AutoSSD: CXL-Enhanced Autonomous SSDs for Low Tail Latency

Mingyao Shen[†], Heewoo Kim[‡], Suyash Mahar[†], Joseph Izraelevitz[‡], Steven Swanson[†]
[†]UC San Diego, [‡]University of Colorado, Boulder

HPDC 2025
2025/07/21

NVSL

UCSDCSE Computer Science and Engineering

Outline

- Background and Motivation
- Problems
- Solutions
- Evaluation & Conclusion

NVSL

Outline

- Background and Motivation
- Problems
- Solutions
- Evaluation & Conclusion

NVSL

Compute Express Link (CXL) is Here

- Standard for high-speed communication
- Three sub-protocols
  - CXL.io
  - CXL.mem
  - CXL.cache

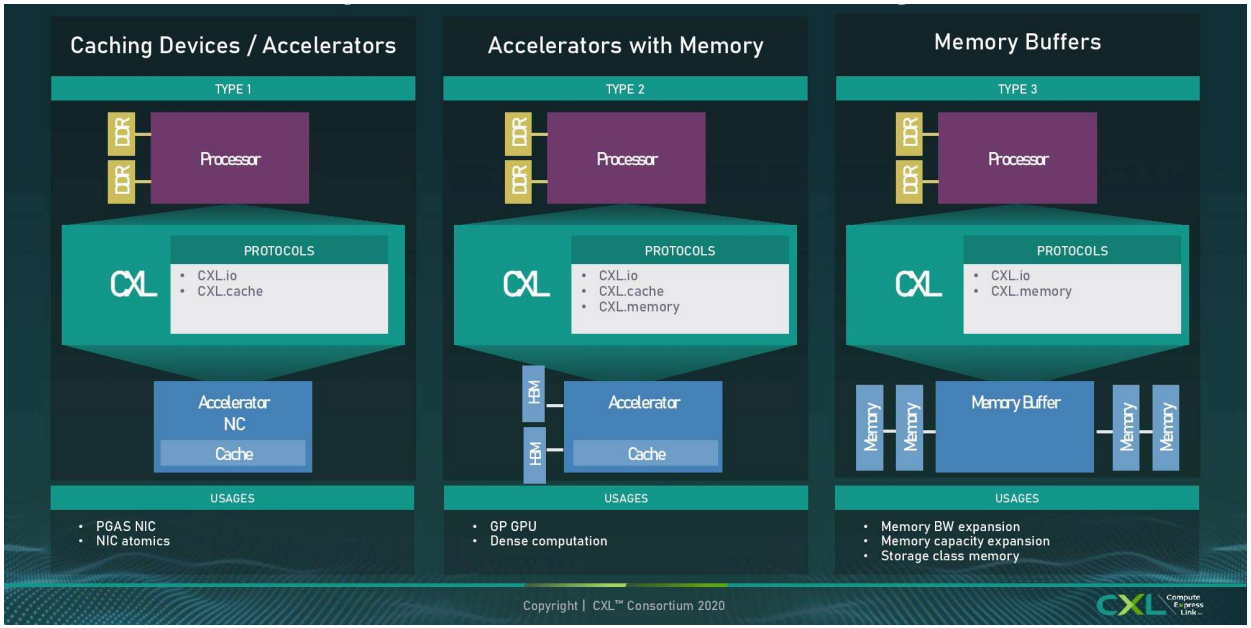

Representative CXL Usage*

NVSL

*https://www.anandtech.com/show/16227/compute-express-link-cxl-2-0-switching-pmem-security

Compute Express Link (CXL) is Here

- Standard for high-speed communication
- Three sub-protocols
- Provides cache-coherent memory access
  - CPU to device
  - Device to CPU
  - Device to device

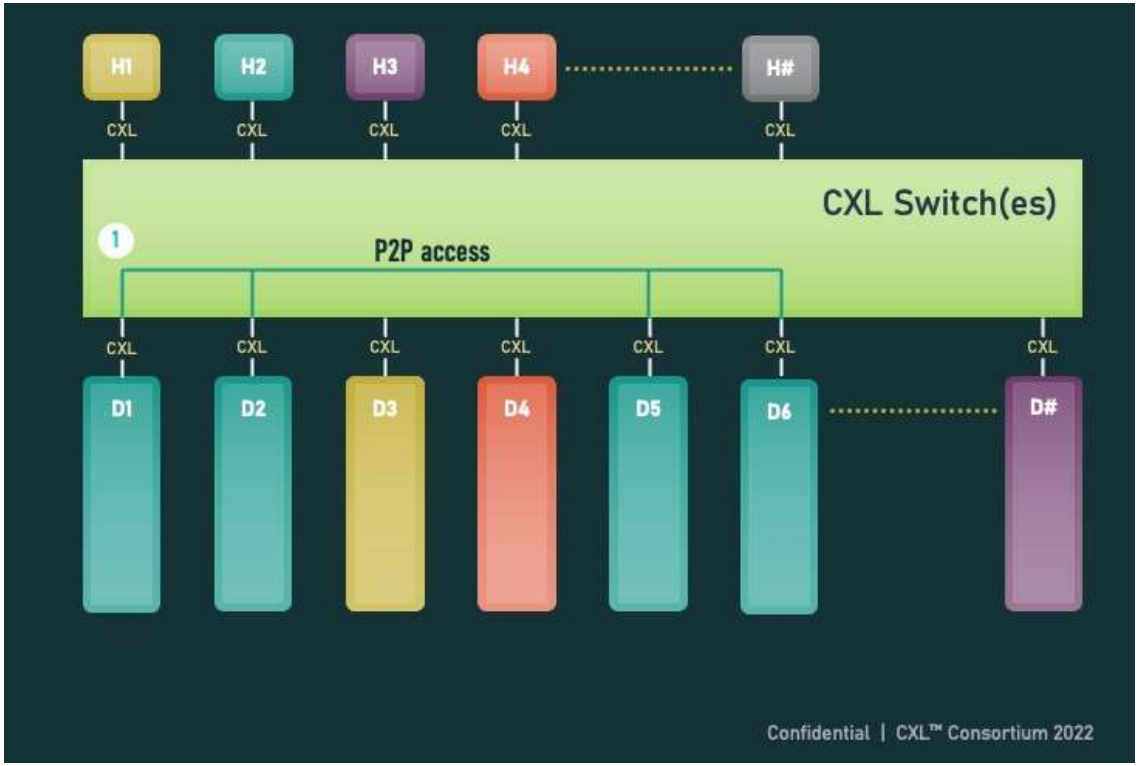

CXL P2P Memory Access*

*https://www.servethehome.com/compute-express-link-cxl-3-0-is-the-exciting-building-block-for-disaggregation/

Device Collaboration Enabled

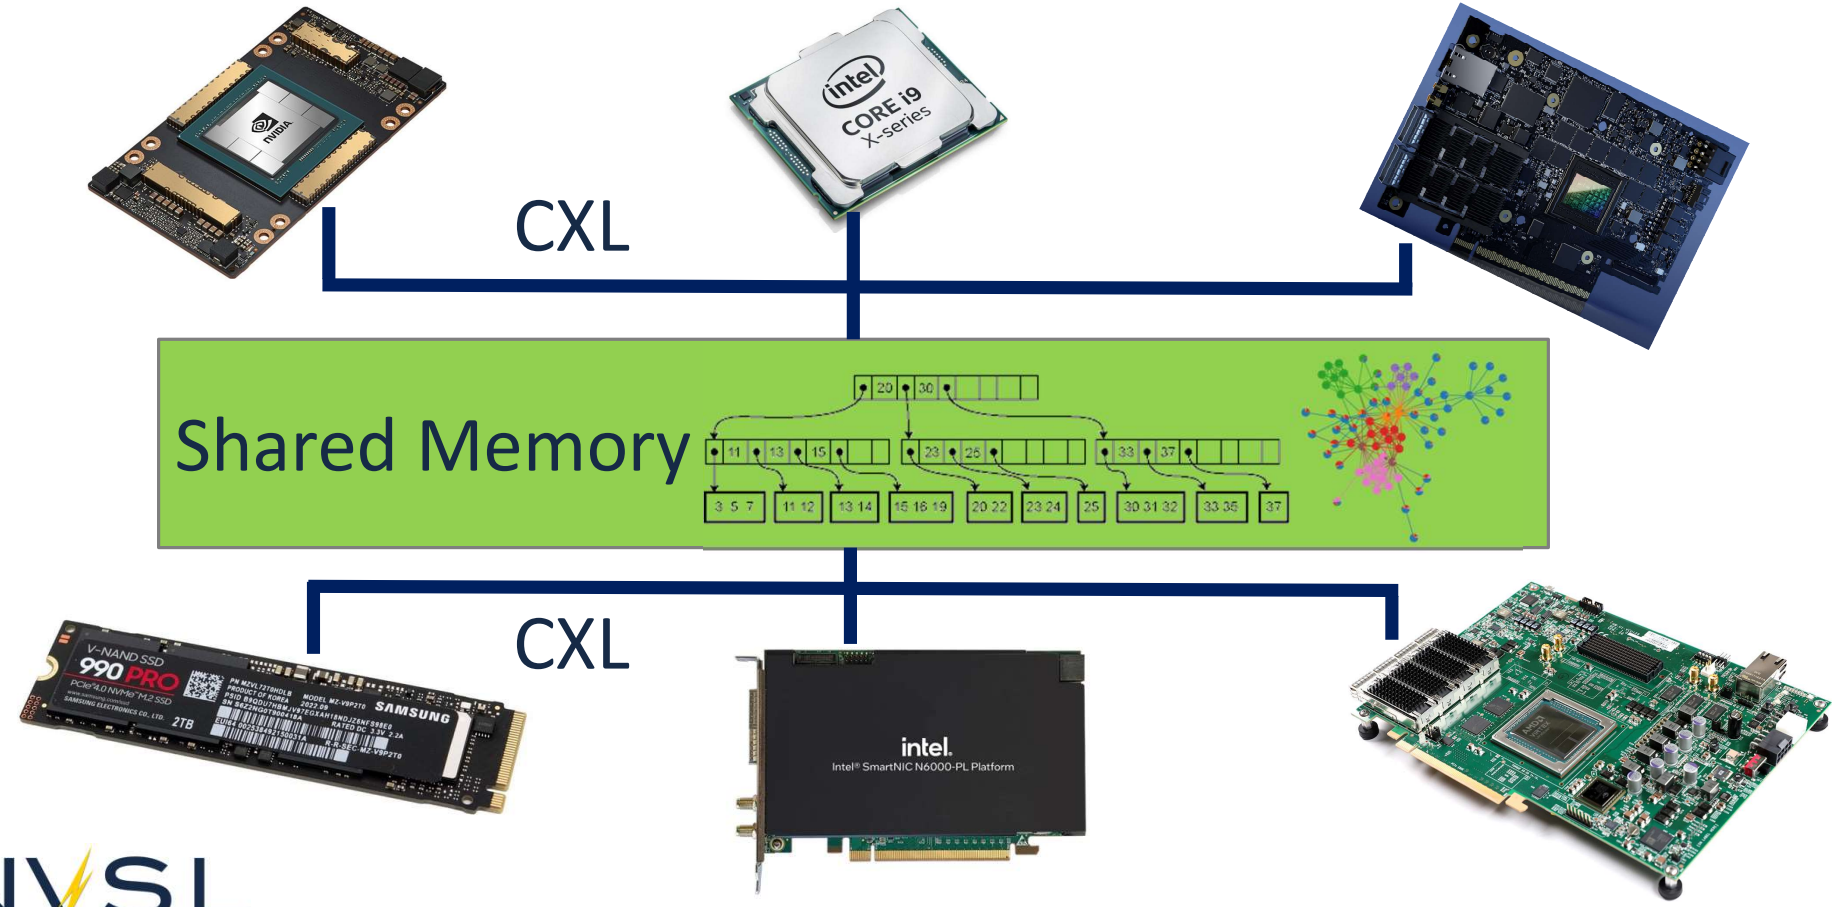

NVSL

But why?

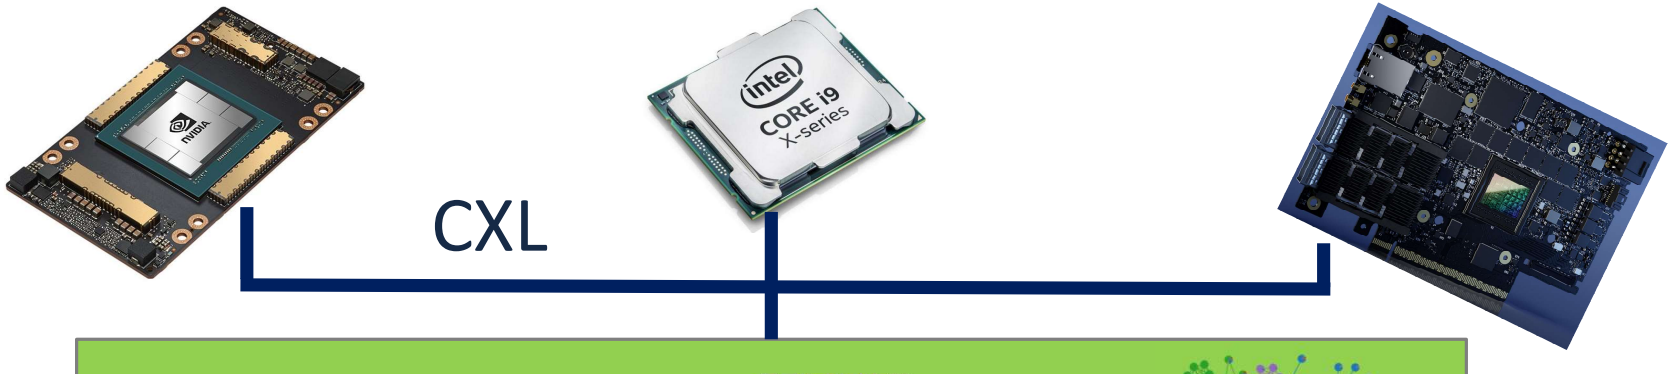

How to build systems to exploit new capability?

NVSL

OS on CPU in Charge Now

- Peripherals are managed as black boxes

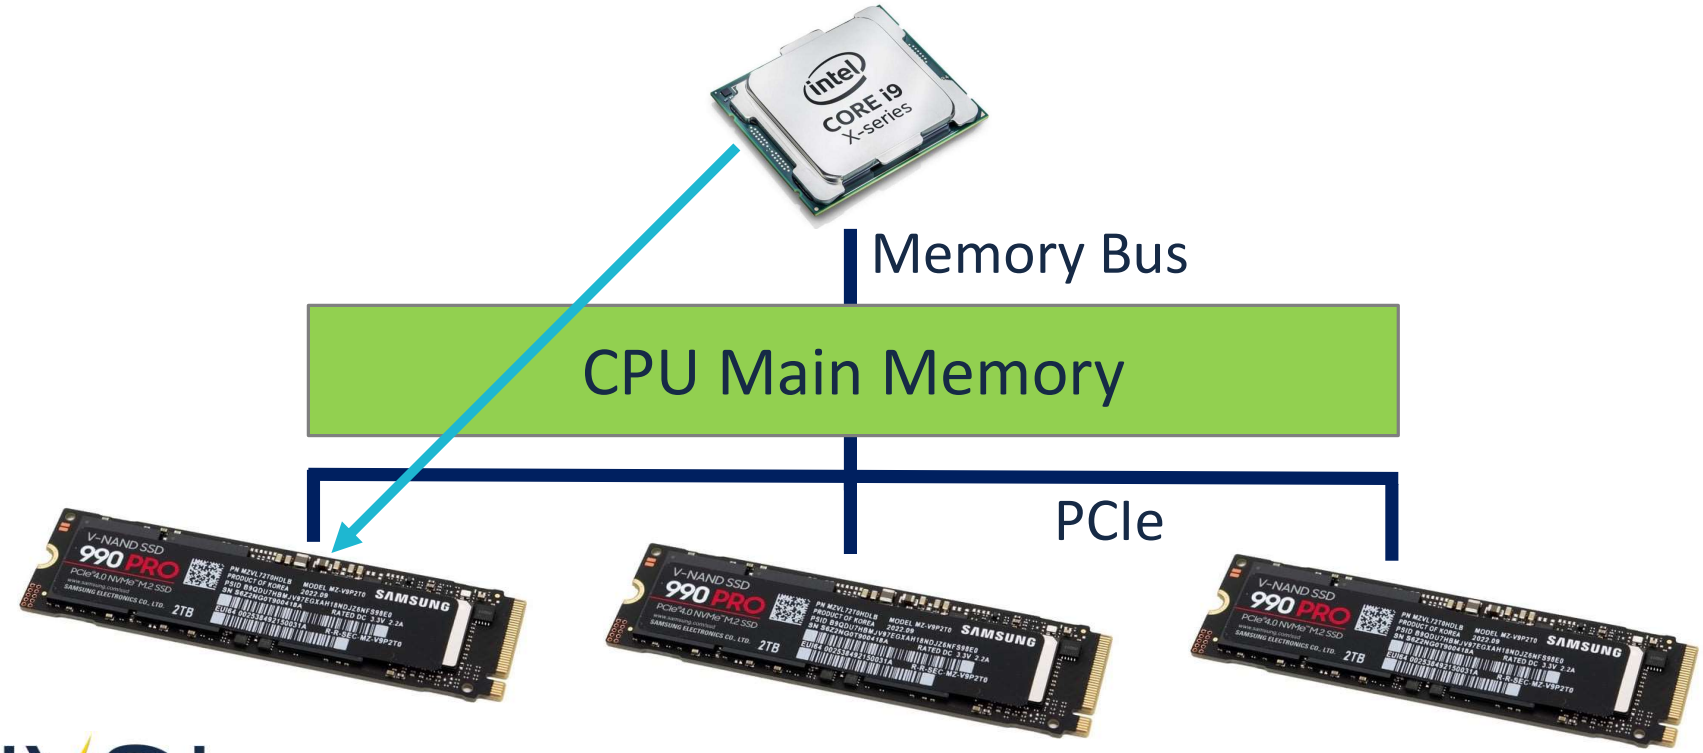

NVSL

OS on CPU in Charge Now

- Peripherals are managed as black boxes

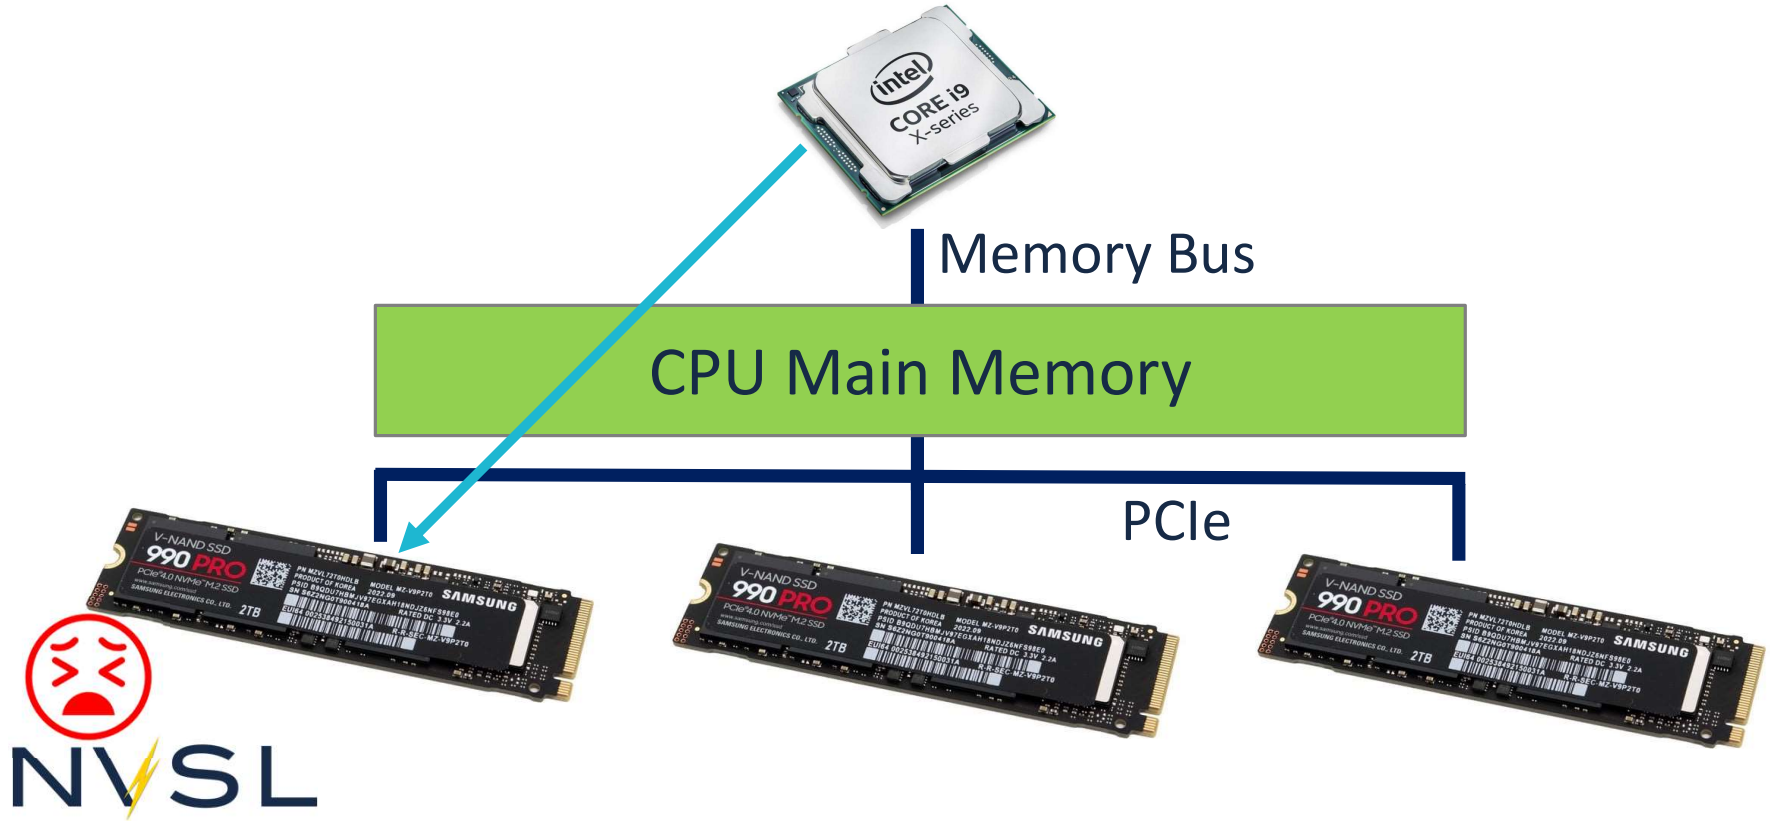

OS on CPU in Charge Now

- Peripherals are managed as black boxes

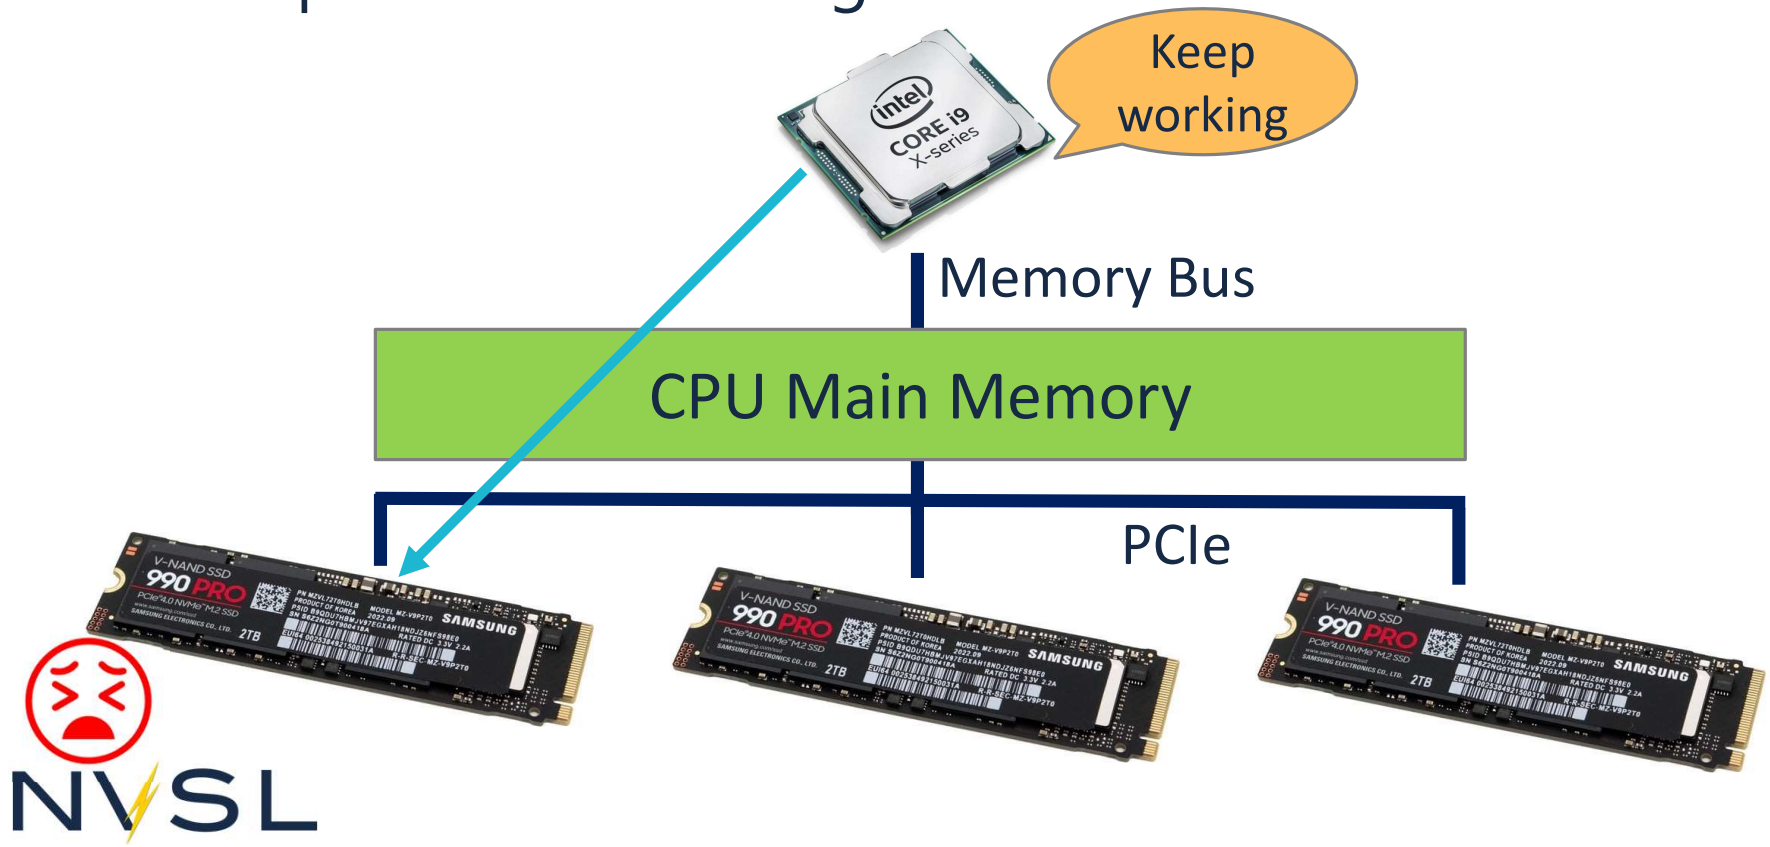

SSD Tail Latency Problem

- SSD's performance jitters
  - Garbage collection
  - Wear leveling
- SSD RAID's tail latency is worse
  - One operation touches multiple SSDs
  - Any SSD's spike latency affects the whole operation

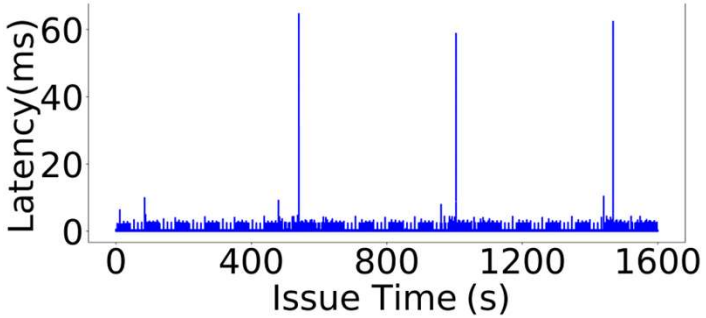

Sample spike behavior in write tests on an SSD (write request size at 64KB).
Source: Figure 4(a) from FusionRAID

SSDs	Median latency (ms)	Avg. latency (ms)	P99 latency (ms)	P999 latency (ms)
SSD 0	0.049	0.68	0.42	24.4
SSD 1	0.049	1.26	0.46	702.24
SSD 2	0.05	0.63	0.39	30.16
SSD 3	0.049	1.64	0.53	895.02
SSD 4	0.05	1.71	0.64	827.91

Application Exchange latency, individual aged SSDs within RAID.
Source: Table 2 from FusionRAID

NVSL

Chance for Collaboration

- Devices can collaborate
  - Information sharing
  - Decision making

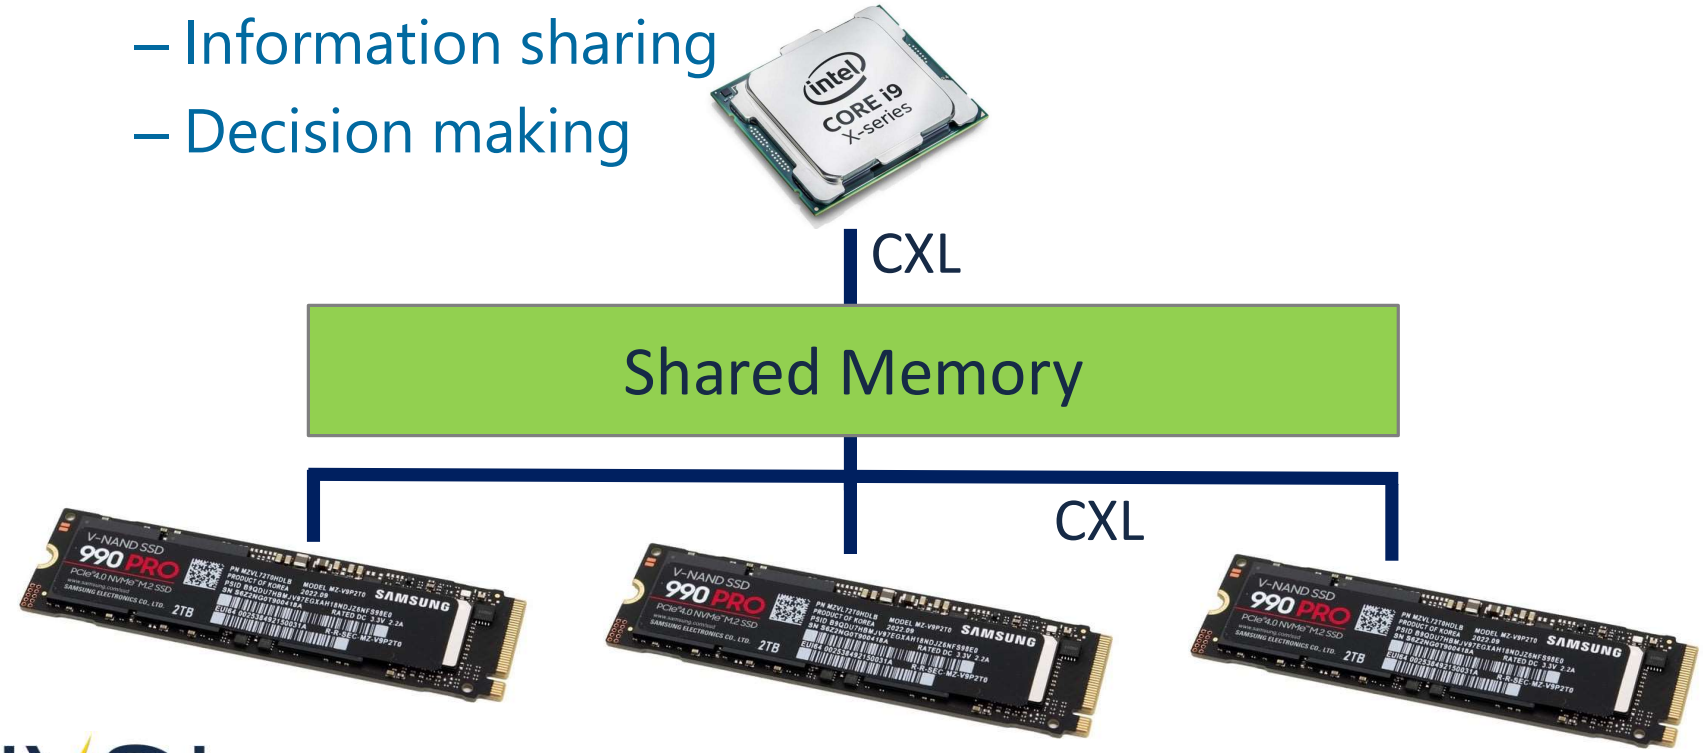

NVSL

Chance for Collaboration

- Devices can collaborate
  - Information sharing
  - Decision making

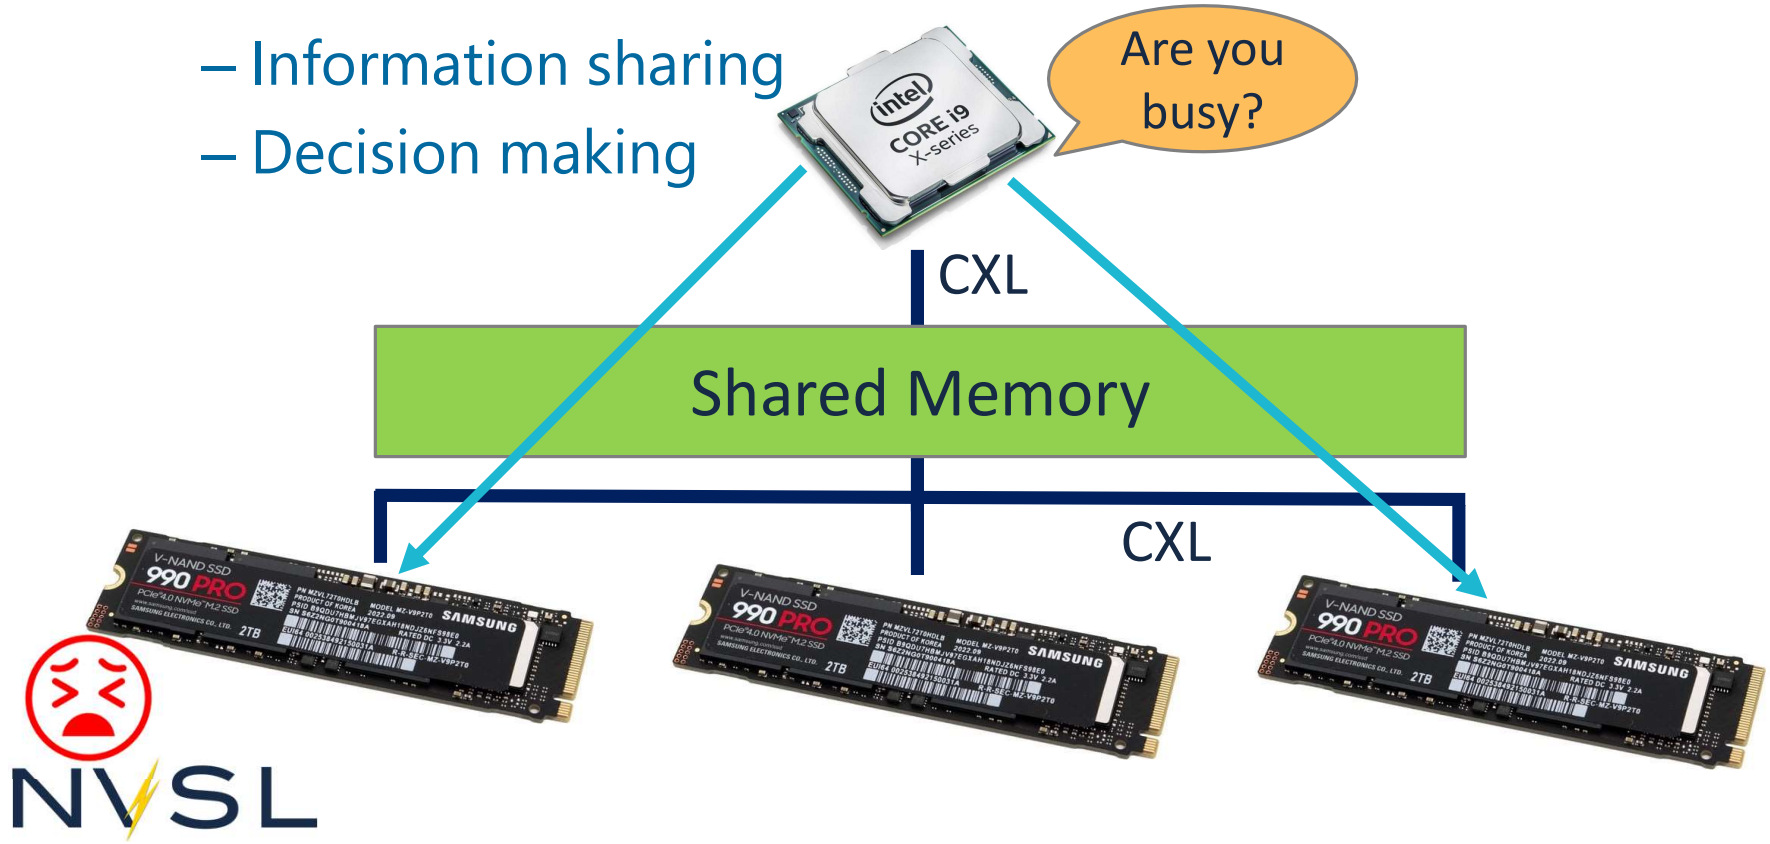

NVSL

Chance for Collaboration

- Devices can collaborate
  - Information sharing
  - Decision making

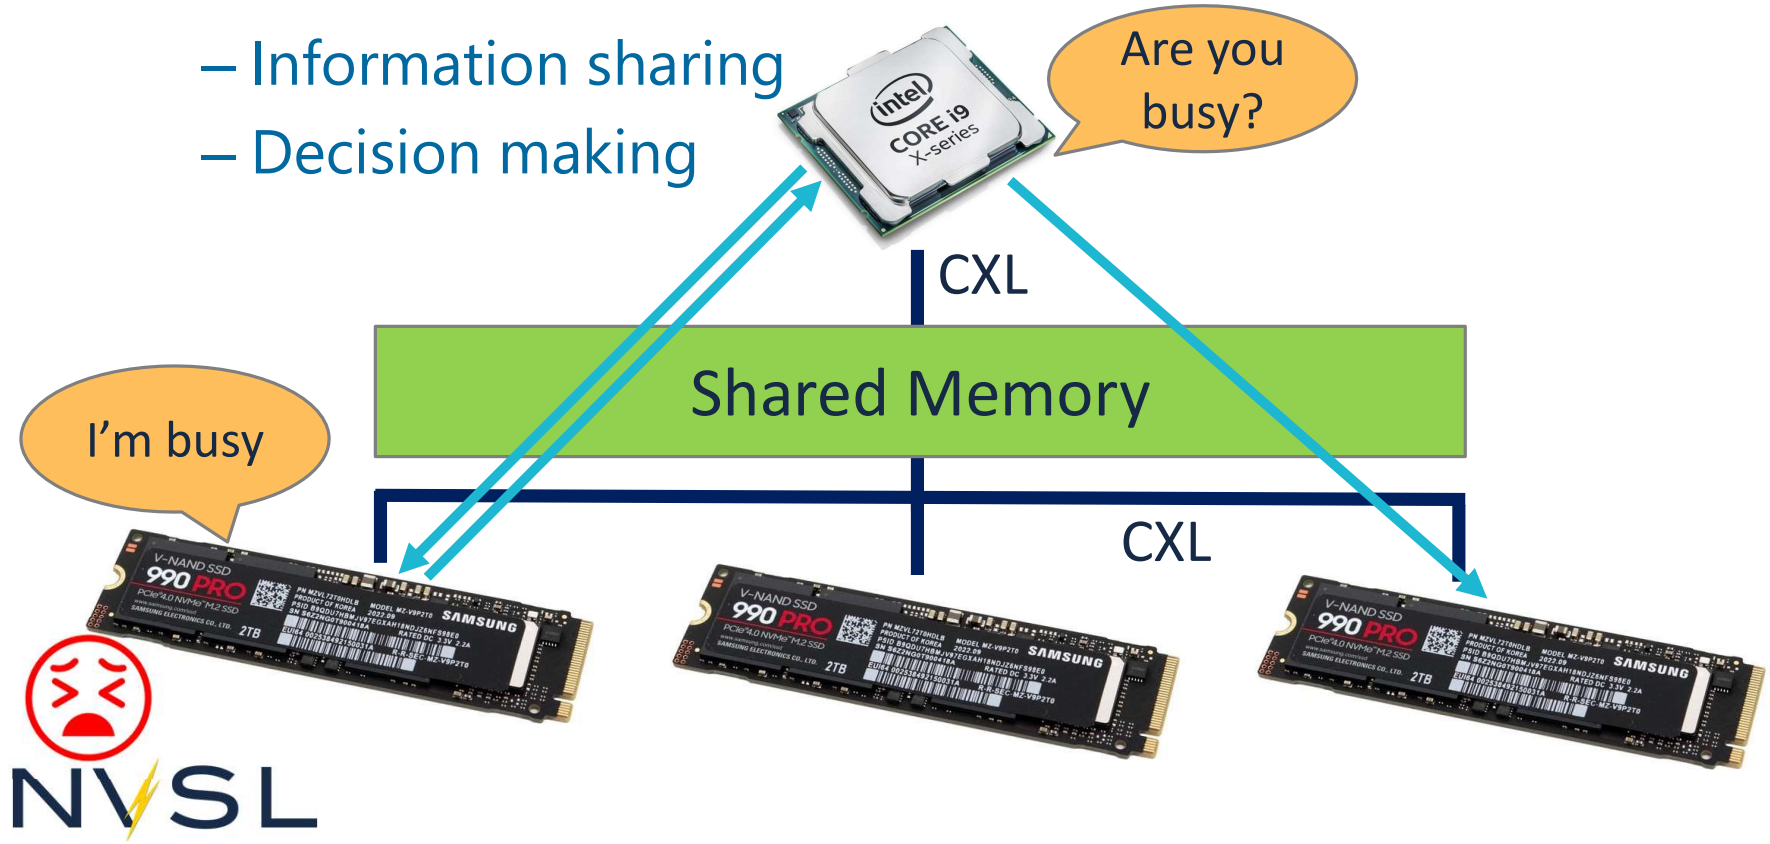

NVSL

Chance for Collaboration

- Devices can collaborate
  - Information sharing
  - Decision making

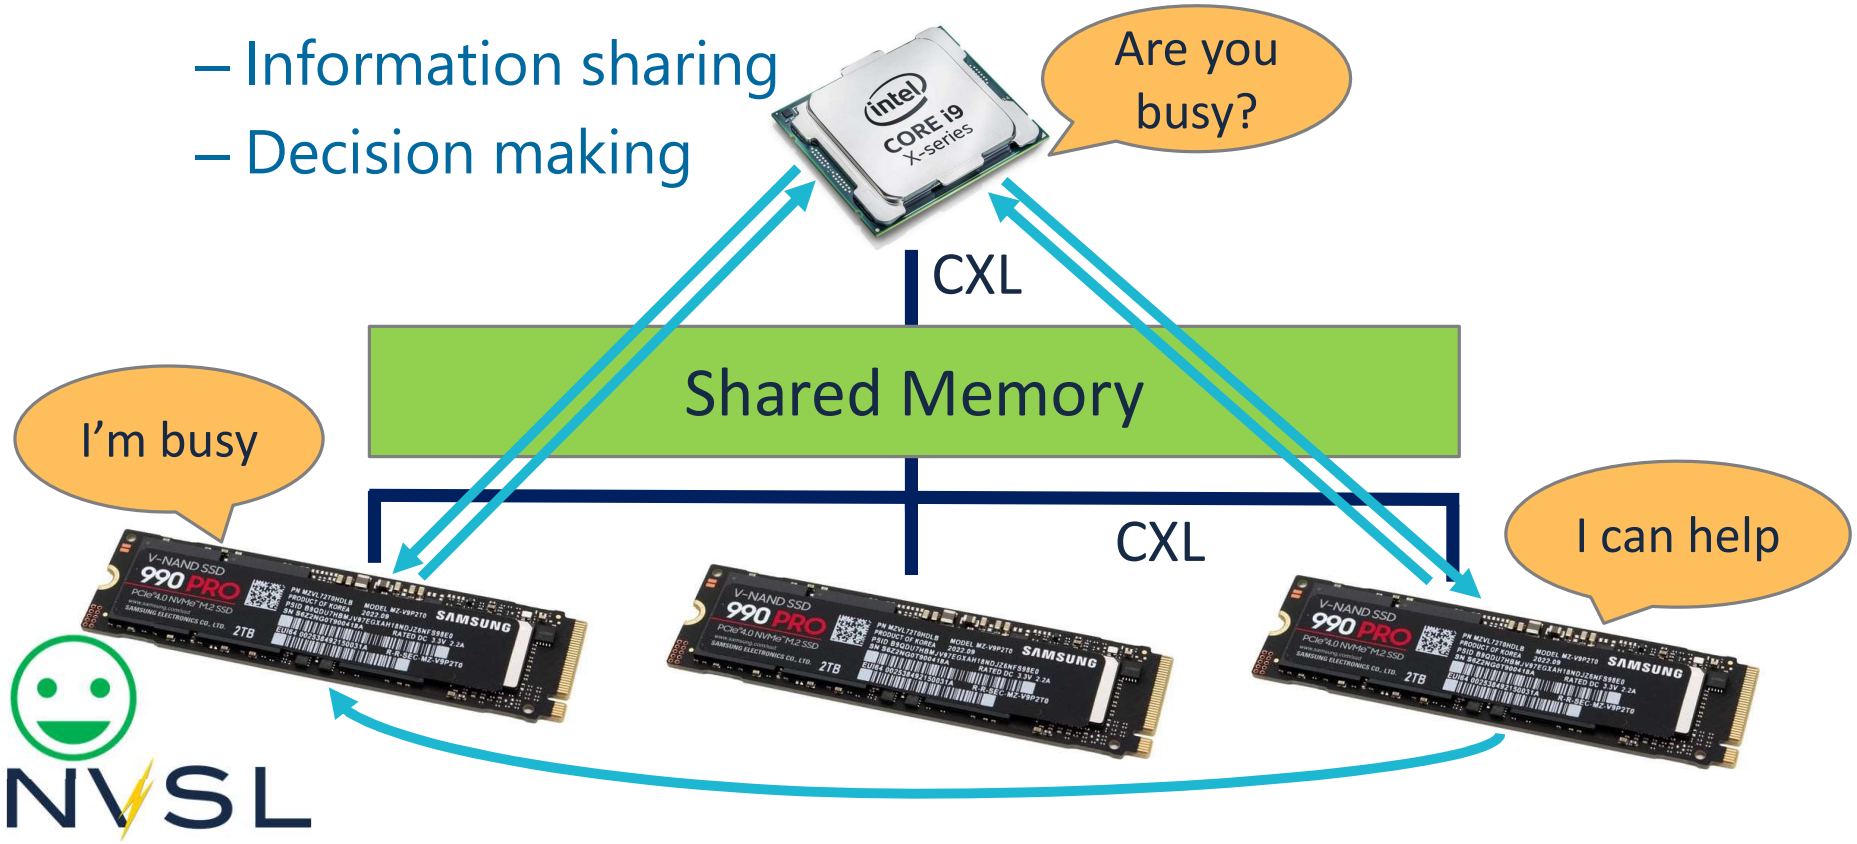

NVSL

Straightforward Solution

Write operation: redirect

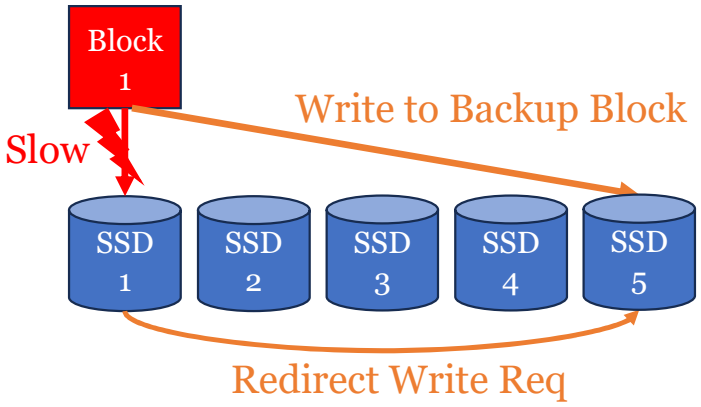

Read operation: redirect

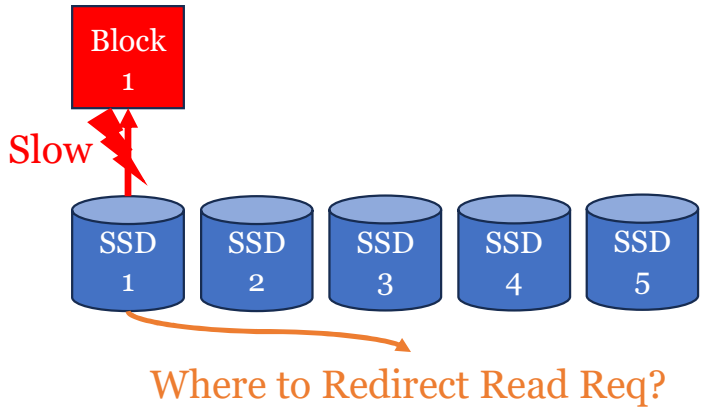

NVSL

Outline

- Background and Motivation
- Problems
- Solutions
- Evaluation & Conclusion

NVSL

Problems

- Replication Overhead
- Block Tracking
- CPU Centric
- Redirection Performance

NVSL

Problem 1: Replication Overhead

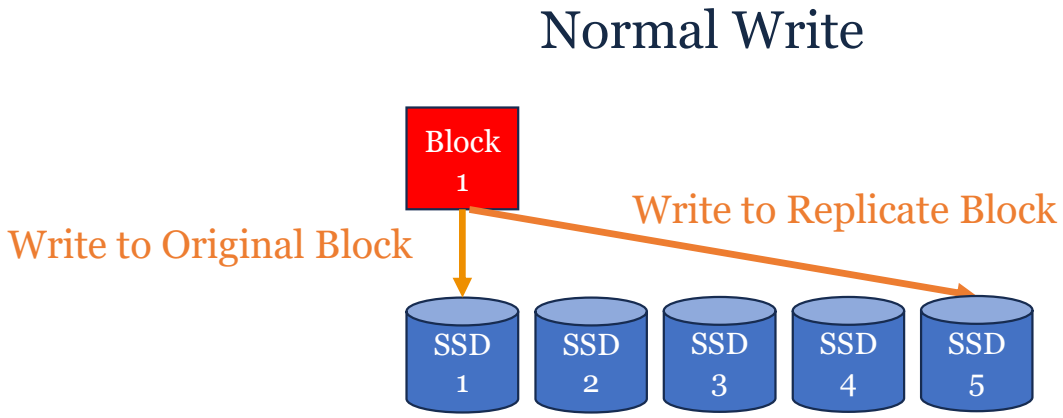

NVSL

Problem 1: Replication Overhead

- Hurt normal write performance
- Duplication wastes space

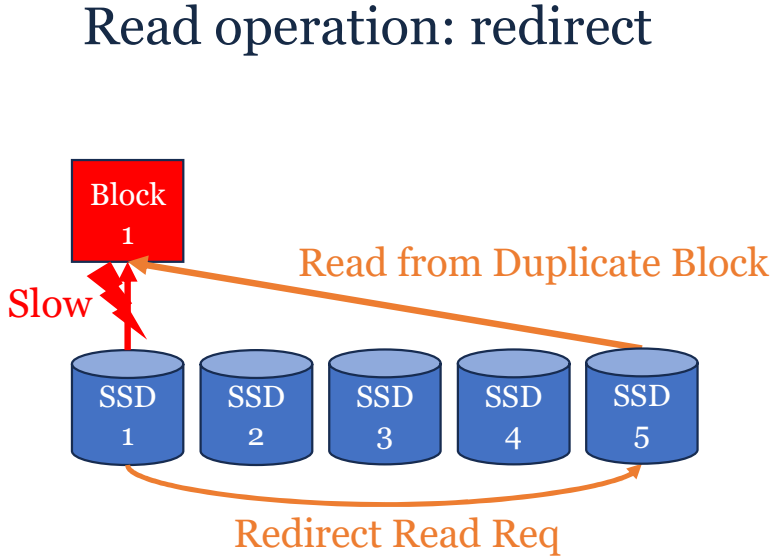

NVSL

Problem 2: Block Tracking

Write operation: redirect

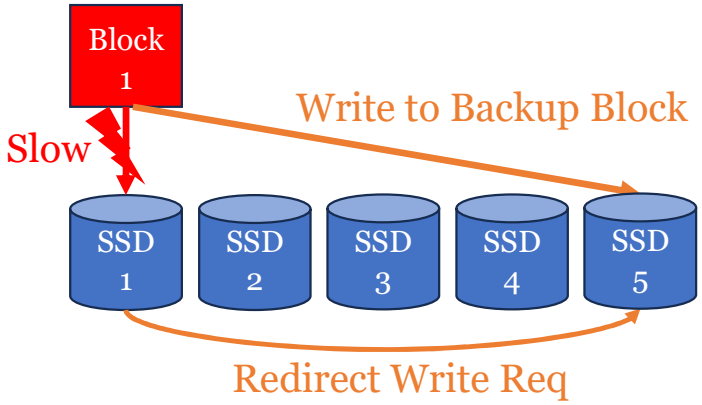

NVSL

Problem 2: Block Tracking

- Memory footprint
- Concurrent map update

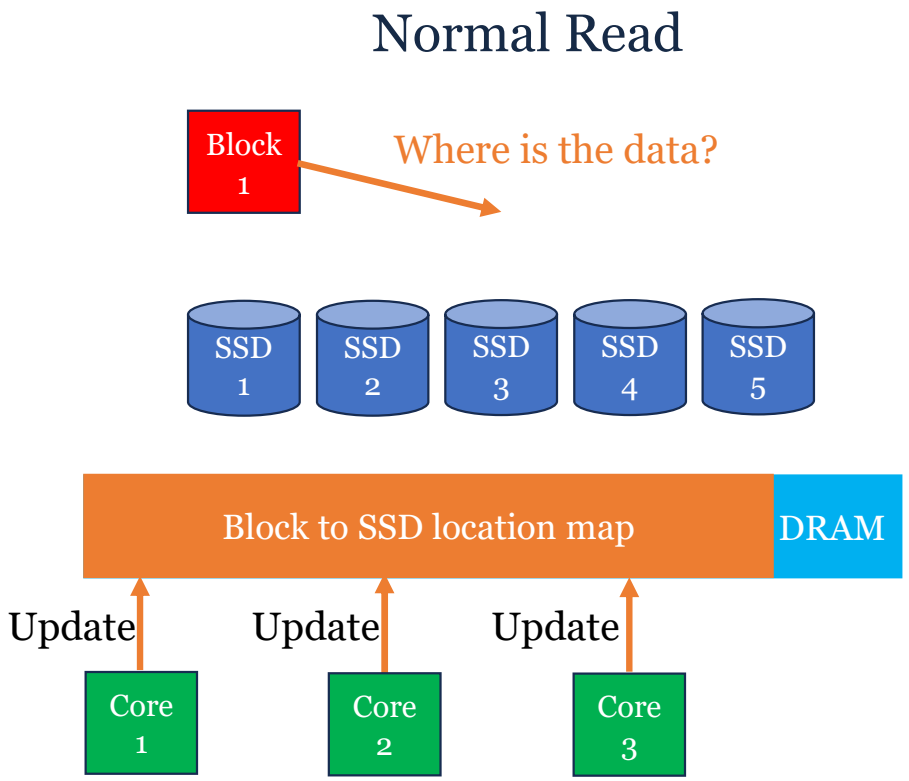

NVSL

Problem 3: CPU Centric

- Interference with normal work
  - Monitoring
  - Redirection
- SSD's internal state unknown
  - Late redirecting
  - Late resuming

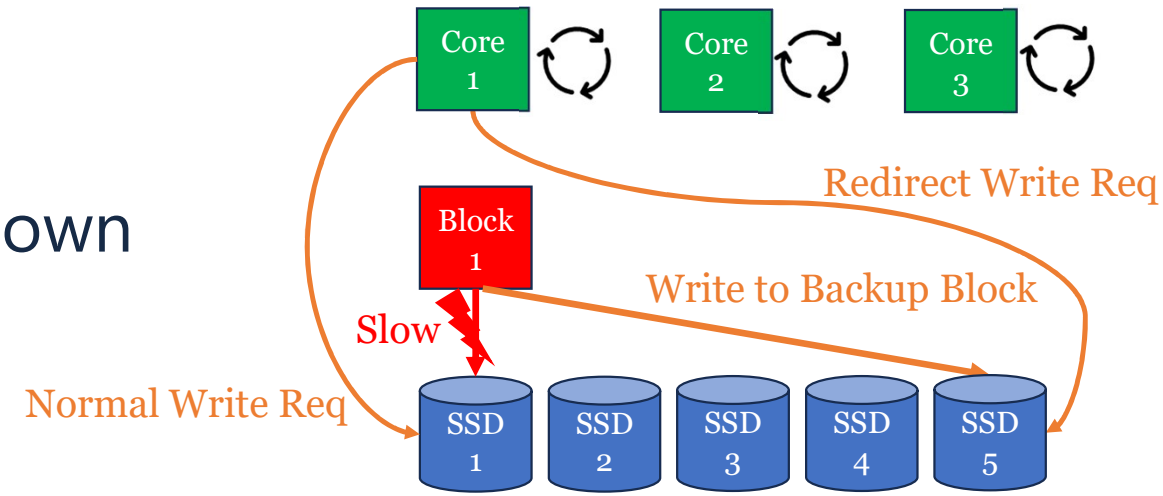

NVSL

Problem 3: CPU Centric

- Interference with normal work
  - Monitoring
  - Redirection
- SSD's internal state unknown
  - Late redirecting
  - Late resuming
- Not scalable

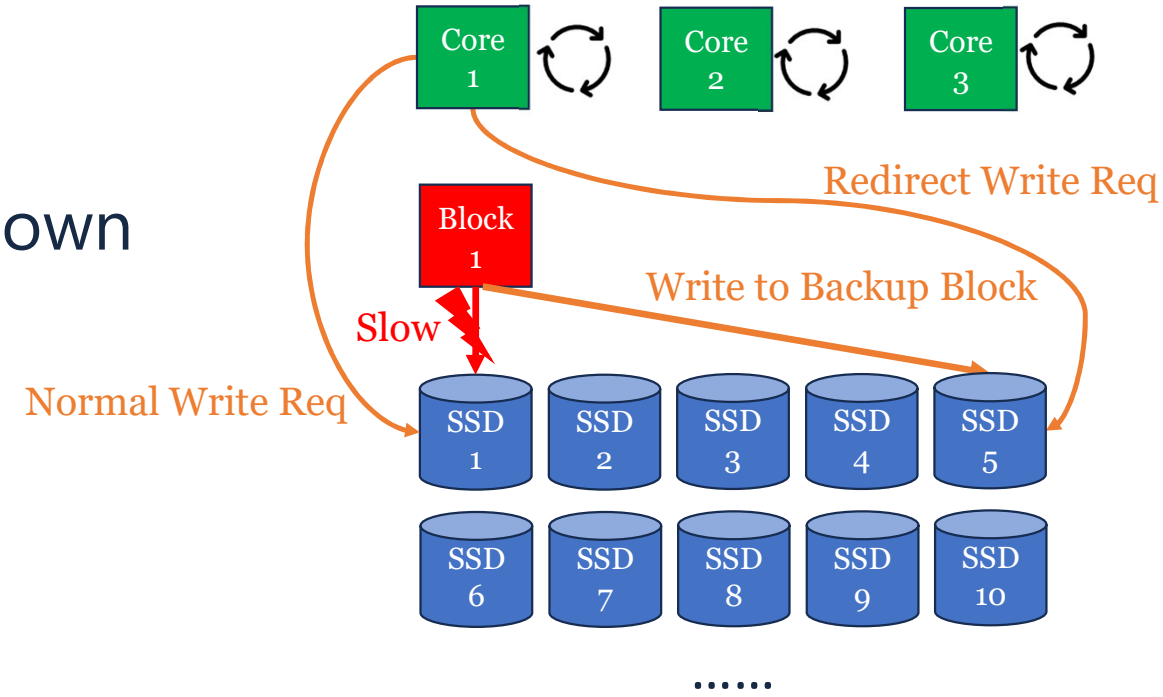

Problem 4: Redirection Performance

- CPU: duplicate request
  - Shouldn't remove requests
  - Waste bandwidth

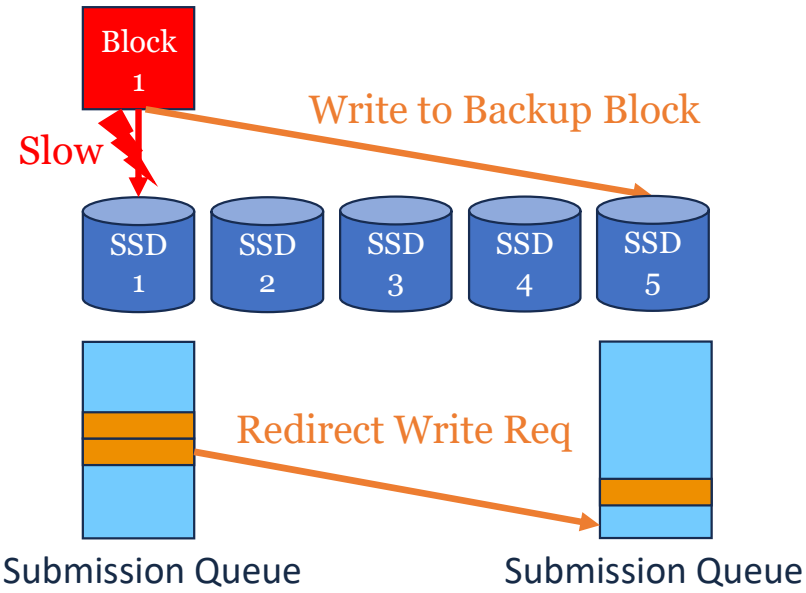

NVSL

Problem 4: Redirection Performance

- CPU: duplicate request
  - Shouldn't remove requests
  - Waste bandwidth
- SSD: DMA not suitable
  - Initialization and setup overhead
  - Interrupt notification

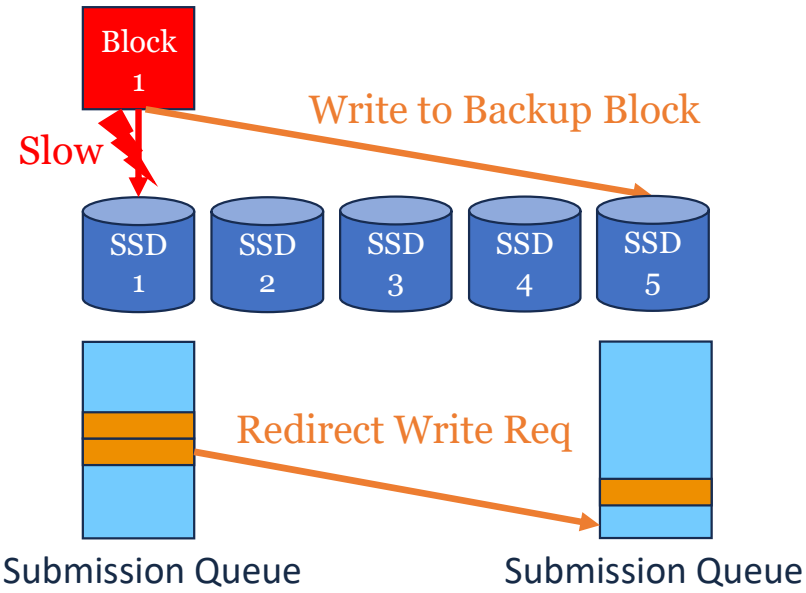

NVSL

Outline

- Background and Motivation
- Problems
- Solutions
- Evaluation & Conclusion

NVSL

CXL-based SSD Autonomic and Collaborative Scheduling

- CXL-based
  - High-performance communication

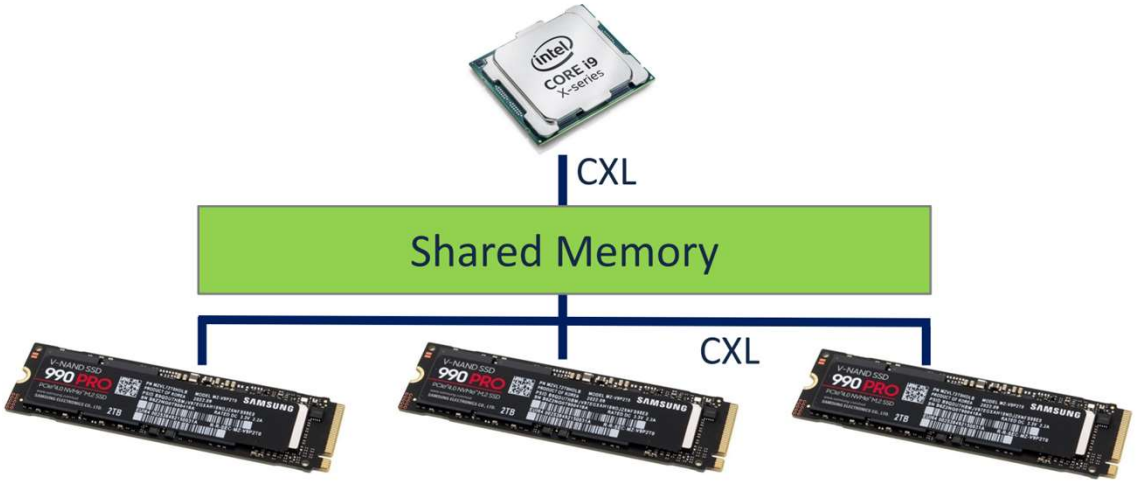

NVSL

CXL-based SSD Autonomic and Collaborative Scheduling

- CXL-based
  - High-performance communication
- SSD autonomic
  - Stop polling when busy
  - Redirect / rebuild

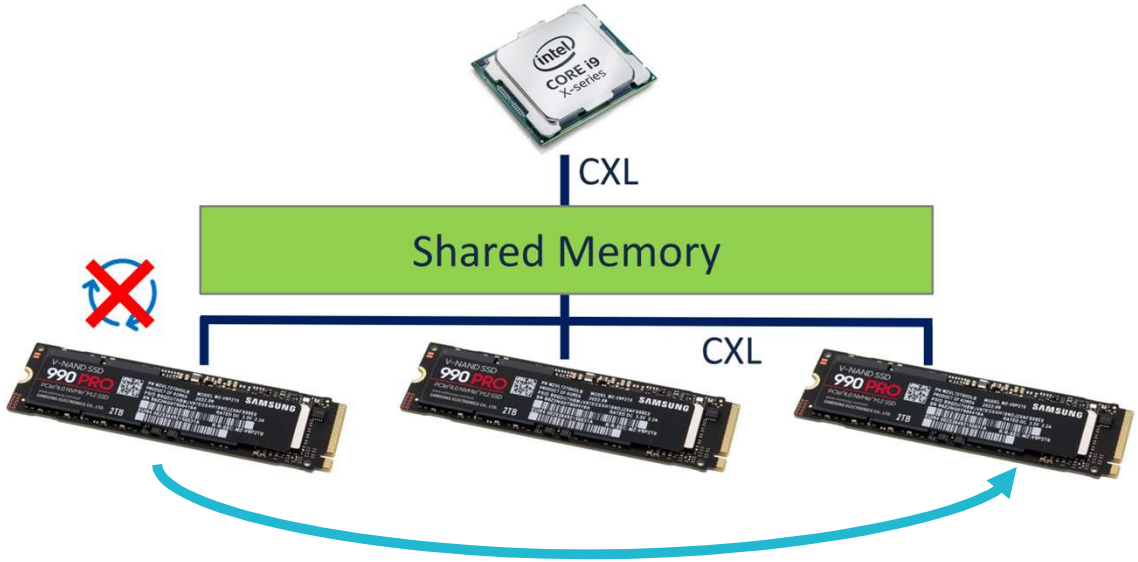

NVSL

CXL-based SSD Autonomic and Collaborative Scheduling

- CXL-based
  - High-performance communication
- SSD autonomic
  - Stop polling when busy
  - Redirect / rebuild
- Share status
  - Flag busy to stop requests
  - No redirection when backup busy

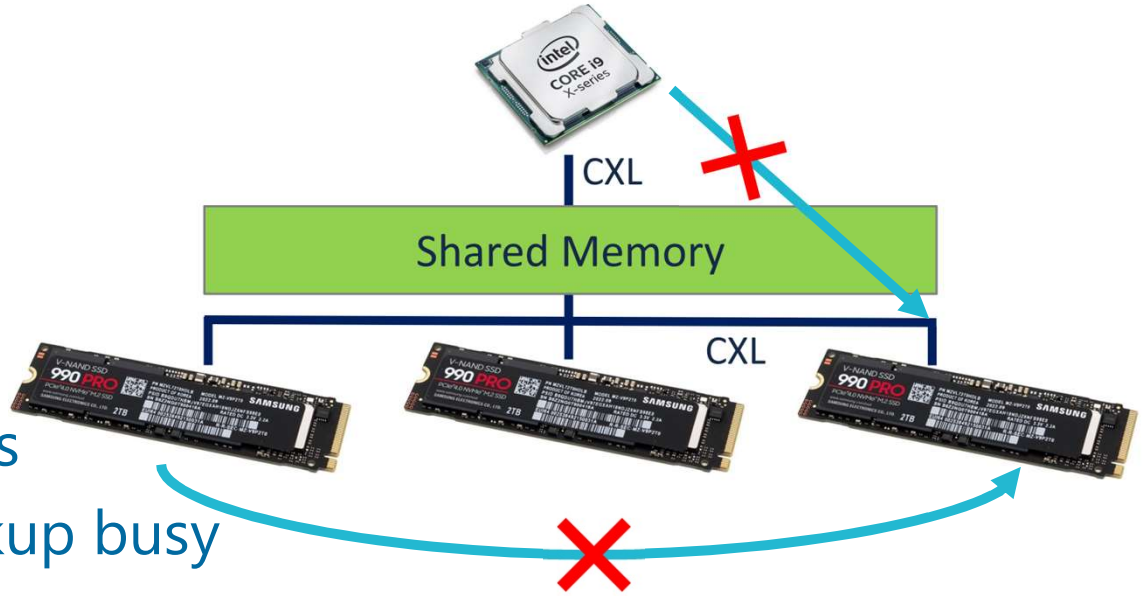

NVSL

Solutions

- Replication Overhead -> RAID Rebuild
- Block Tracking -> Stack + Dynamic Block Mapping
- CPU Centric -> CPU + SSD Collaboration
- Redirection Performance -> CXL Based Data Sharing

NVSL

Solutions

- Replication Overhead -> RAID Rebuild
- Block Tracking -> Stack + Dynamic Block Mapping
- CPU Centric -> CPU + SSD Collaboration
- Redirection Performance -> CXL Based Data Sharing

NVSL

Using RAID Rebuild for Read

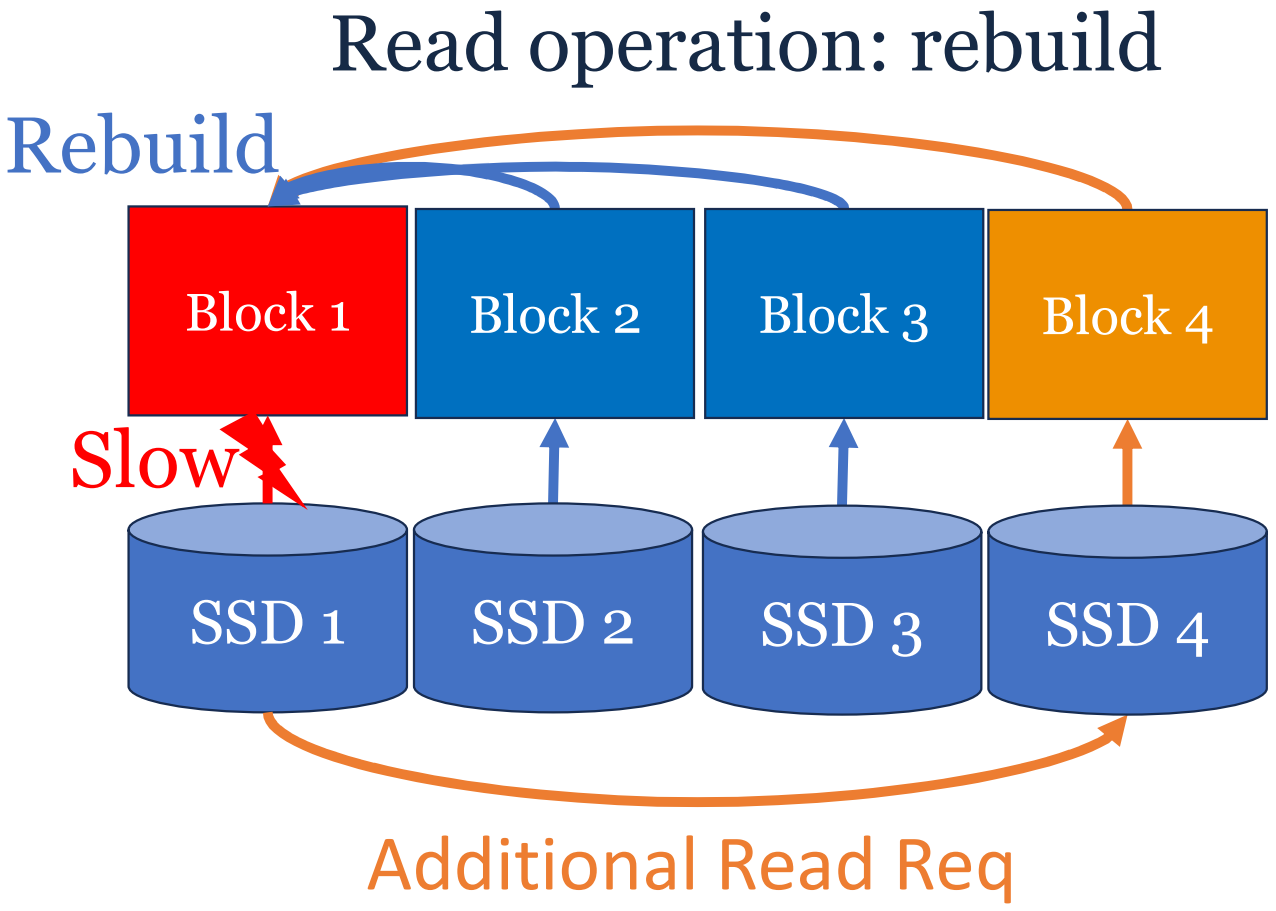

NVSL

Solutions

- Replication Overhead -> RAID Rebuild
- Block Tracking -> Stack + Dynamic Block Mapping
- CPU Centric -> CPU + SSD Collaboration
- Redirection Performance -> CXL Based Data Sharing

NVSL

Static + Dynamic Block Mapping

- Deterministic block mapping calculation
- N (N <= stripe size) backup blocks for one stripe

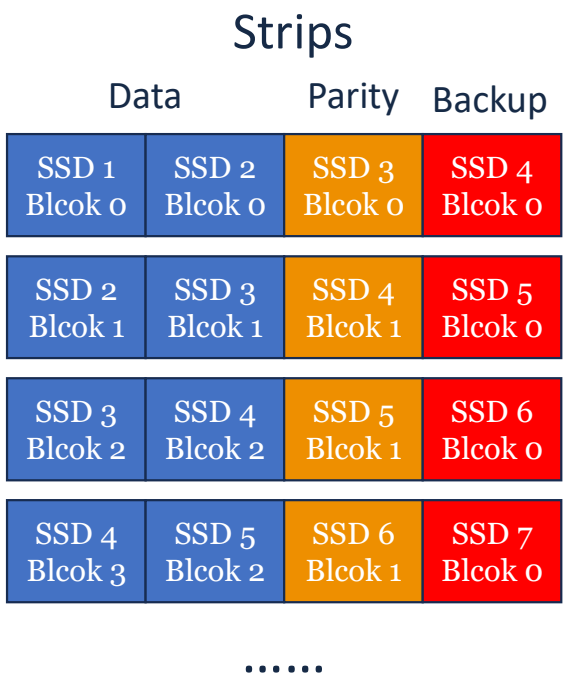

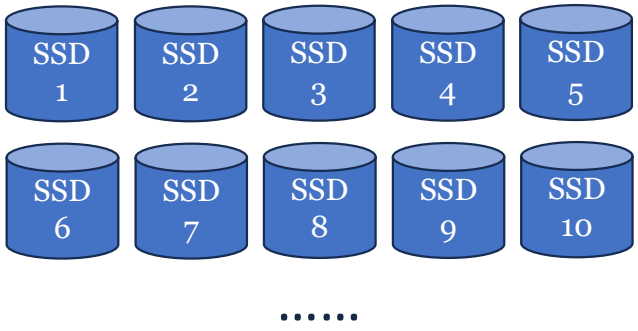

NVSL

Static + Dynamic Block Mapping

- Deterministic block mapping calculation
- N (N <= stripe size) backup blocks for one stripe
- Dynamically tracking redirected data

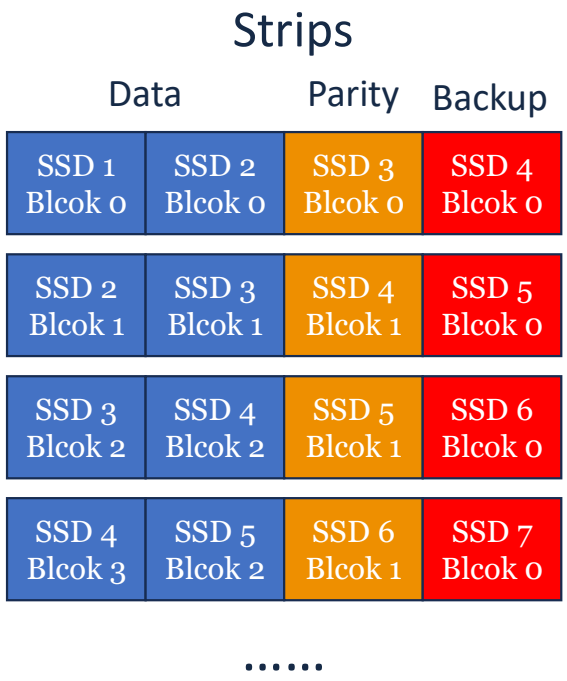

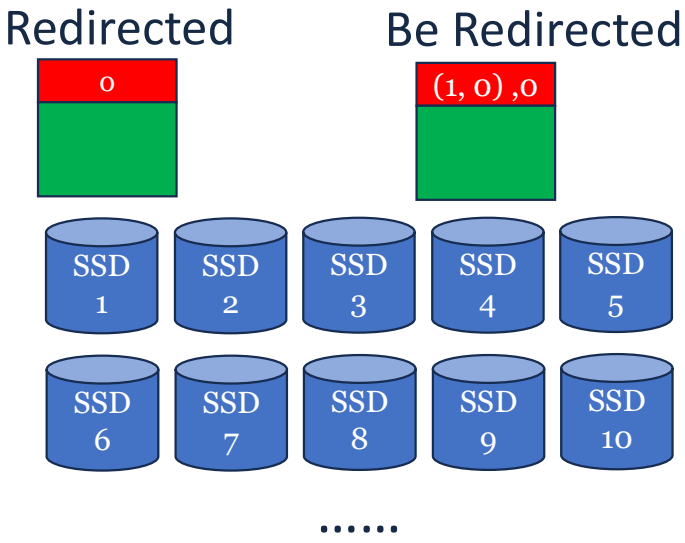

Solutions

- Replication Overhead -> RAID Rebuild
- Block Tracking -> Stack + Dynamic Block Mapping
- CPU Centric -> CPU + SSD Collaboration
- Redirection Performance -> CXL Based Data Sharing

NVSL

CPU + SSD Collaboration

- SSD autonomic

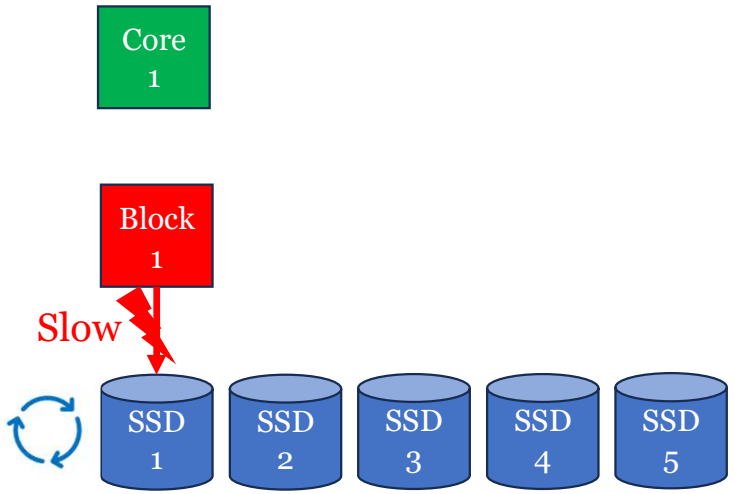

NVSL

CPU + SSD Collaboration

- SSD autonomic
  - Stop polling when busy
  - Redirect / rebuild

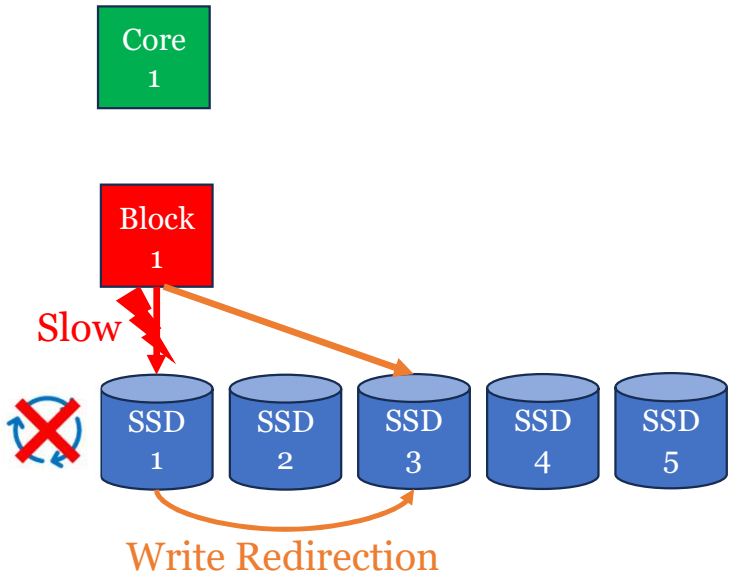

NVSL

CPU + SSD Collaboration

- SSD autonomic
  - Stop polling when busy
  - Redirect / rebuild
- Share status

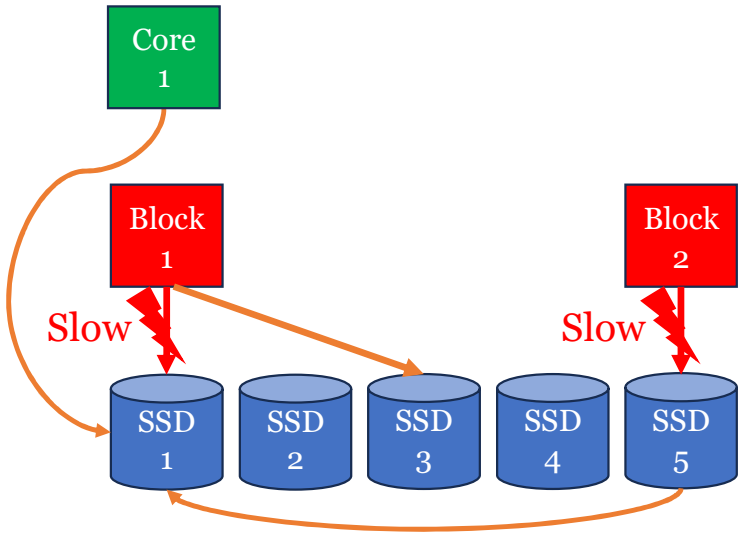

NVSL

CPU + SSD Collaboration

- SSD autonomic
  - Stop polling when busy
  - Redirect / rebuild
- Share status
  - Flag busy to stop requests
  - No redirection when backup busy

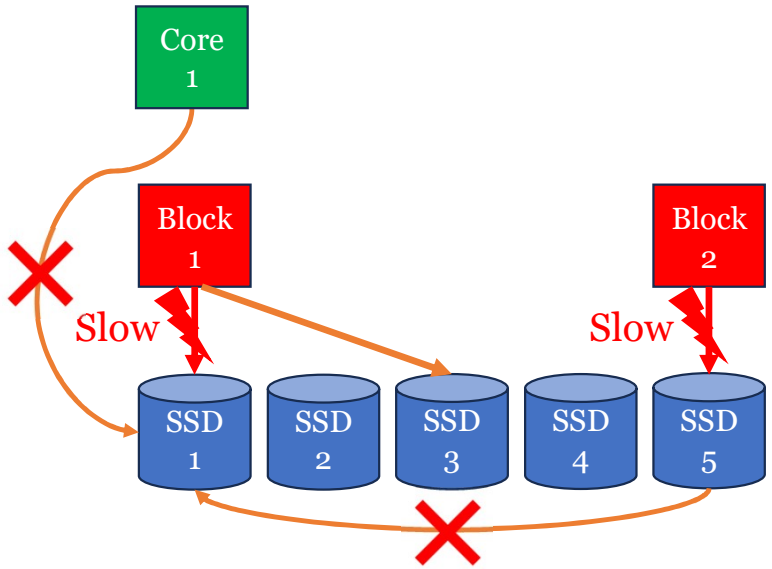

NVSL

Solutions

- Replication Overhead -> RAID Rebuild
- Block Tracking -> Stack + Dynamic Block Mapping
- CPU Centric -> CPU + SSD Collaboration
- Redirection Performance -> CXL Based Data Sharing

NVSL

CXL Based Data Sharing

- Communication through high-performance shared memory
- Optimized data structures

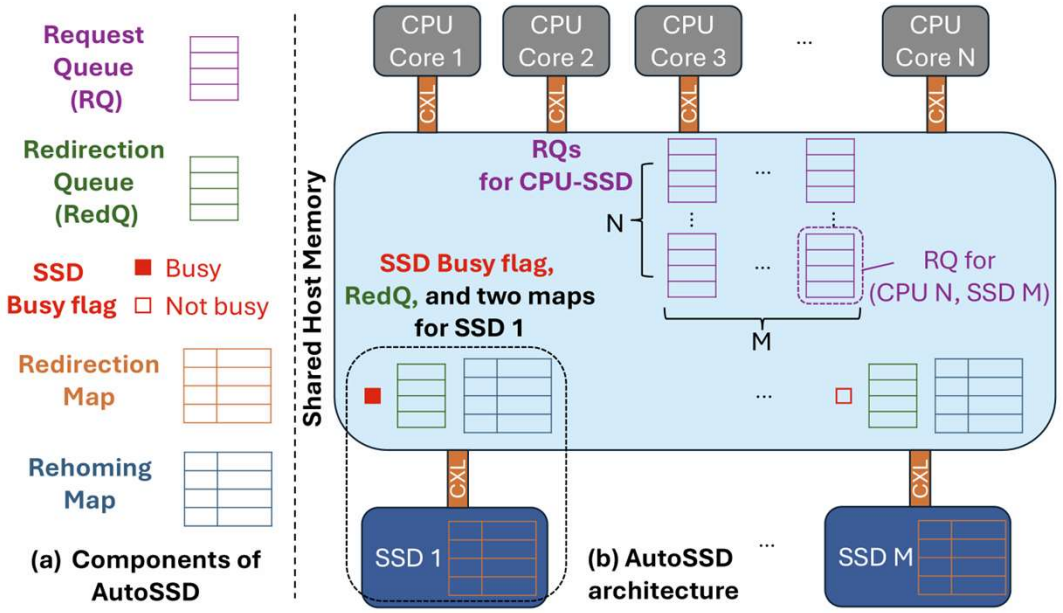

Outline

- Background and Motivation
- Problems
- Solutions
- Evaluation & Conclusion

NVSL

Evaluation

- Dual socket machine
- Remote socket as CXL memory

NVSL

Evaluation

- Dual socket machine
- Remote socket as CXL memory
- ~15% and ~45% decreased P99 latency for random reads and writes

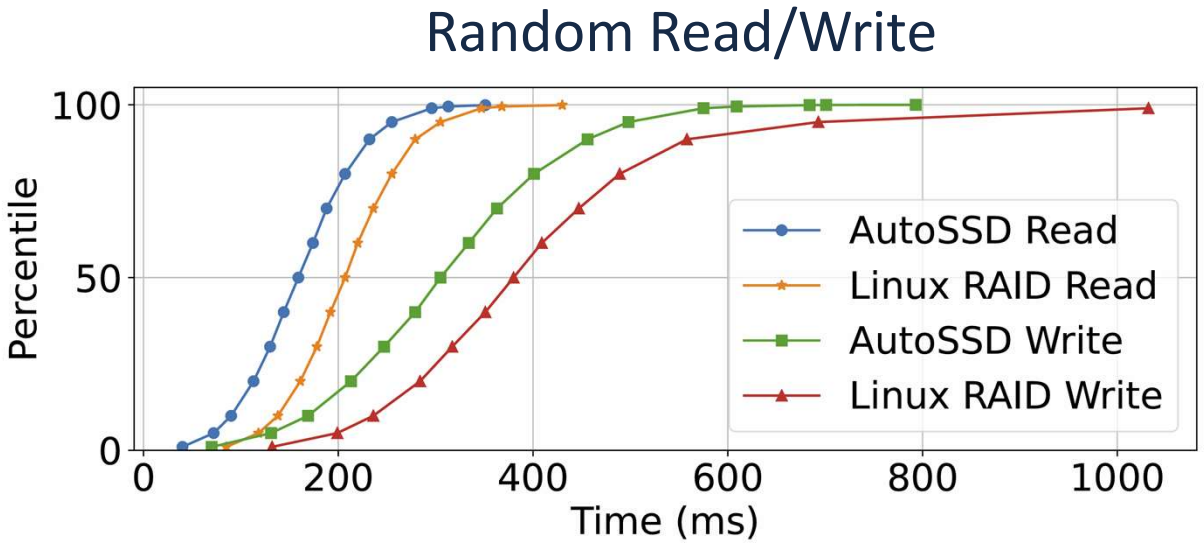

NVSL

Evaluation

- Dual socket machine
- Remote socket as CXL memory
- ~15% and ~45% decreased P99 latency for random reads and writes

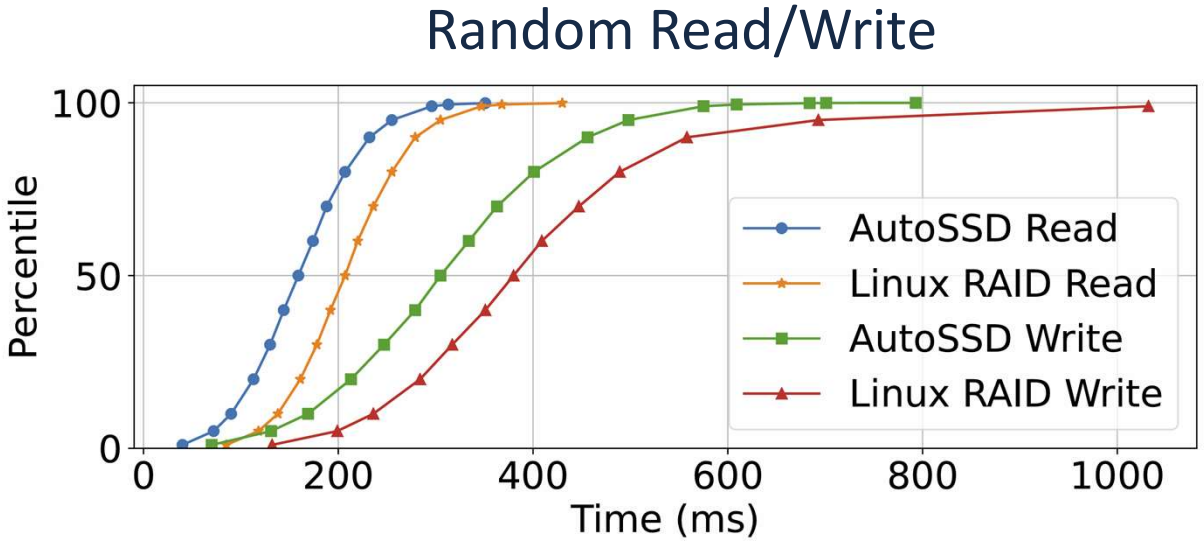

NVSL

Conclusion

- CXL shared memory provides chances for devices' collaboration

NVSL

Conclusion

- CXL shared memory provides chances for devices' collaboration
- SSDs can utilize it to improve their tail latency
  - Internal status sharing
  - Autonomic decision making
  - Collaboration

NVSL

Conclusion

- CXL shared memory provides chances for devices' collaboration
- SSDs can utilize it to improve their tail latency
  - Internal status sharing
  - Autonomic decision making
  - Collaboration
- Other peripherals could also exploit it

NVSL

Questions

Q&A