Ultra-Low Latency SSDs' Impact on Overall Energy Efficiency

Bryan Harris and Nihat Altiparmak
Computer Science & Eng. Department
University of Louisville

HotStorage '20

What is **Ultra-Low Latency** (ULL) data access?

Device	**Latency**	**Performance Gap** *(relative to DRAM)*
HDD	~ 10 ms	100,000×
Traditional SSD	~ 100 µs	1000×
Ultra-Low Latency SSD	**~ 10 µs**	**100×**
DRAM	~ 100 ns	1×

What is **Optane SSD**?

Based on **3D Xpoint** technology by Intel and Micron, using phase change memory (PCM)

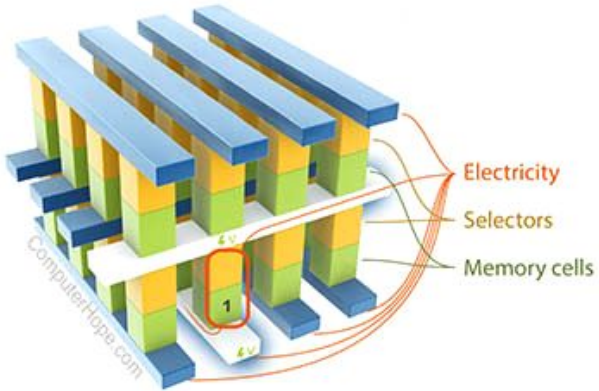

Ideal for applications that require low latency access

Also available as byte addressable DIMM

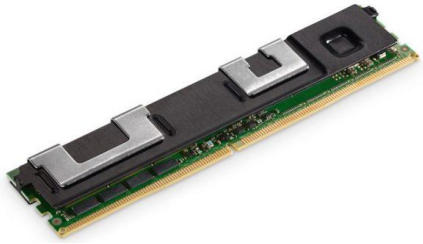

Motivation

Energy characterization

- Existing studies of Optane SSDs primarily focuses on performance
- Energy studies look at only individual device or CPU usage

Impact on overall system software

- ULL disk IO puts increased pressure on system software
- How does this affect power consumption?

Energy Characterization & Impact

Test devices

Technology	Interface	Model	Capacity
Magnetic	SATA 3.1	WD Black 7200 rpm	4 TB
Flash	SATA 3.1	Samsung 850 EVO	1 TB
Flash	PCIe 3	Samsung 960 EVO	500 GB
3D XPoint	PCIe 3	Intel Optane 900P	280 GB

Experimental setup

HOBO plug meter logs power, current, joules, etc., every second.

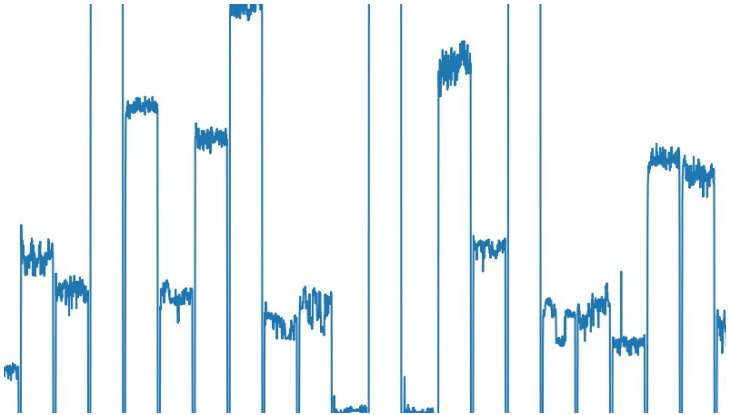

Image from ONSET online catalog
https://www.onsetcomp.com/products/data-loggers/ux120-018/

fio (“Flexible IO tester”) generates storage IO workloads

Storage IO workloads

Larger data requests

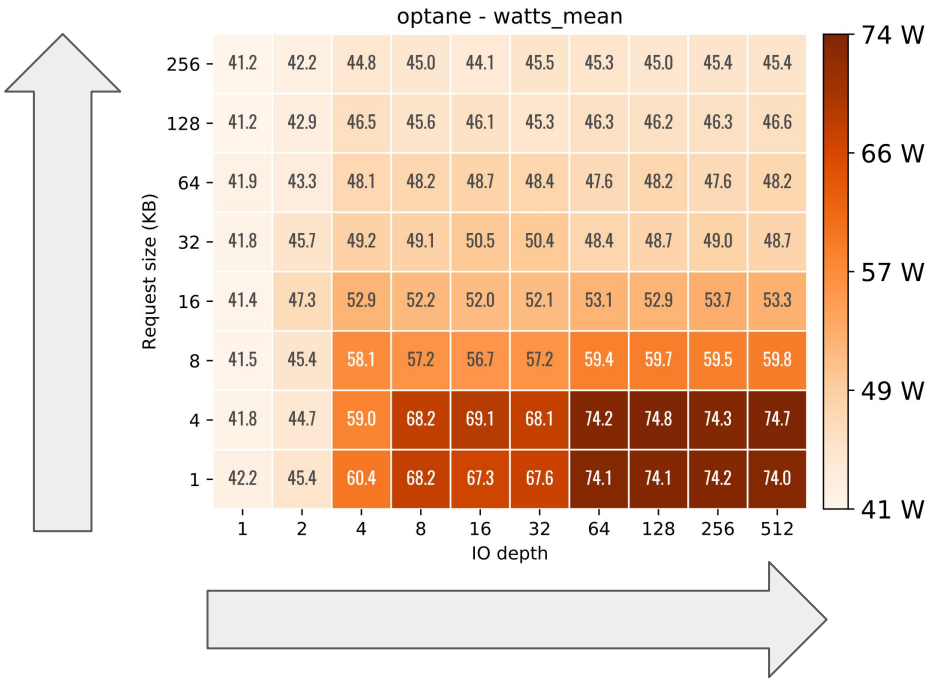

More simultaneous requests

Observations

Idle vs. active power consumption

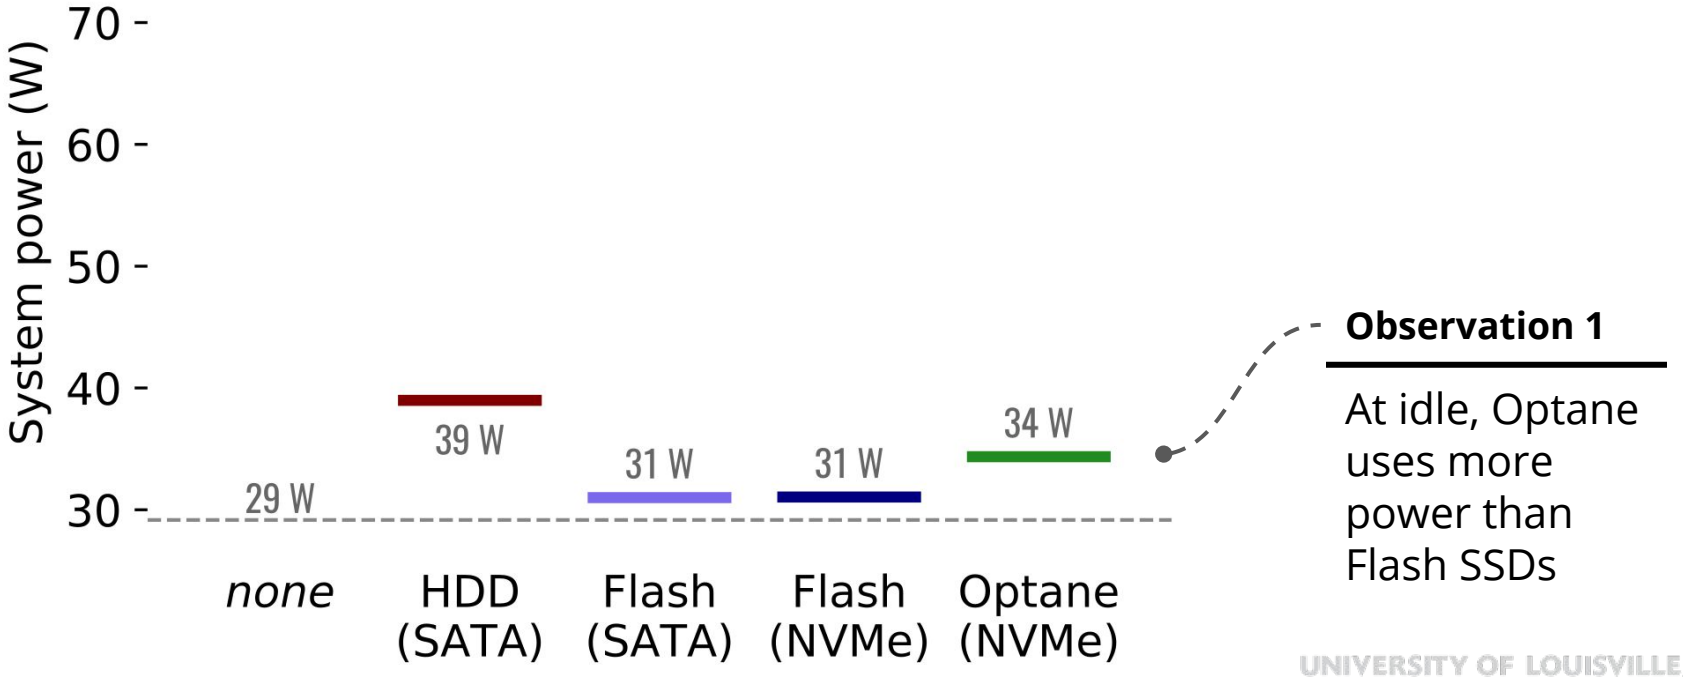

Idle vs. active power consumption

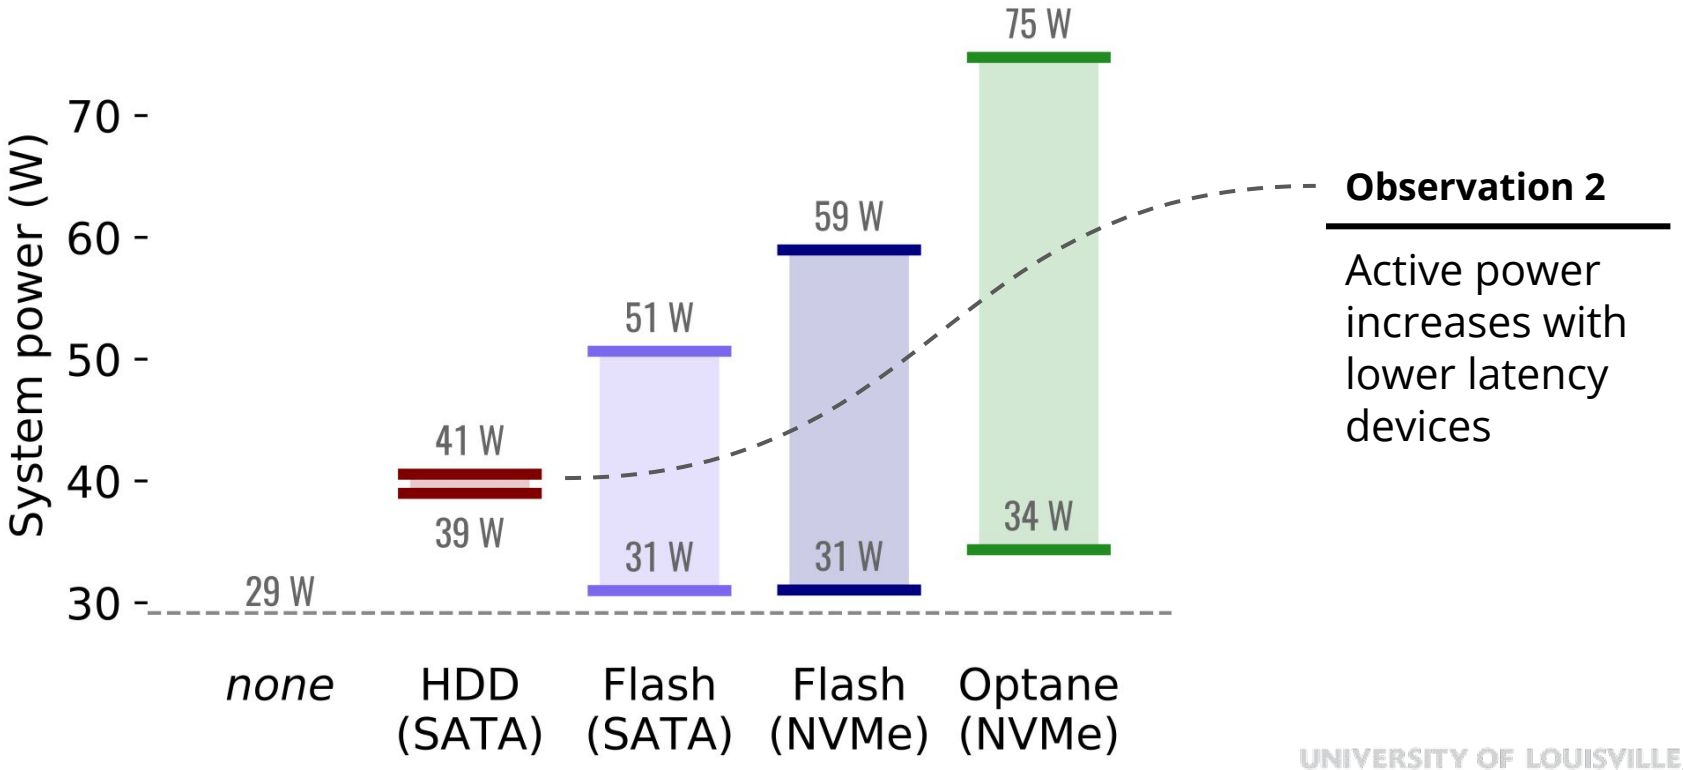

Observation 2

Active power increases with lower latency devices

Read vs. write power consumption

Observation 3

Newer storage generations have greater differences in power consumption between reads and writes

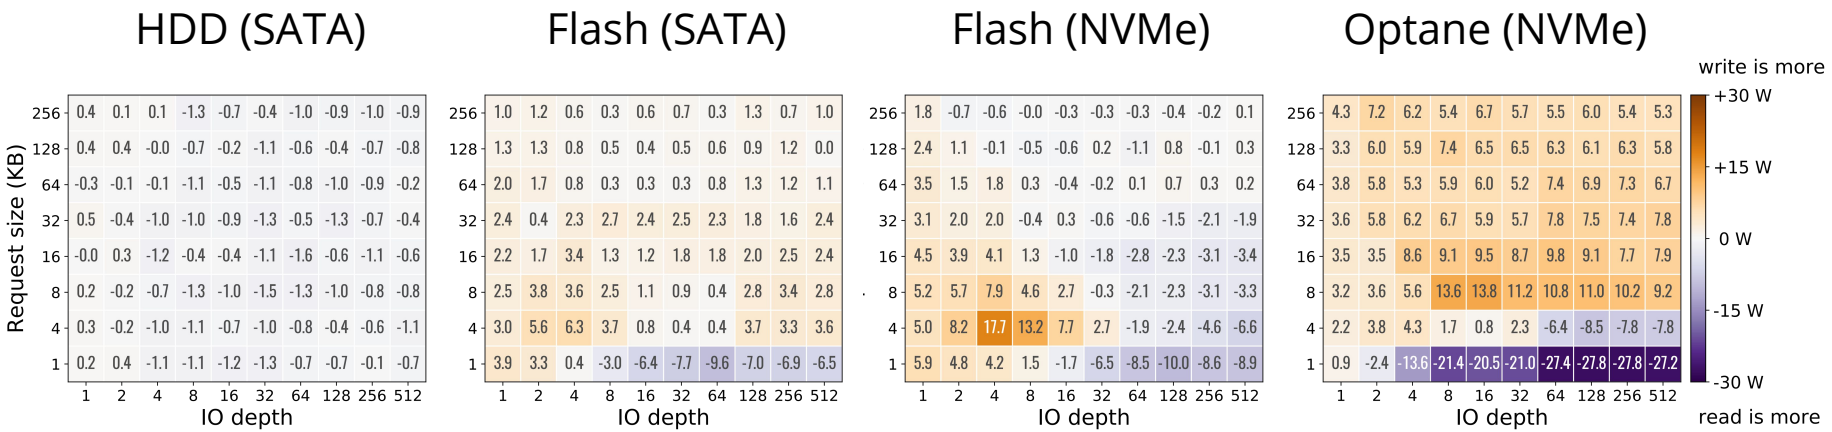

Energy proportionality

Consuming energy in proportion to the amount of work performed

System power
Power normalized to peak power observed

Storage performance utilization
IOPS normalized to peak IOPS for each device

Energy proportionality

Observation 4

Newer advancements lead to better energy proportionality

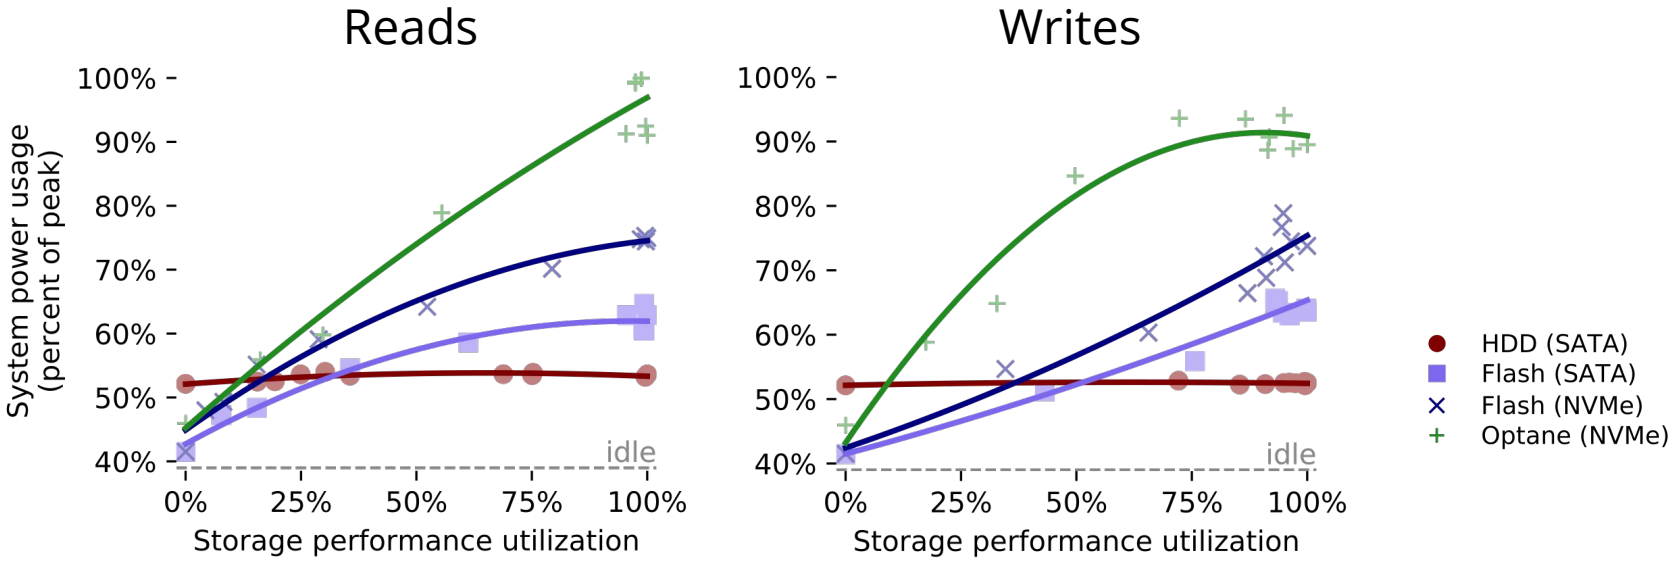

Impact on system software

Observation 5

As latency decreases, pressure on the system software increases, resulting in more overall energy consumption.

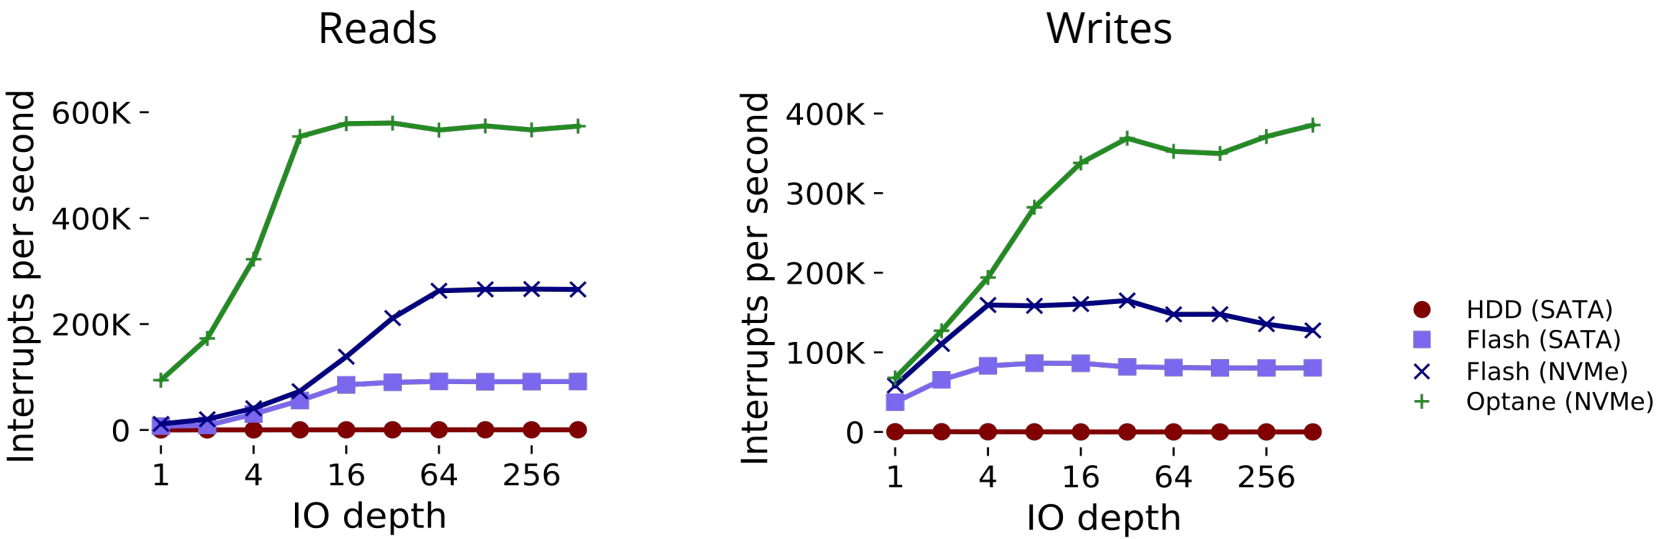

Impact on system software

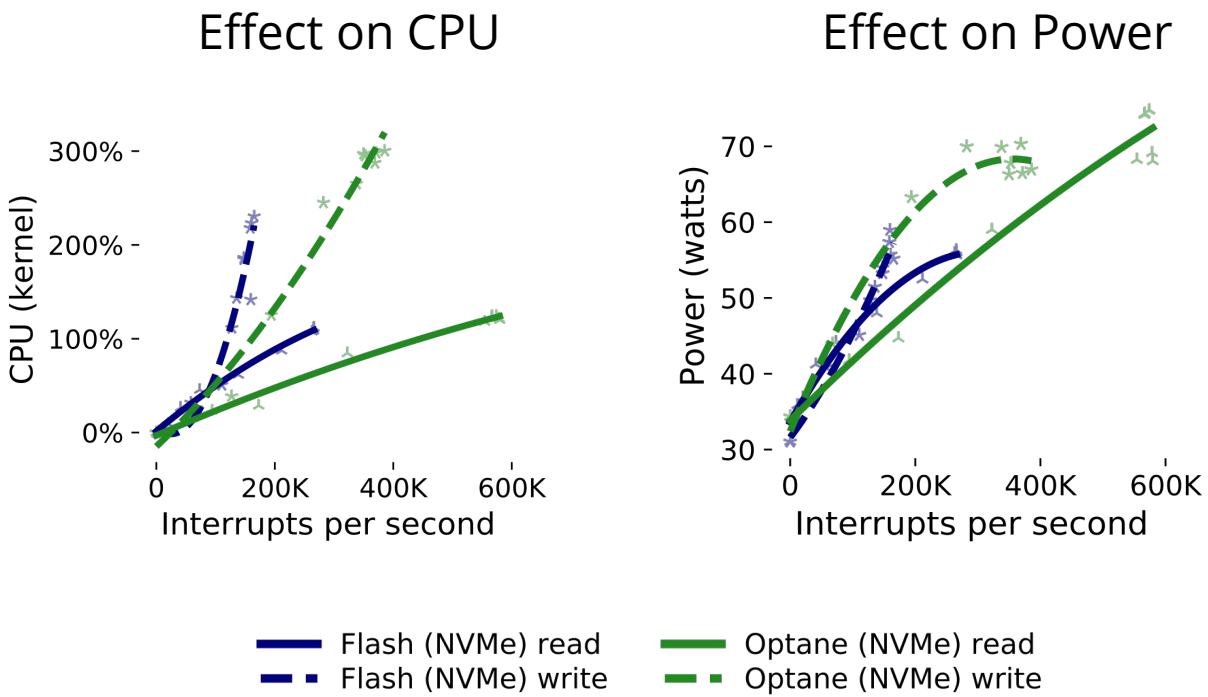

Overall energy efficiency metrics

Performance per unit energy

$$\textit{Bandwidth:}\quad \frac{\text{Bytes/second}}{\text{watt (joule/second)}} = \textbf{Bytes per joule}$$

$$\textit{Throughput:}\quad \frac{\text{IOPS}}{\text{watt (joule/second)}} = \textbf{IOs per joule}$$

Impact on overall energy efficiency

Observation 6

Energy efficiency as *bytes per joule* increases as request size increases.

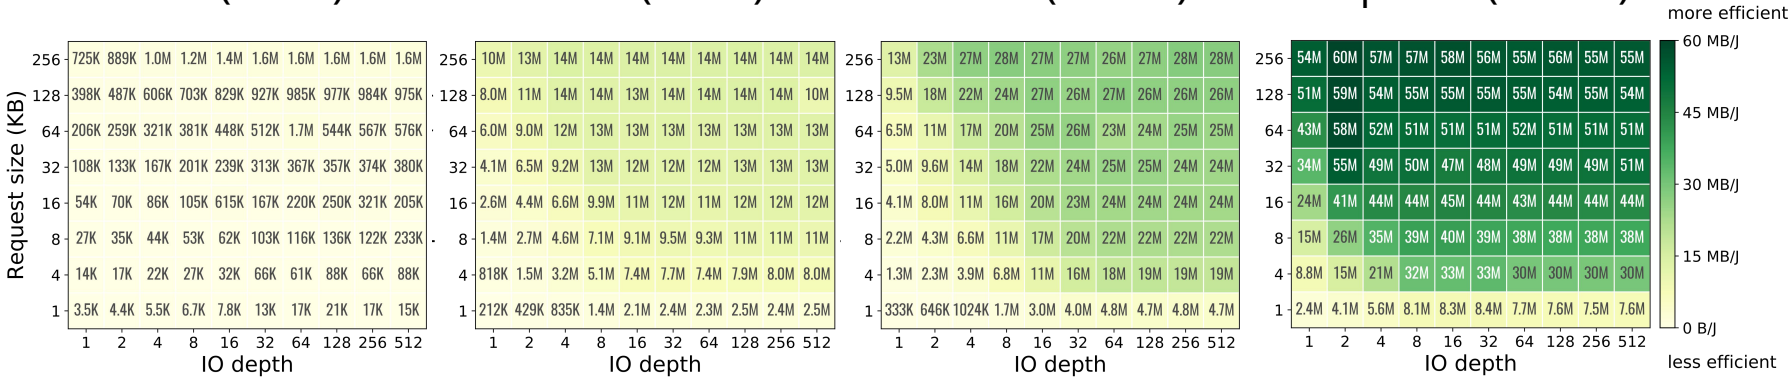

Bytes transferred per joule

Impact on overall energy efficiency

Observation 7

Energy efficiency as *IOs per joule* is coupled to internal parallelism.

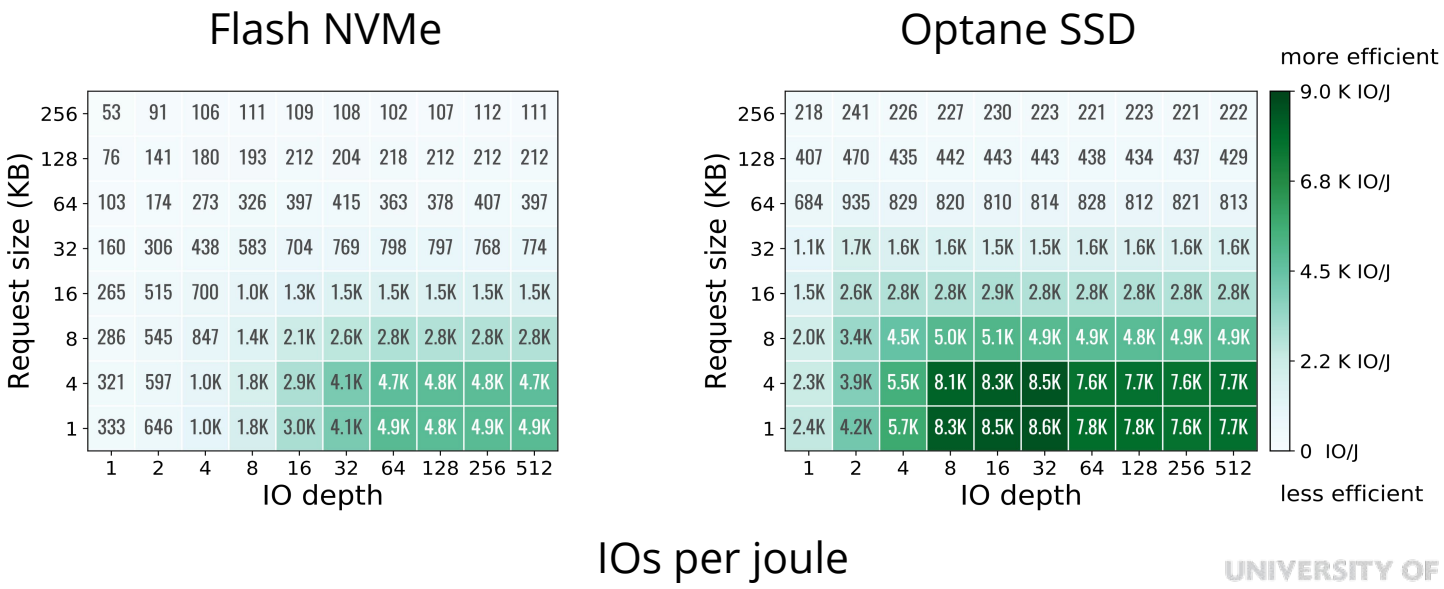

IOs per joule

Conclusion

Optane SSD has the greatest energy proportionality since its power consumption scales better than previous storage generations based on its range of throughput. Although Optane's peak power consumption is higher, it yields better energy efficiency as measured in bytes per joule and IOs per joule.

Discussion and future research

Rethink system software for energy efficiency:

- Polling vs. interrupts?
- New IO interfaces: *io_uring*, *SPDK*, etc.?
- Simpler mechanisms for *blk-mq* and *NVMe*?
- IO scheduling?
- Merging, prefetching, buffering, log structuring?

Discussion and future research

Investigate **R/W asymmetry**:

- Why do writes use more CPU but fewer interrupts?
- Do existing design choices get in the way of 3D XPoint?
- Strategies for hybrid drives for energy efficiency

Discussion and future research

Energy as a whole:

- Do performance optimizations lose sight of energy efficiency?
- Compared to CPU and memory, is the impact of the storage subsystem underestimated?

Thank you!

Questions? Please contact us:

Bryan Harris *bryan.harris.1@louisville.edu*

Nihat Altiparmak *nihat.altiparmak@louisville.edu*

References

[1] Luiz André Barroso and Urs Hölzle. The case for energy-proportional computing. Computer, 40(12):33–37, 2007. https://doi.org/10.1109/MC.2007.443.

[2] Kan Wu, Andrea Arpaci-Dusseau, and Remzi Arpaci-Dusseau. Towards an unwritten contract of Intel Optane SSD. In *11th USENIX Workshop on Hot Topics in Storage and File Systems (HotStorage '19)*, Renton, WA, July 2019. USENIX Association. https://www.usenix.org/conference/hotstorage19/presentation/wu-kan

[3] Sehgal, Priya, Vasily Tarasov, and Erez Zadok. "Evaluating Performance and Energy in File System Server Workloads." FAST 2010.