

# DI-SSD: Desymmetrized Interconnection Architecture and Dynamic Timing Calibration for Solid-State Drives

Ren-Shuo Liu and **Jian-Hao Huang**


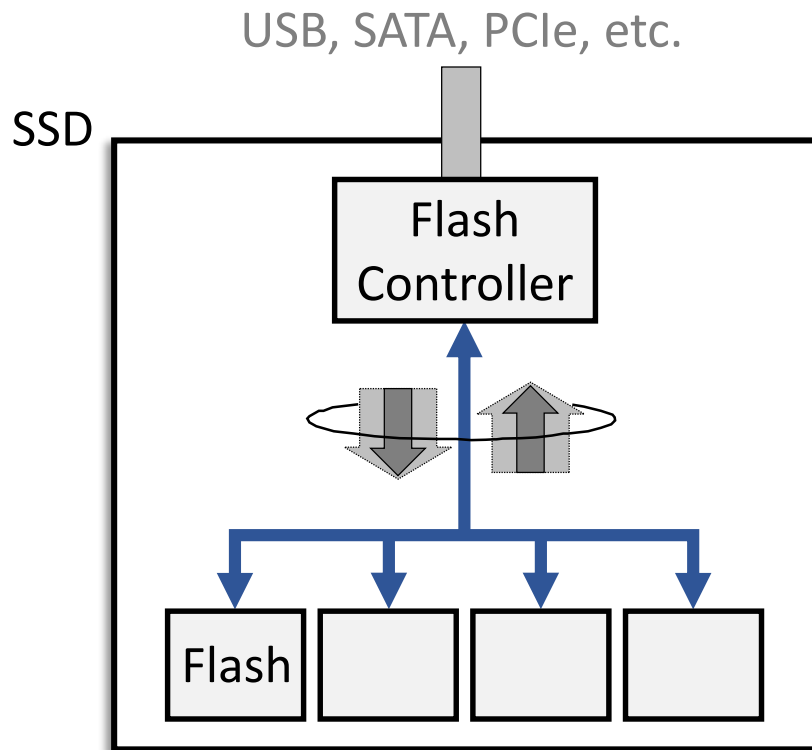


System and Storage Design Lab  
Department of Electrical Engineering  
**National Tsing Hua University, Taiwan**



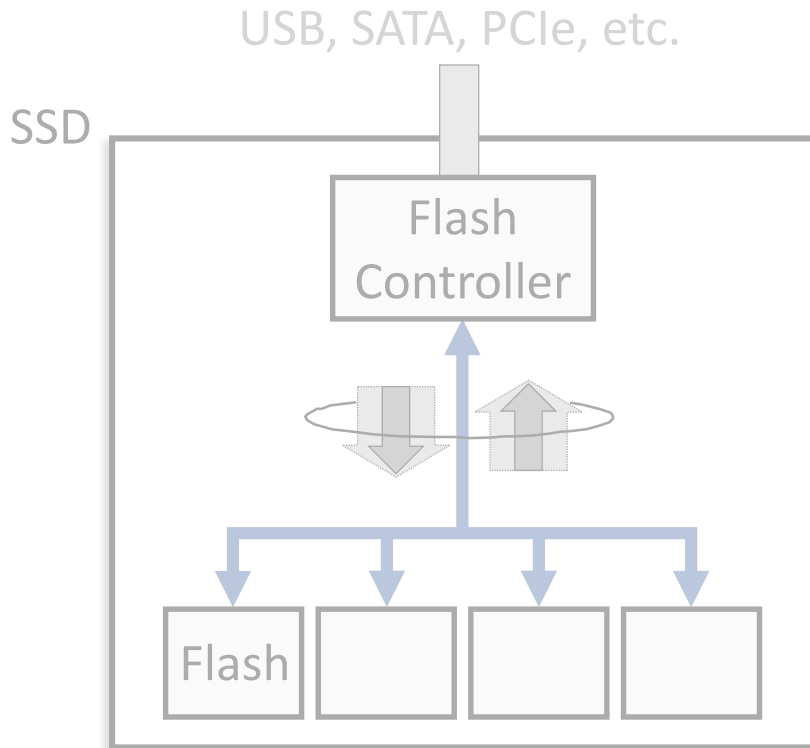
# Conventional SSD Architecture

## Symmetric Interconnection (SI) Architecture

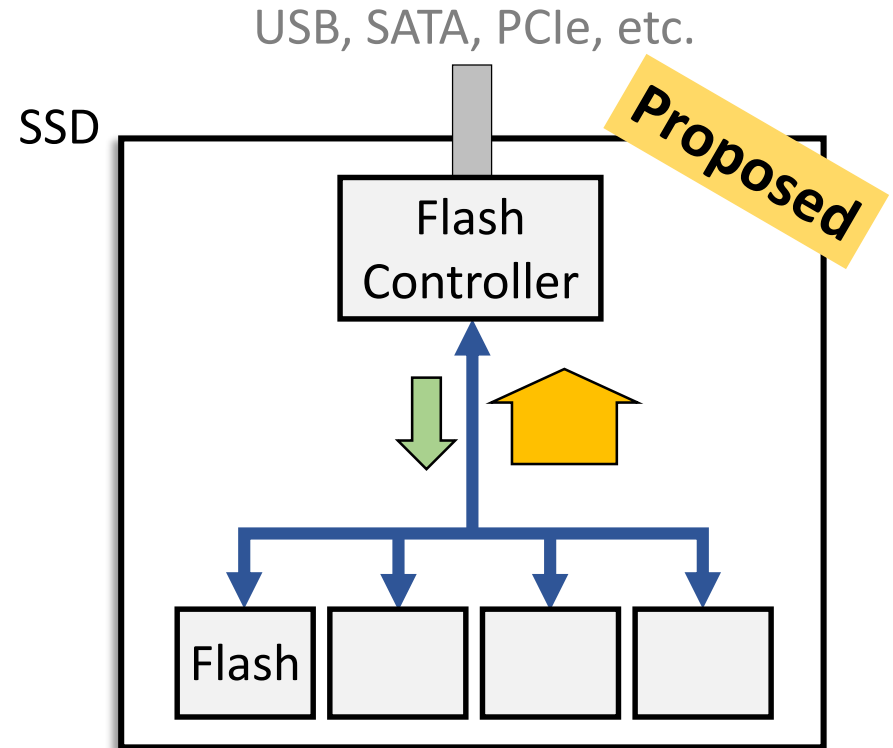


# Our Key Contribution

## Symmetric Interconnection (SI) Architecture



## Desymmetrized Interconnection (DI) Architecture



# Train Station Stair Analogy


- Desymmetrized widths
  - Up direction is wider
  - Down is narrower
- ✓ Can accommodate the rush of passengers heading upstairs when a train arrives
- ✓ Best utilizes the limited space of the train platform



# Outline

---




- Contributions
- SI vs. DI SSD architecture
- Proof of concept — dynamic timing calibration
- Evaluation
- Conclusions



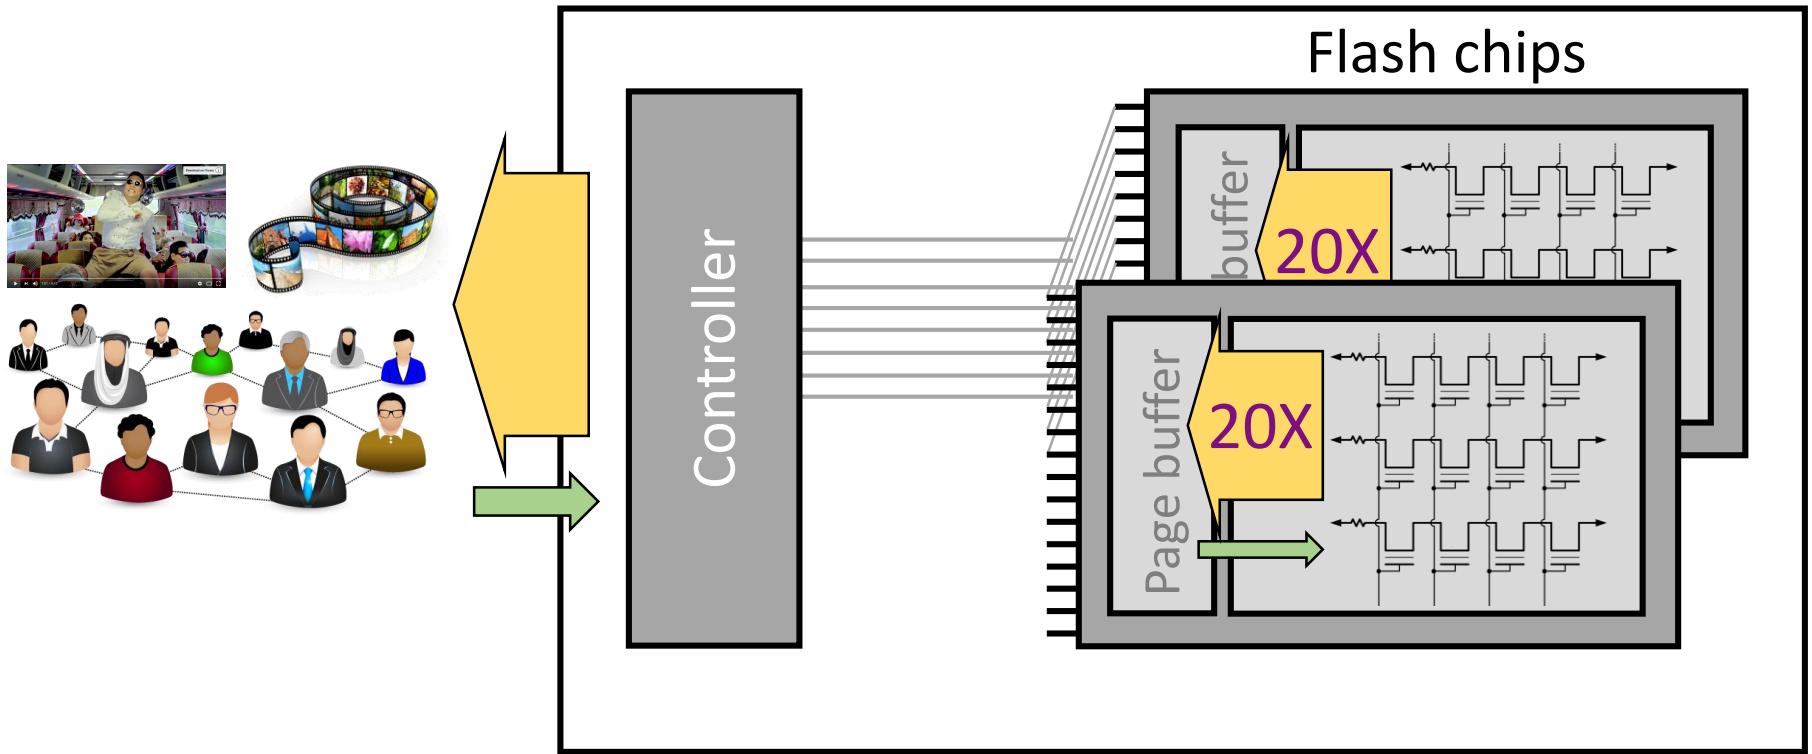
SSDs



# Read-Dominant Usage is important to SSDs



# Flash Memory Characteristics

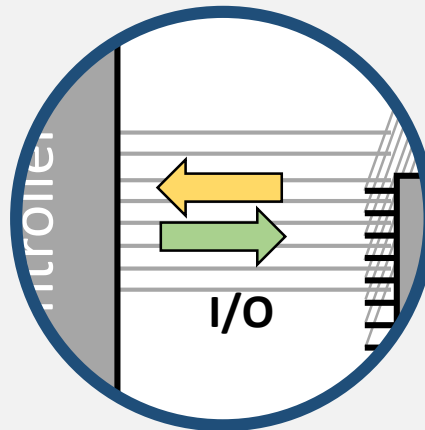


- Sensing a page: 0.1 ms
  - Programming a page: 1.3~2.6 ms
- } 20x gap

SSD



# Symmetric Interconnection (SI)




- Interestingly, the interconnections are symmetric
- Standardized by JEDEC and ONFi

Clock Frequency Grades (MHz)											
SDR	10	20	28	33	40	50					
DDR	20	33	50	66	83	100					
DDR2/DDR3	33	40	66	83	100	133	166	200	266	333	400



# SI's Dilemma & Compromise



- Dilemma
  - High bandwidth → overdesign for writes
  - Low bandwidth → bottleneck for reads
- Compromise
  - Select a speed somewhere between the 20x gap
  - Sacrifice the read strength of flash memory


# Real USB Drive Measurement

16 GB, USB 3.0

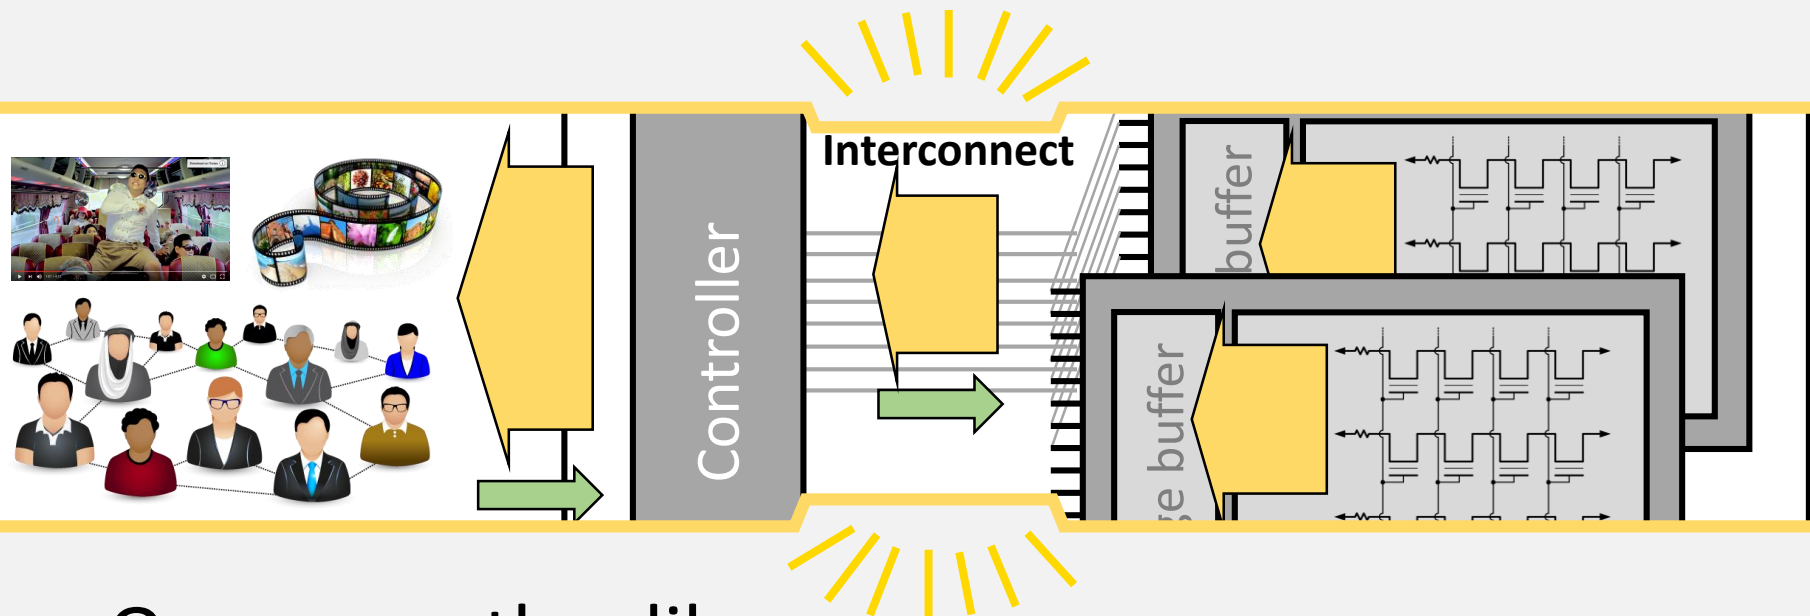


DDR 100 MHz clock  
(200 MT/s)

Sequential Read {  
16% Memory sensing  
72% Interconnection



# Desymmetrized Interconnection (DI)




- Overcome the dilemma
- Allow dedicating resources to **speeding up the flash-to-controller direction**
  - Avoid the cost of speeding up the other
- DI **breaks the norm** of SSD design for decades!

# DI Implementation

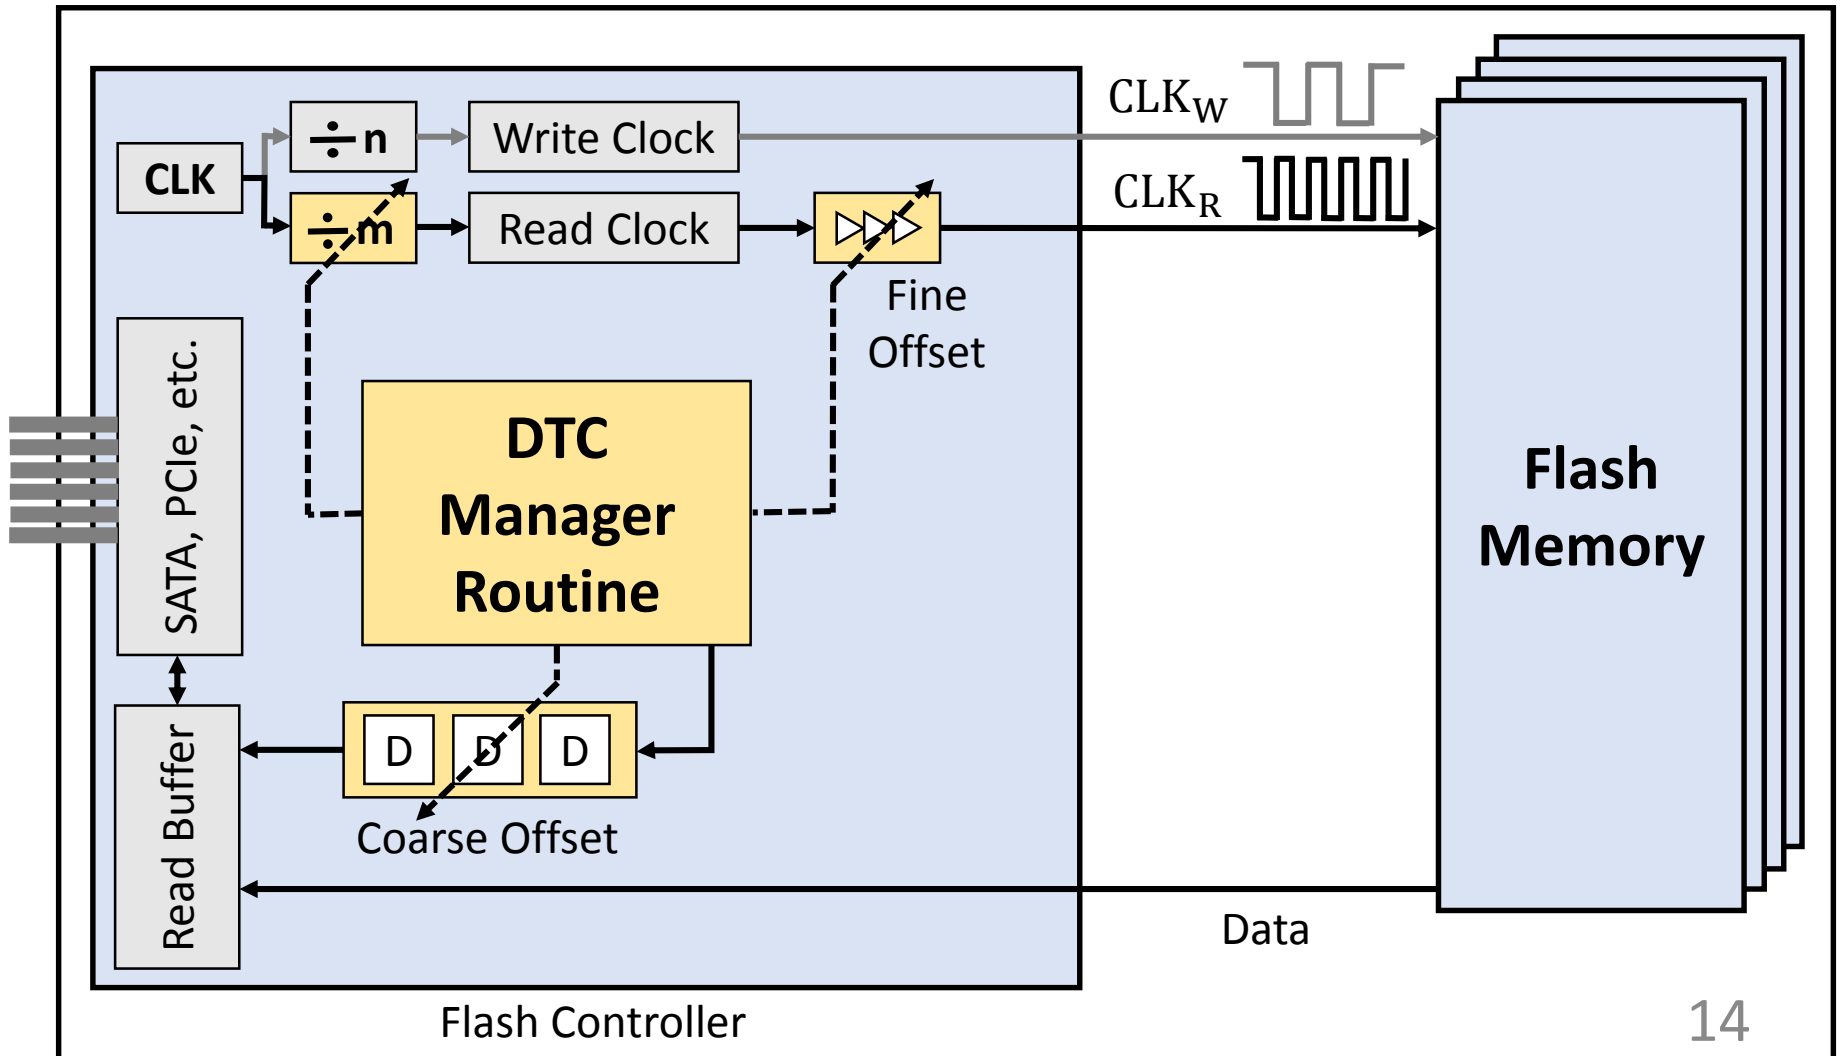
---

- DI is **architecture-level** design
- DI applies to various technologies
  - Serial, parallel, SDR, DDR, raw flash, managed flash, etc.

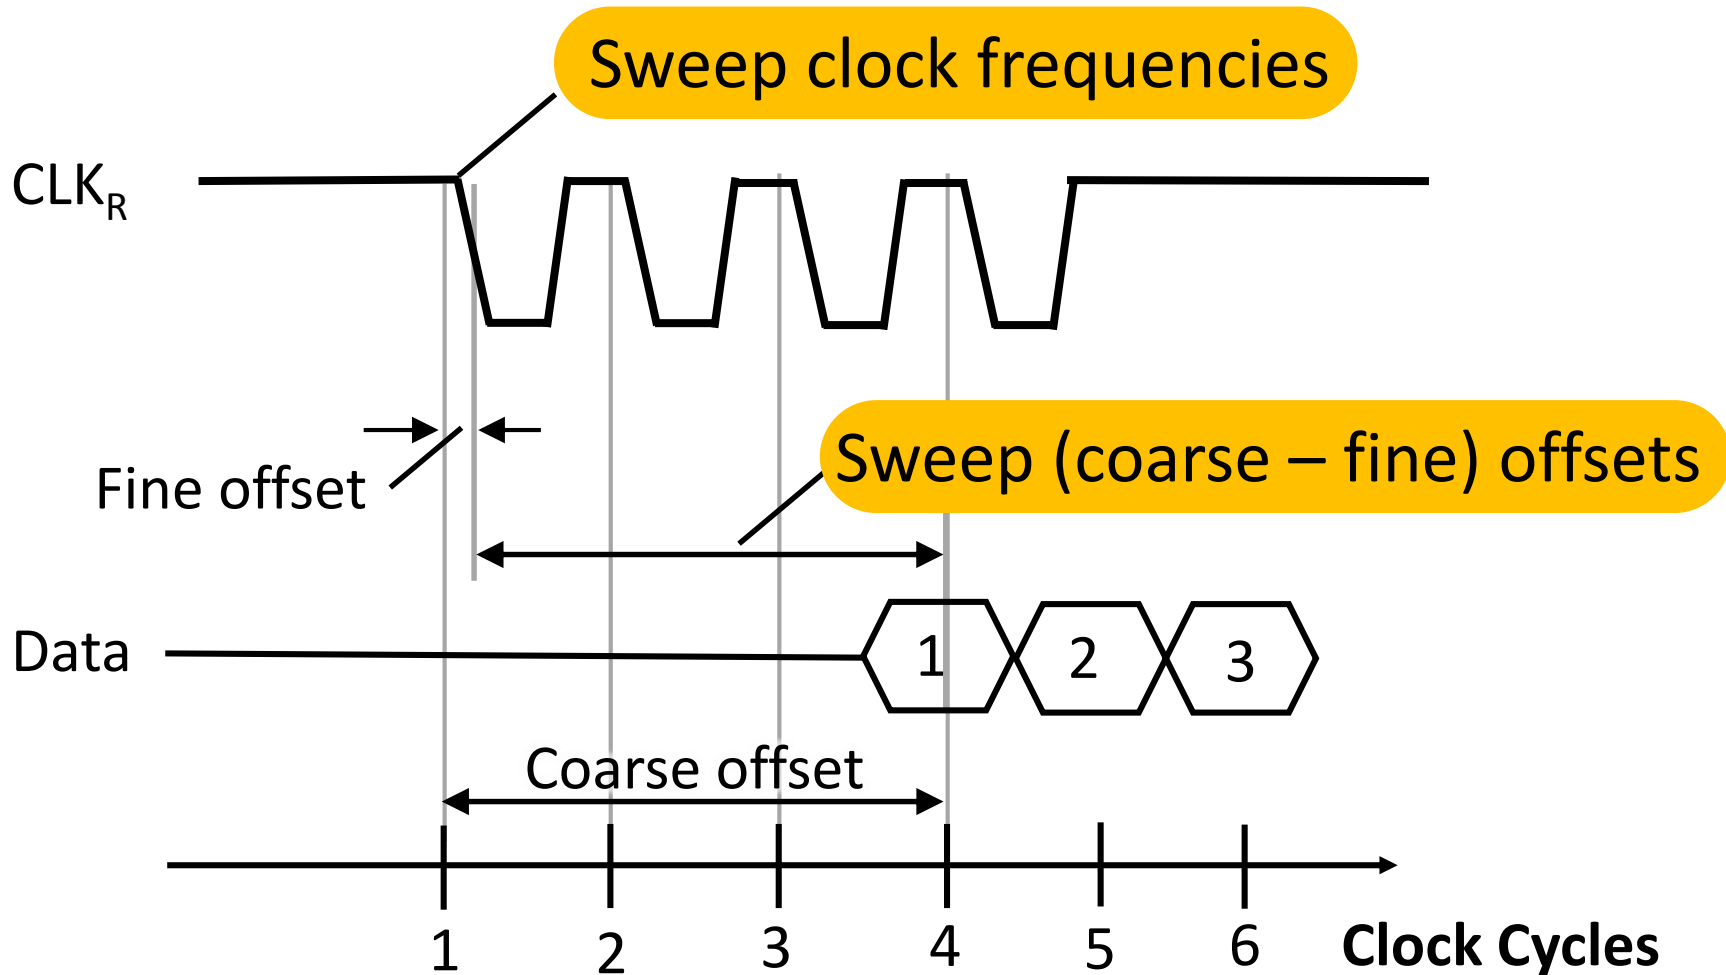


# Dynamic Timing Calibration (DTC)

- Proof-of-concept of DI-SSD on top of SDR flash

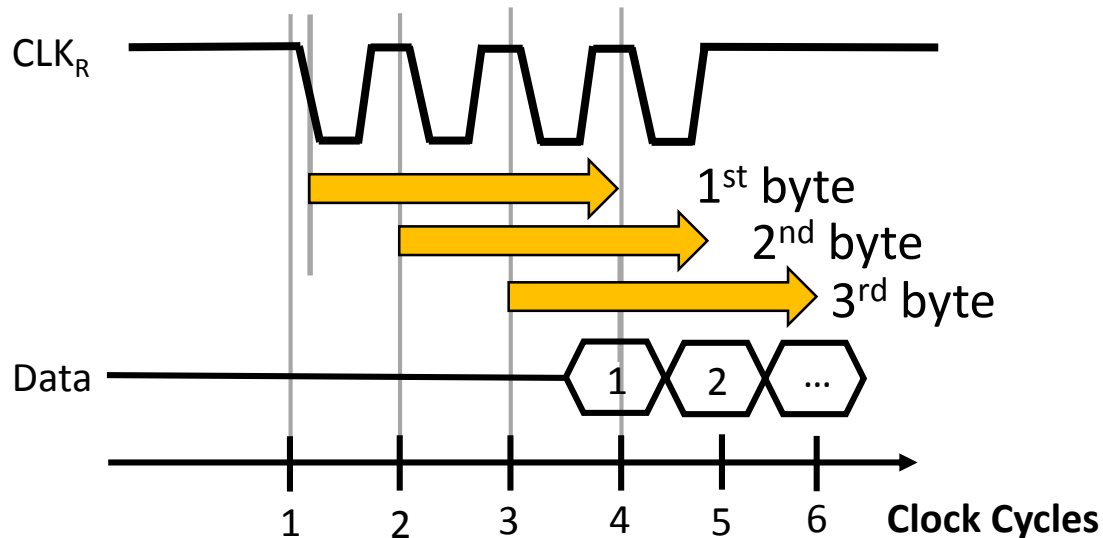


# DTC's Goals



# DTC's Benefits

- Reclaim flash's underutilized **frequency margins**
  - Fast/slow process corners
  - $\pm 10\%$  supply voltage
  - $0\sim 70^{\circ}\text{C}$  environments
- Exploit the fact that outputs can be **overlapped**





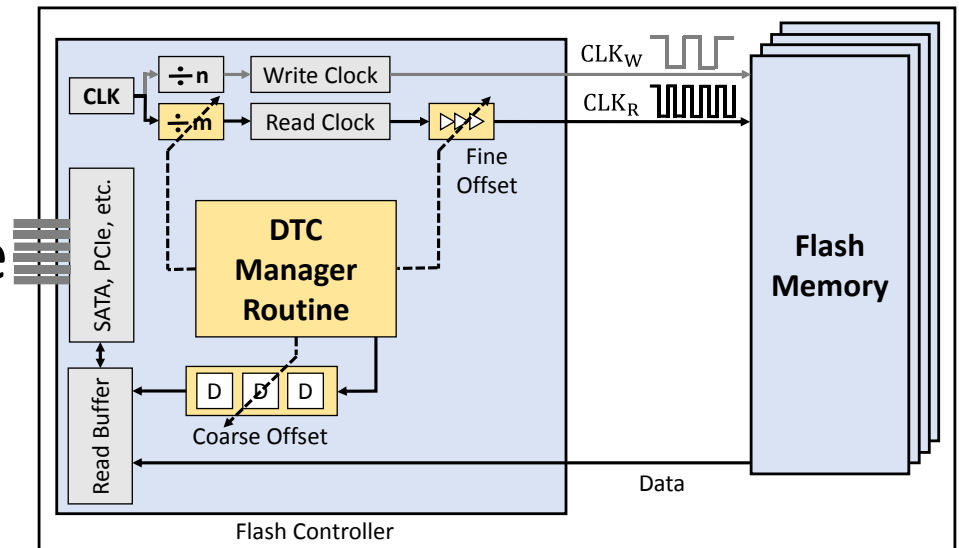
# Correctness Guarantees

---

- DTC only raises flash output frequency
  - Data transferring into the flash are intact
  - Data retention is intact
- Fallback mechanism
  - Controller can always fall back to a conservative output frequency ( $\text{CLK}_R$ )
  - Raw bit error rate is intact

# Insignificant Overhead

- Coarse and fine offsets
  - Few tens of logic gates and a multiplexer
  - Few flip flops
- Clock divider
  - Counter
- DTC routine
  - Upon power-on or significant temperature changes
  - On the processor that perform the FTL



# Outline

---

- Contributions
- SI vs. DI SSD architecture
- Proof of concept — dynamic timing calibration
- **Evaluation**
- Conclusions

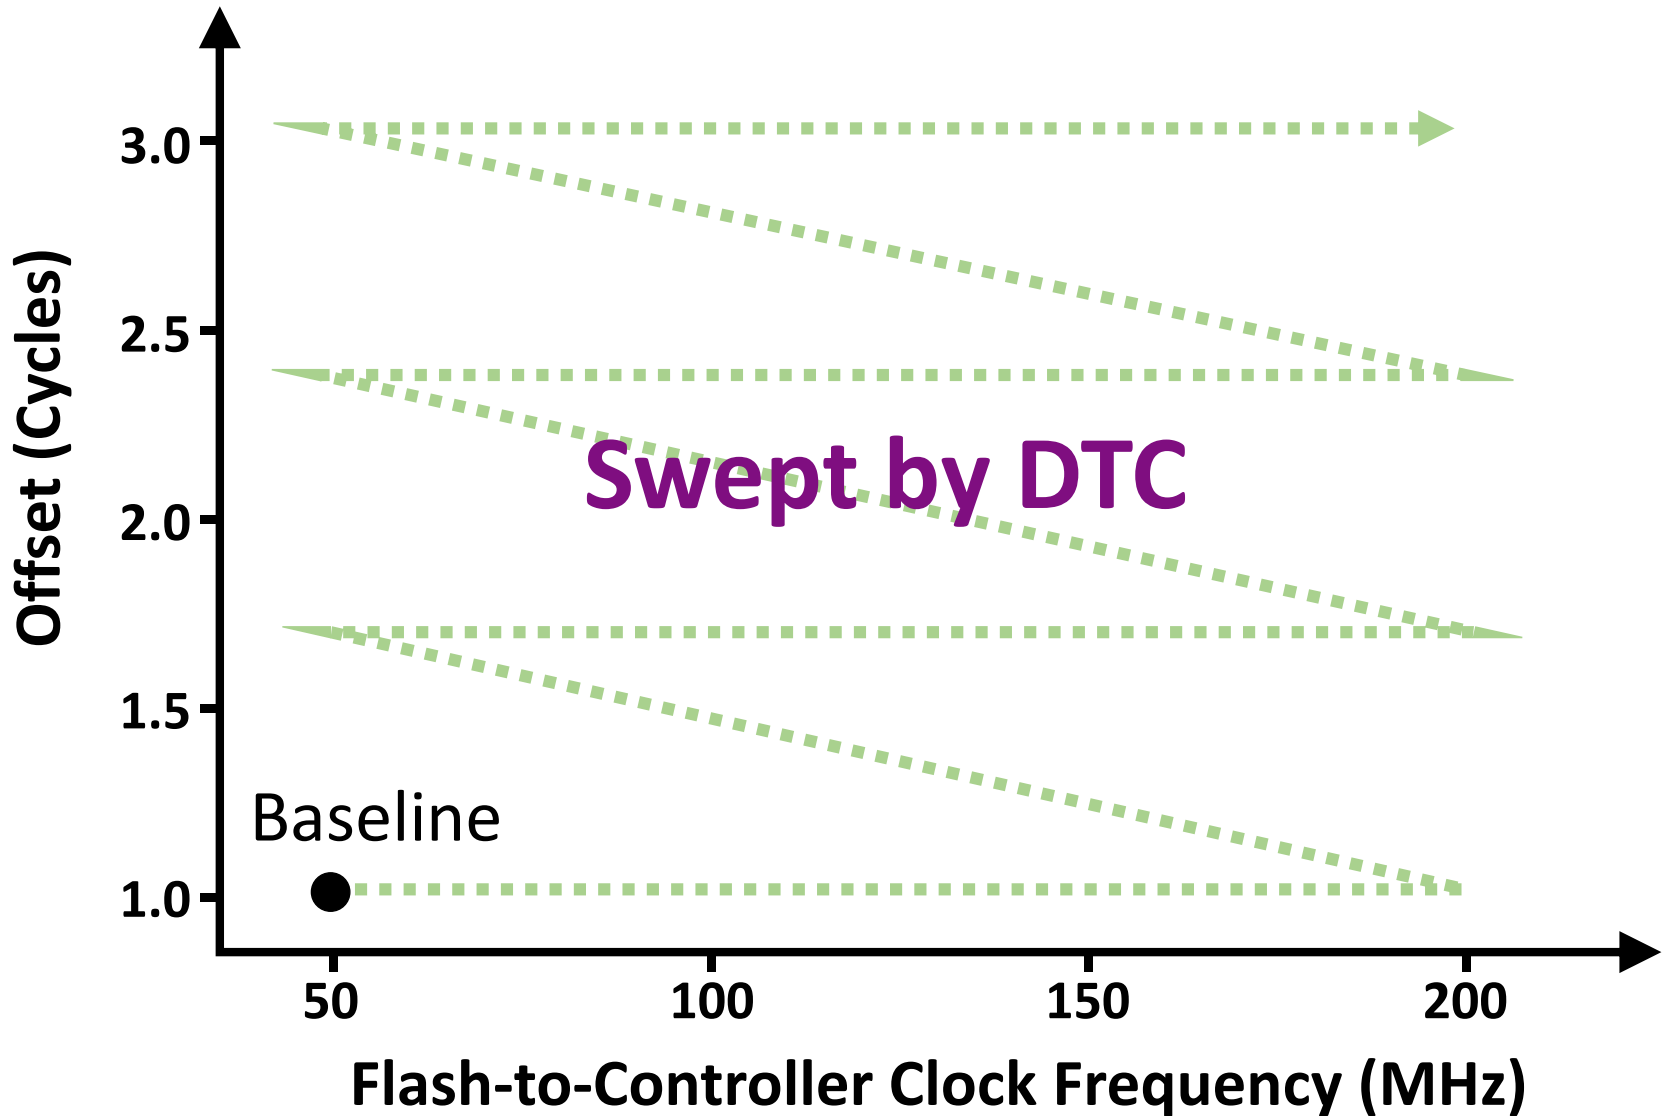
# Experimental Setup

- Chip level (DTC)
  - Industrial strength IC tester
  - Real 20 nm MLC NAND flash
  - Emulate DTC and measure its benefits
- System level (DI-SSD)
  - Extended SSDSim simulator
  - Synthetic workloads
  - Datacenter workloads

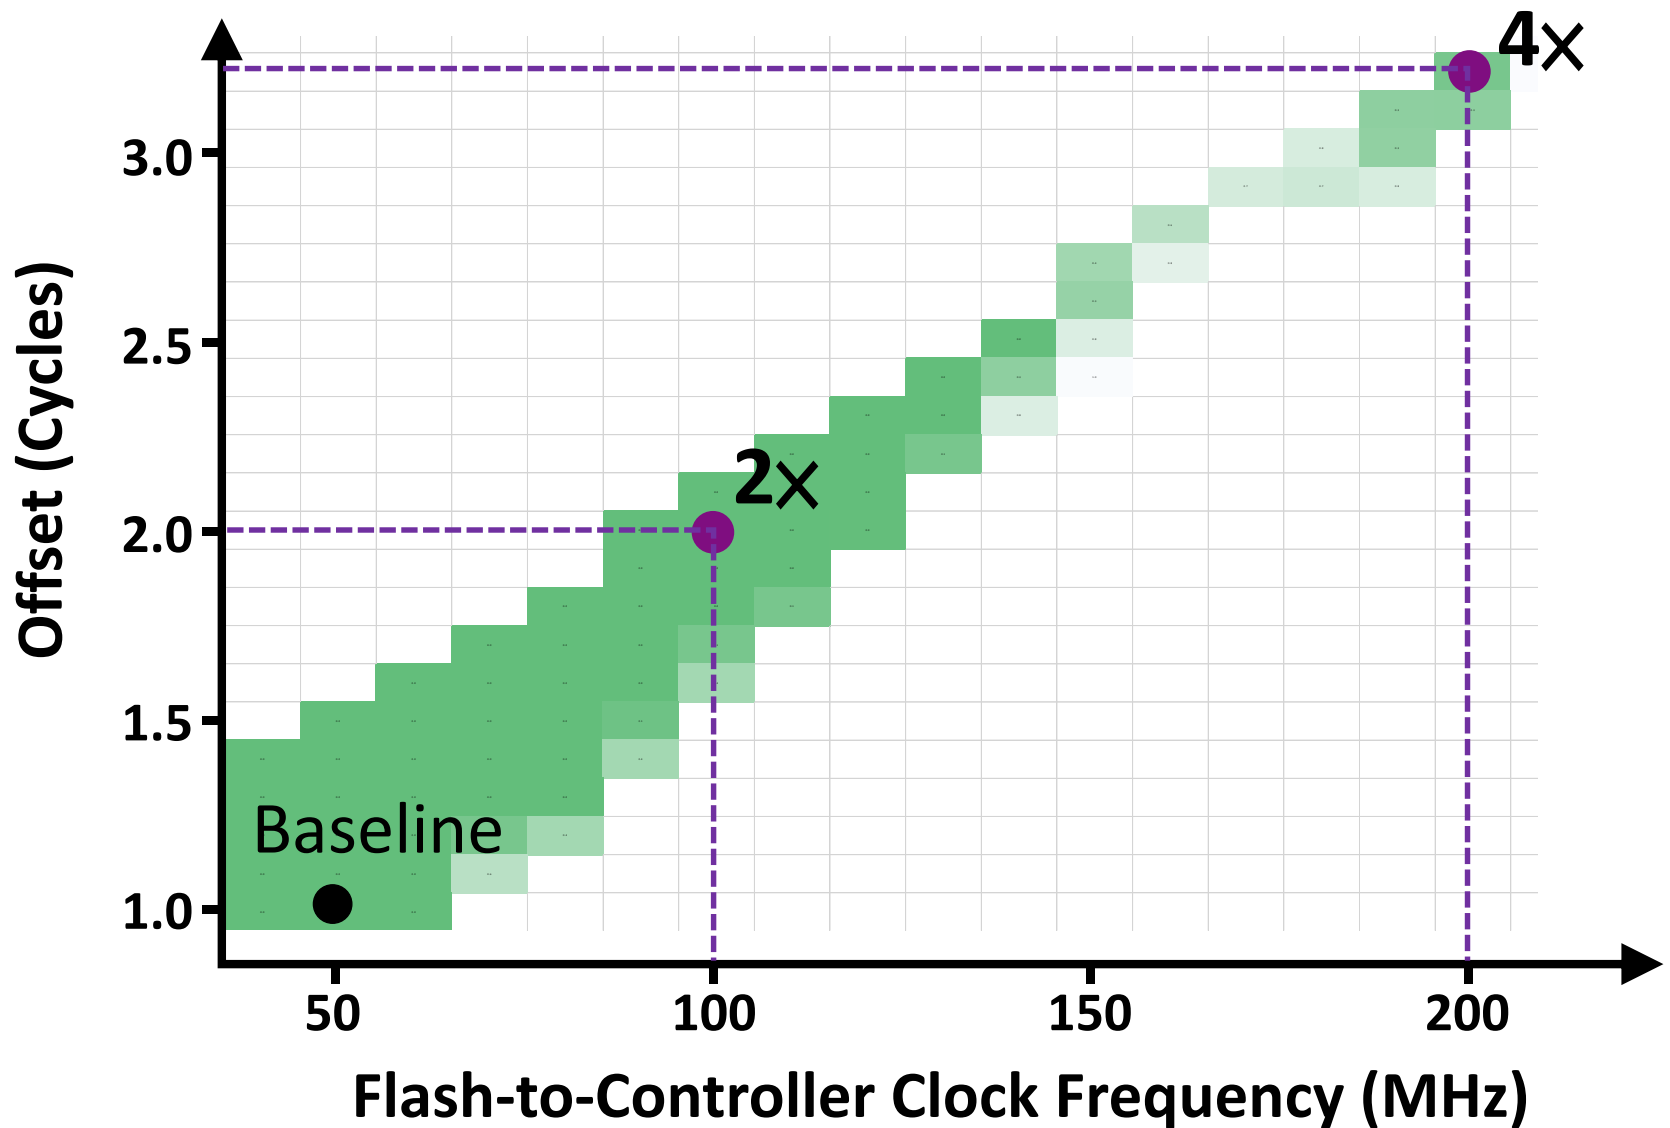


SSD Capacity	1 TB	Interconnection Frequency	50 MHz
#Interconnections	8	Page Read Latency	100 $\mu$ s
#Flash Chips	64	Page Write Latency	1.3 ms
#Planes/Chip	2	Erase Latency	3 ms
#Blocks/Plane	2000	SSD Interface	PCIe 2.0 x8
#Pages/Block	512	Over-Provisioning	15%
Page Size	8 KB	GC Threshold	5%

# Chip-Level (DTC) Results

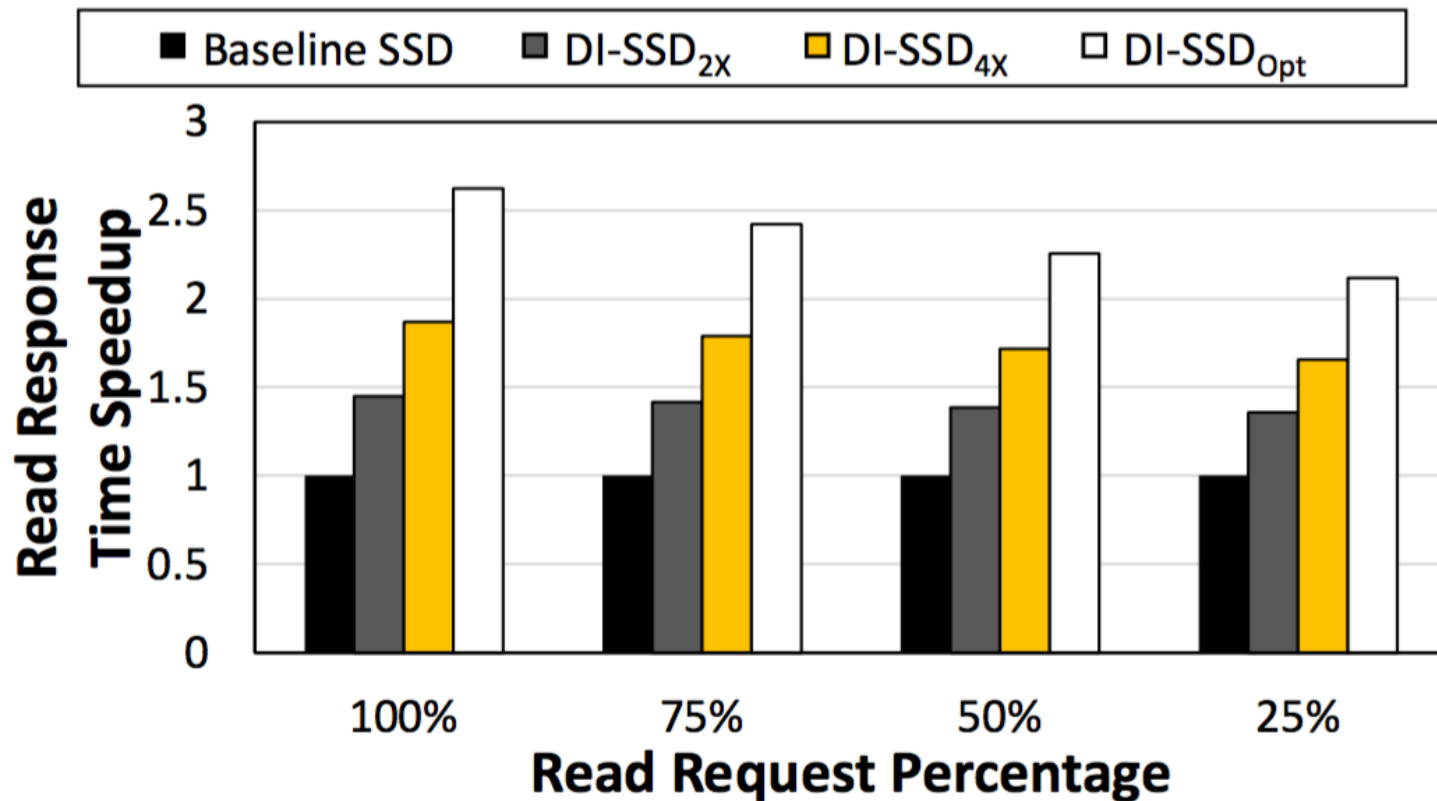


# Chip-Level (DTC) Results




# Synthetic Workloads

- DI-SSD can yield up to 1.9x read response time speedup



# Datacenter Workloads

- DI-SSD can yield 1.7x read response time speedup on average





# Outline

---

- Contributions
- SI vs. DI SSD architecture
- Proof of concept — dynamic timing calibration
- Evaluation
- **Conclusions**

# Conclusions

---

- **Symmetric interconnections (SI)** are widely used and accepted in SSD architecture design
- However, SI is suboptimal
  - Flash sensing is 20× faster than programming
- We propose **desymmetrized interconnection SSD architecture (DI-SSD)**
  - Dynamic timing calibration (DTC) as a proof of concept
  - 4× higher flash-to-controller speed
  - 1.7 to 1.9× read response time improvement

# DI-SSD: Desymmetrized Interconnection Architecture and Dynamic Timing Calibration for Solid-State Drives

Ren-Shuo Liu and Jian-Hao Huang

THANK YOU

