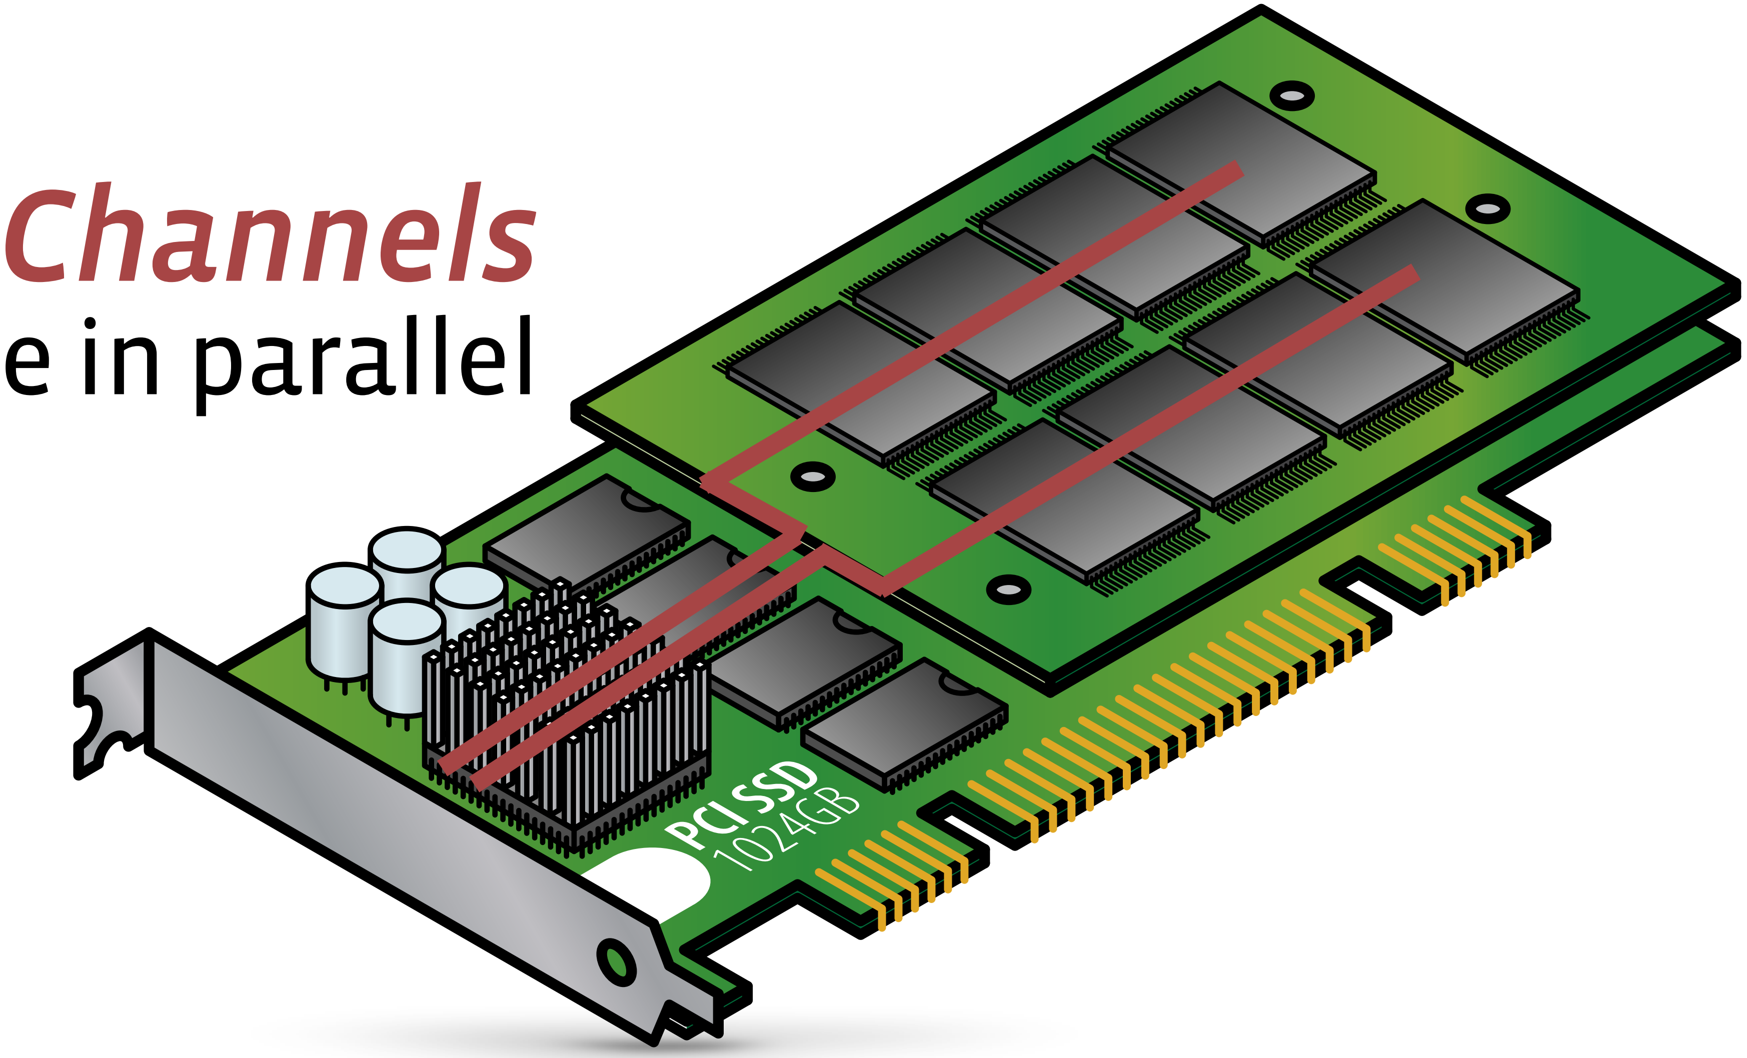

DRAM buffer

- stores address translations
- may buffer writes

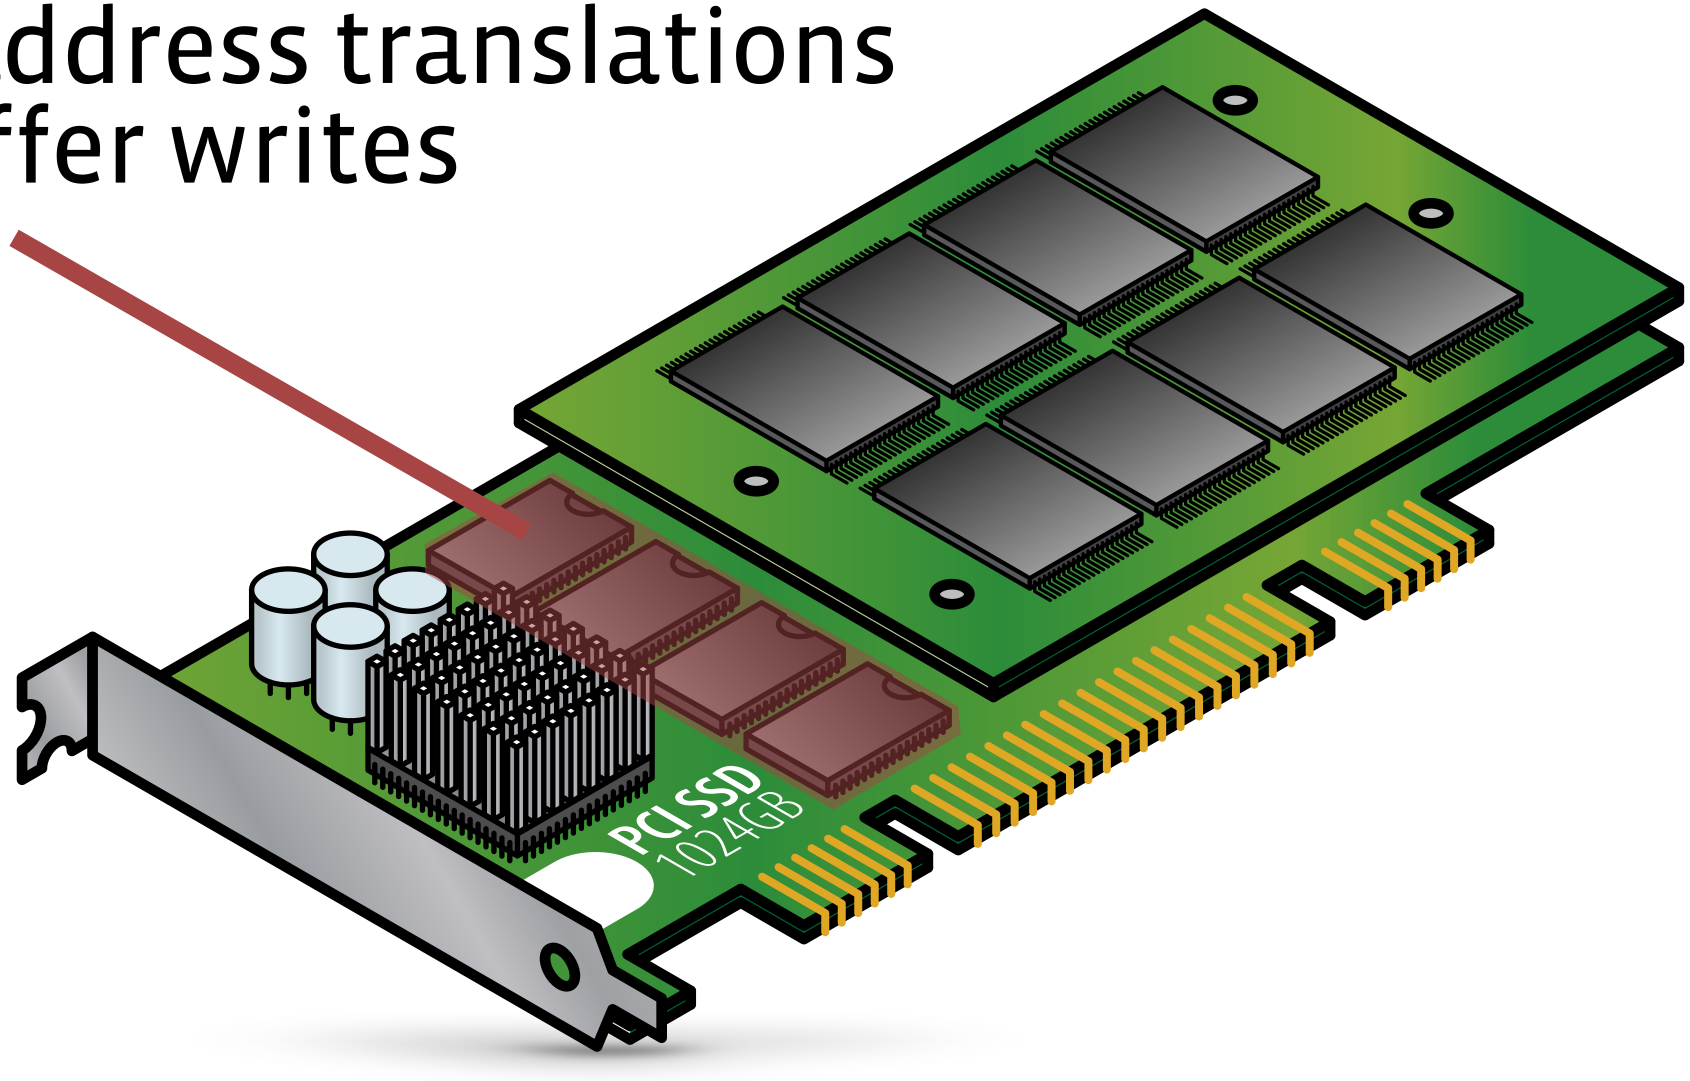

1.2TB, 2 SSDs

3.2TB, 2 SSDs

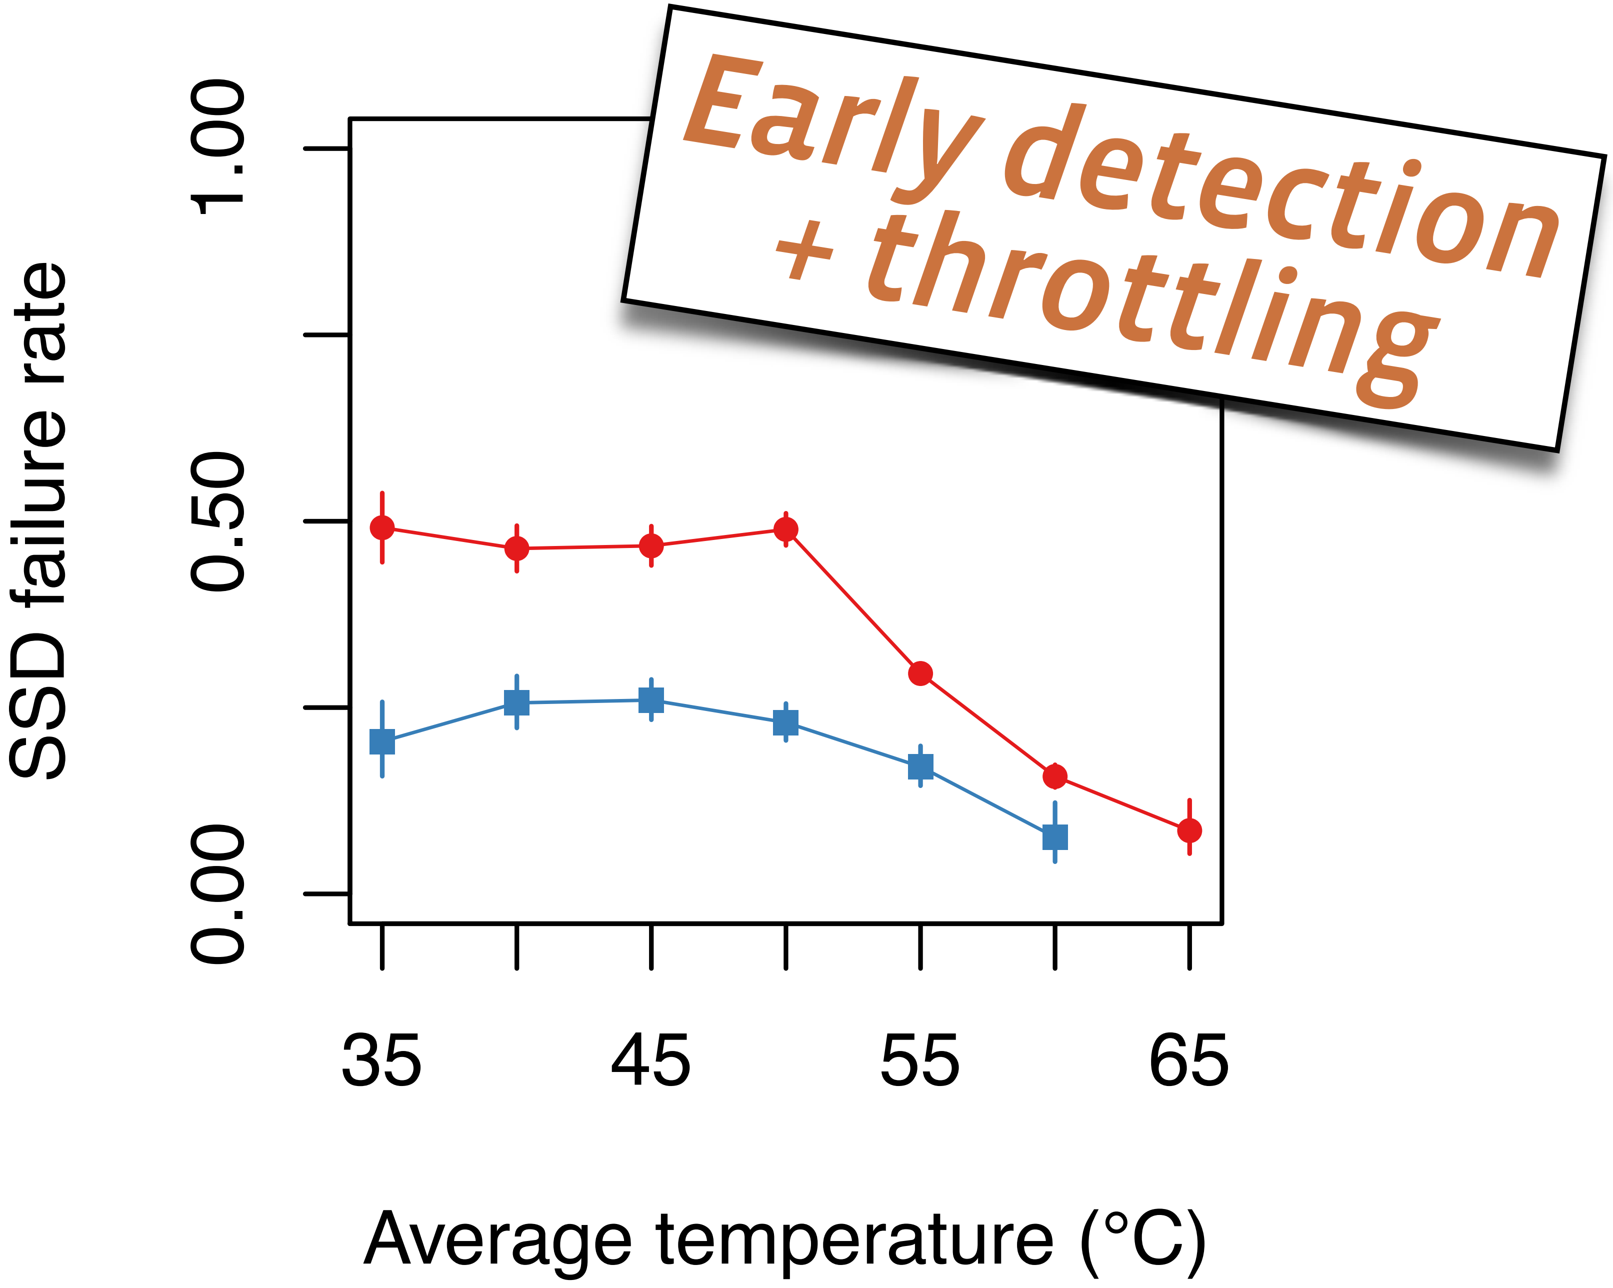

Early detection + throttling