

# The Bleak Future of NAND Flash Memory

Laura M. Grupp<sup>\*</sup>, John D. Davis<sup>γ</sup>, Steven Swanson<sup>\*</sup>

<sup>\*</sup>Non-volatile Systems Laboratory  
Department of Computer Science and Engineering  
University of California, San Diego

<sup>γ</sup>Microsoft Research

Microsoft<sup>®</sup>  
**Research**



**NVSL**  
Non-volatile Systems Laboratory



**UCSD CSE**  
Computer Science and Engineering


# Flash's Future: Bright

---

**Reliability**



**Performance**




**Cost Per Capacity**




# Flash's Future: ~~Bright~~ Bleak

---

## Reliability



## Performance

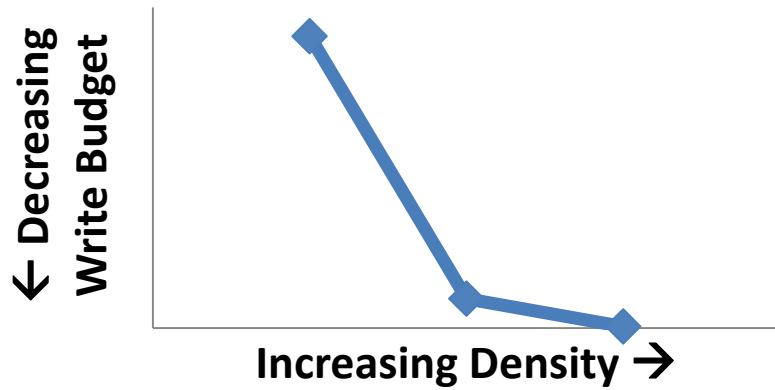


## Cost Per Capacity




# Flash's Future: ~~Bright~~ Bleak

## Reliability



## Performance




## Cost Per Capacity




# Flash's Future: ~~Bright~~ Bleak

## Reliability



## Performance




## Cost Per Capacity

Will the price decline be enough?



# Predicting Future Flash-Based SSDs


Fixed SSD Architecture





Model's Equations

$$\int z \, dV = \frac{\pi r^2}{VH^2} \int_0^h (z^3 - 2z^2H + zH^2) \, dz$$
$$= \frac{\pi r^2}{VH^2} \left[ \frac{z^4}{4} - \frac{2z^3H}{3} + \frac{z^2H^2}{2} \right]_0^h$$
$$= \frac{\pi r^2}{VH^2} \left[ \frac{1}{4} - \frac{2H}{3h} + \frac{H^2}{2h^2} \right]$$

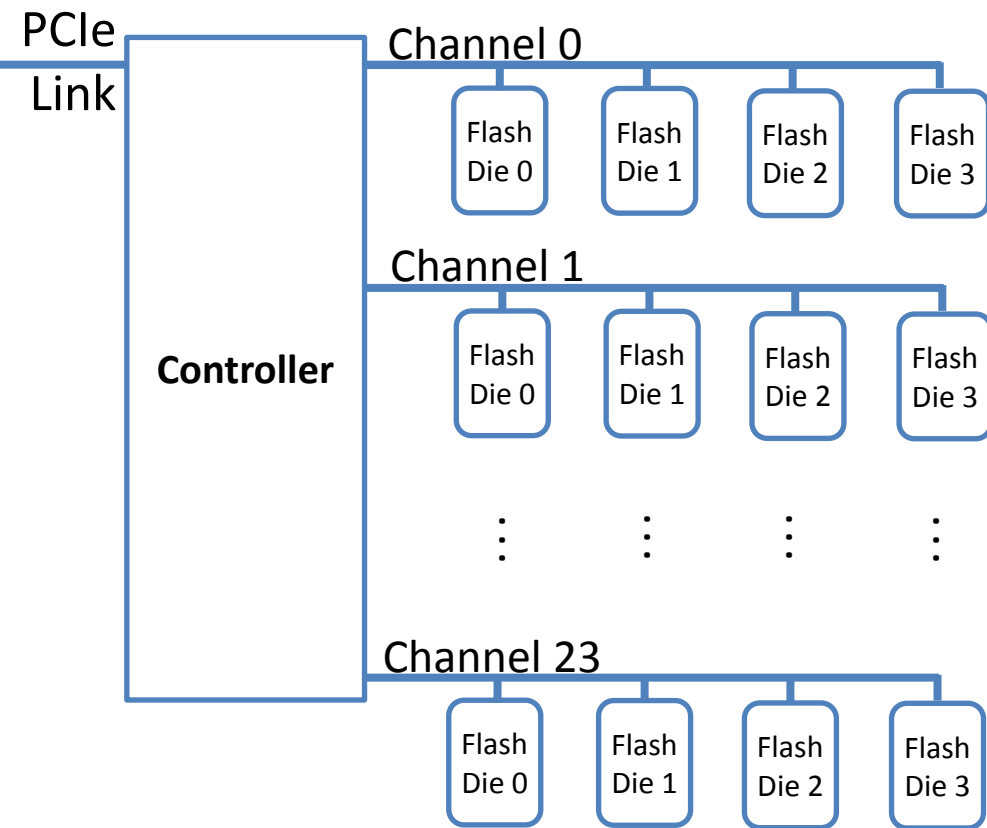
Flash Chip Trends




SSD Trends



# The Constant-Die-Count SSD (SSD-CDC)




- Represents High-End (FusionIO, Virident, OCZ)
- Baseline
  - 96 dies
  - 320 GB
  - 34nm, MLC
- Assumptions
  - Constant die count
  - Unlimited PCIe Link
  - Channel Speed: 400MB/s



# The Metrics

---

- Capacity



- Latency




- Throughput






# Increasing Density: Multi-bit Cells

---




Floating Gate  
(modifies  $V_{TH}$ )


SLC  
Single-Level Cell  
(1 bit)



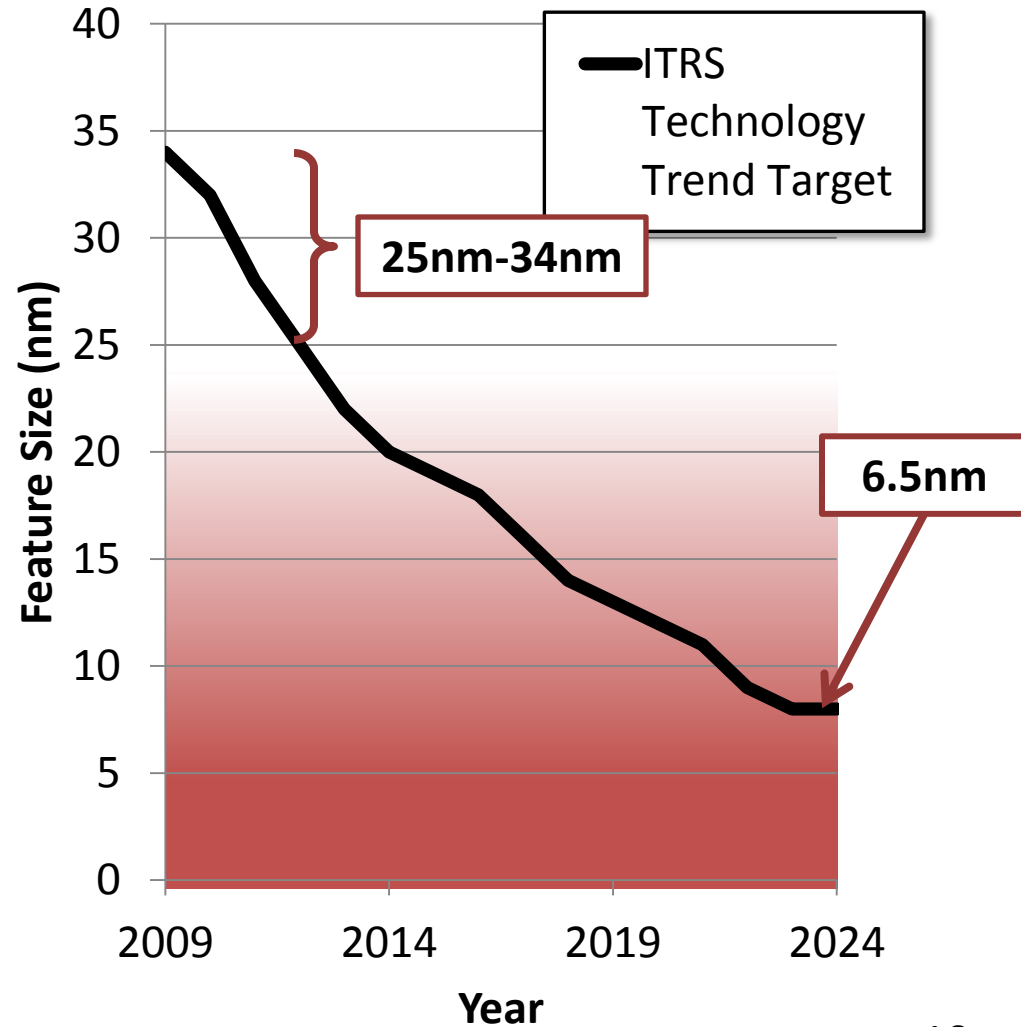

MLC  
Multi-Level Cell  
(2 bits)



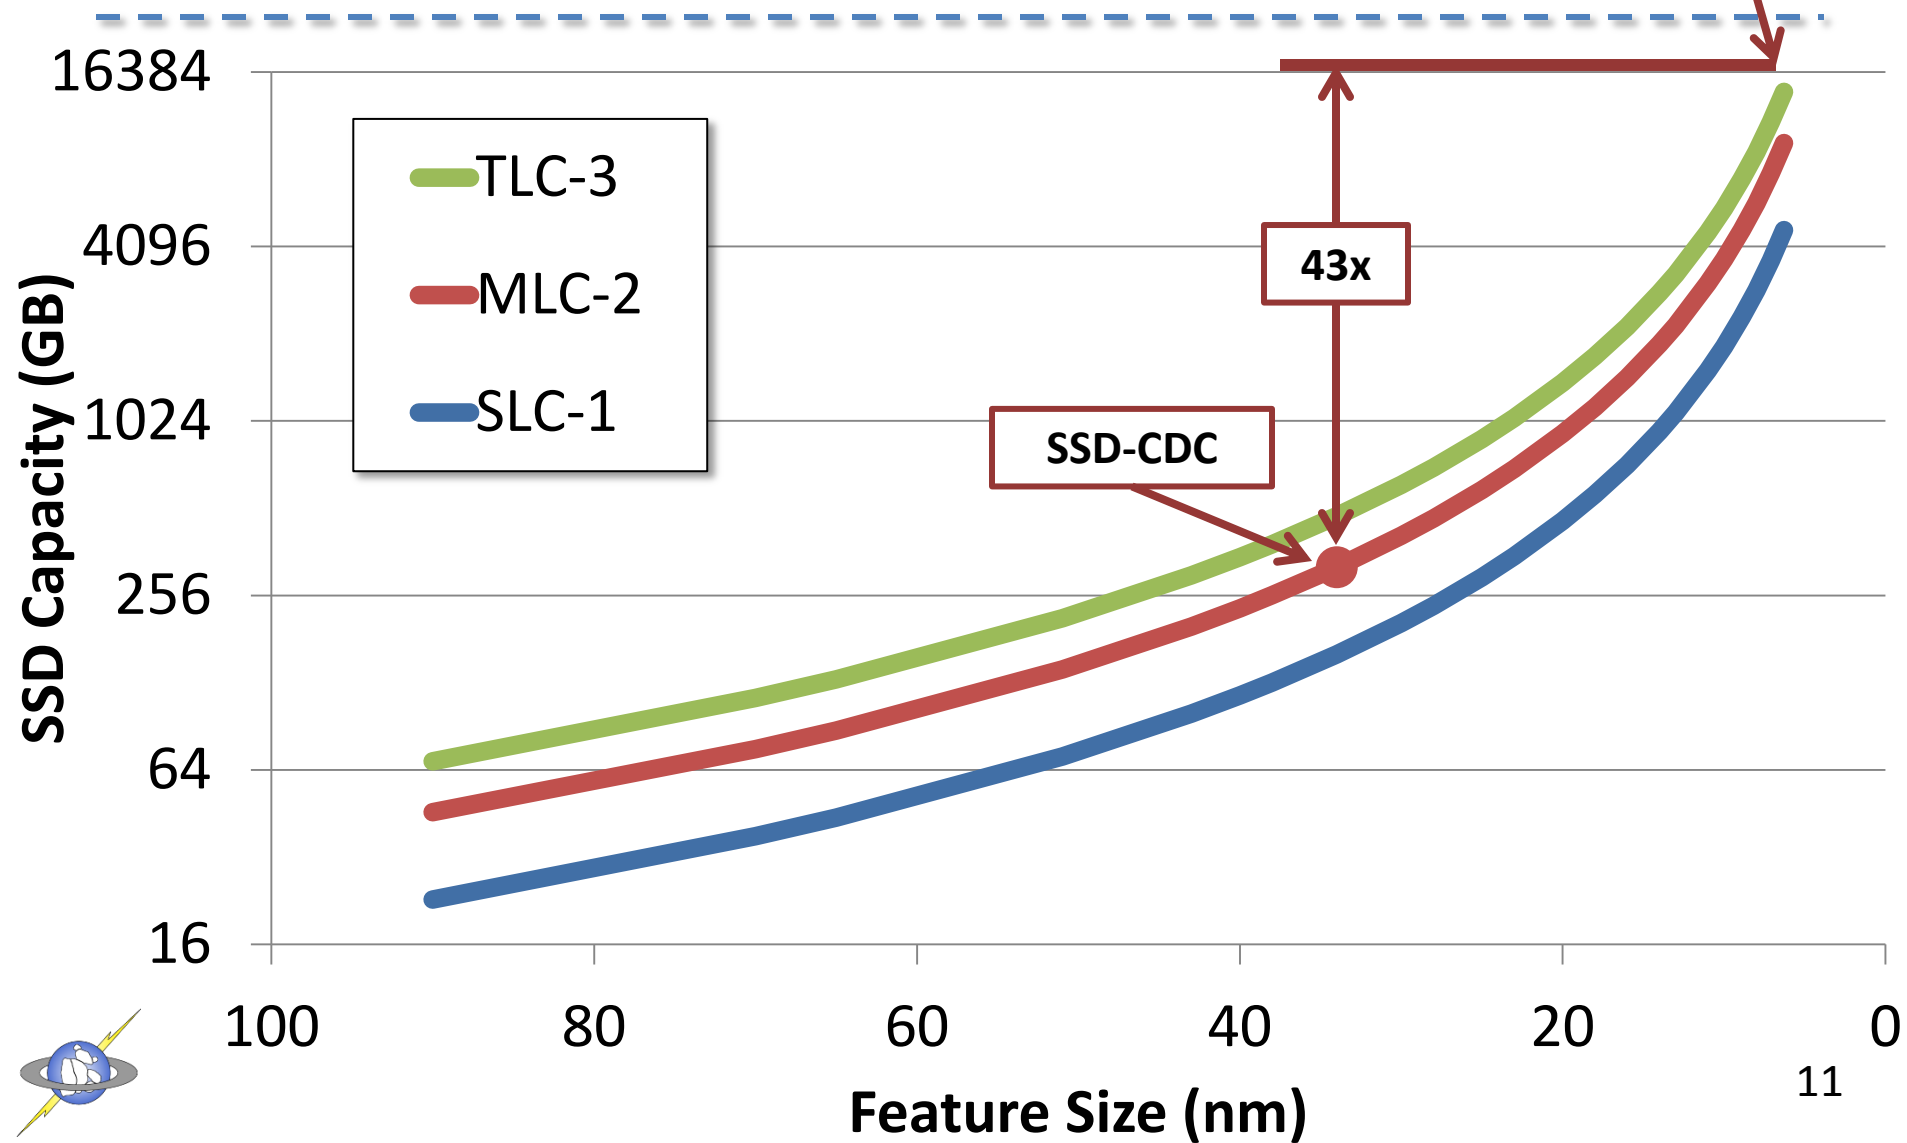
TLC  
Triple-Level Cell  
(3 bits)



# Increasing Density: Moore's Law



# Capacity



Best Possible  
by 2024

43x

SSD-CDC

TLC-3

MLC-2

SLC-1



# The Metrics

---

- Capacity: 43x

- Latency




- Throughput



# Collecting Flash Latency Trends


---

- In-house flash testing rig
  - XUP Virtex-II
  - Daughter board
  - 10ns resolution
- Chip Collection
  - 45 chips
  - 6 companies
  - 25nm-72nm
  - SLC, MLC, TLC




# Empirical Data

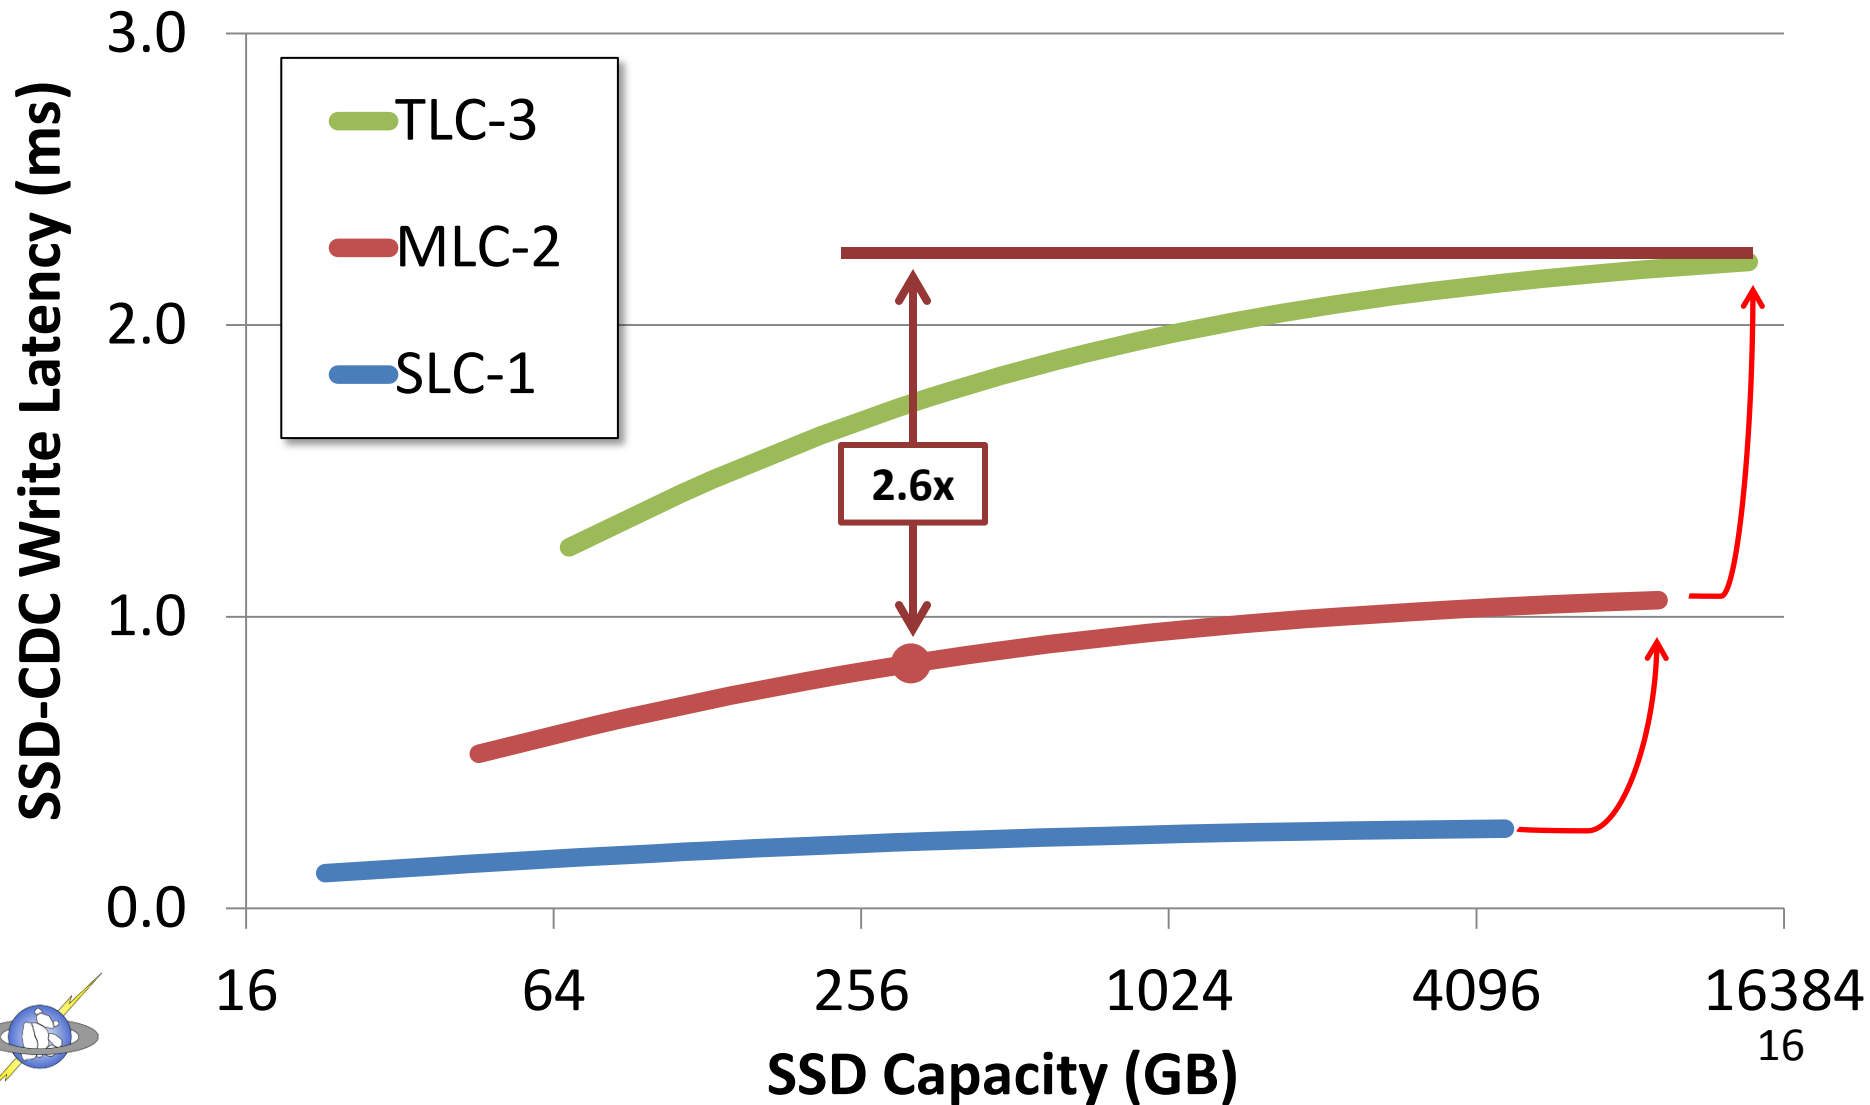
---



# Scaling Trends in Empirical Data



# Write Latency of SSD-CDC






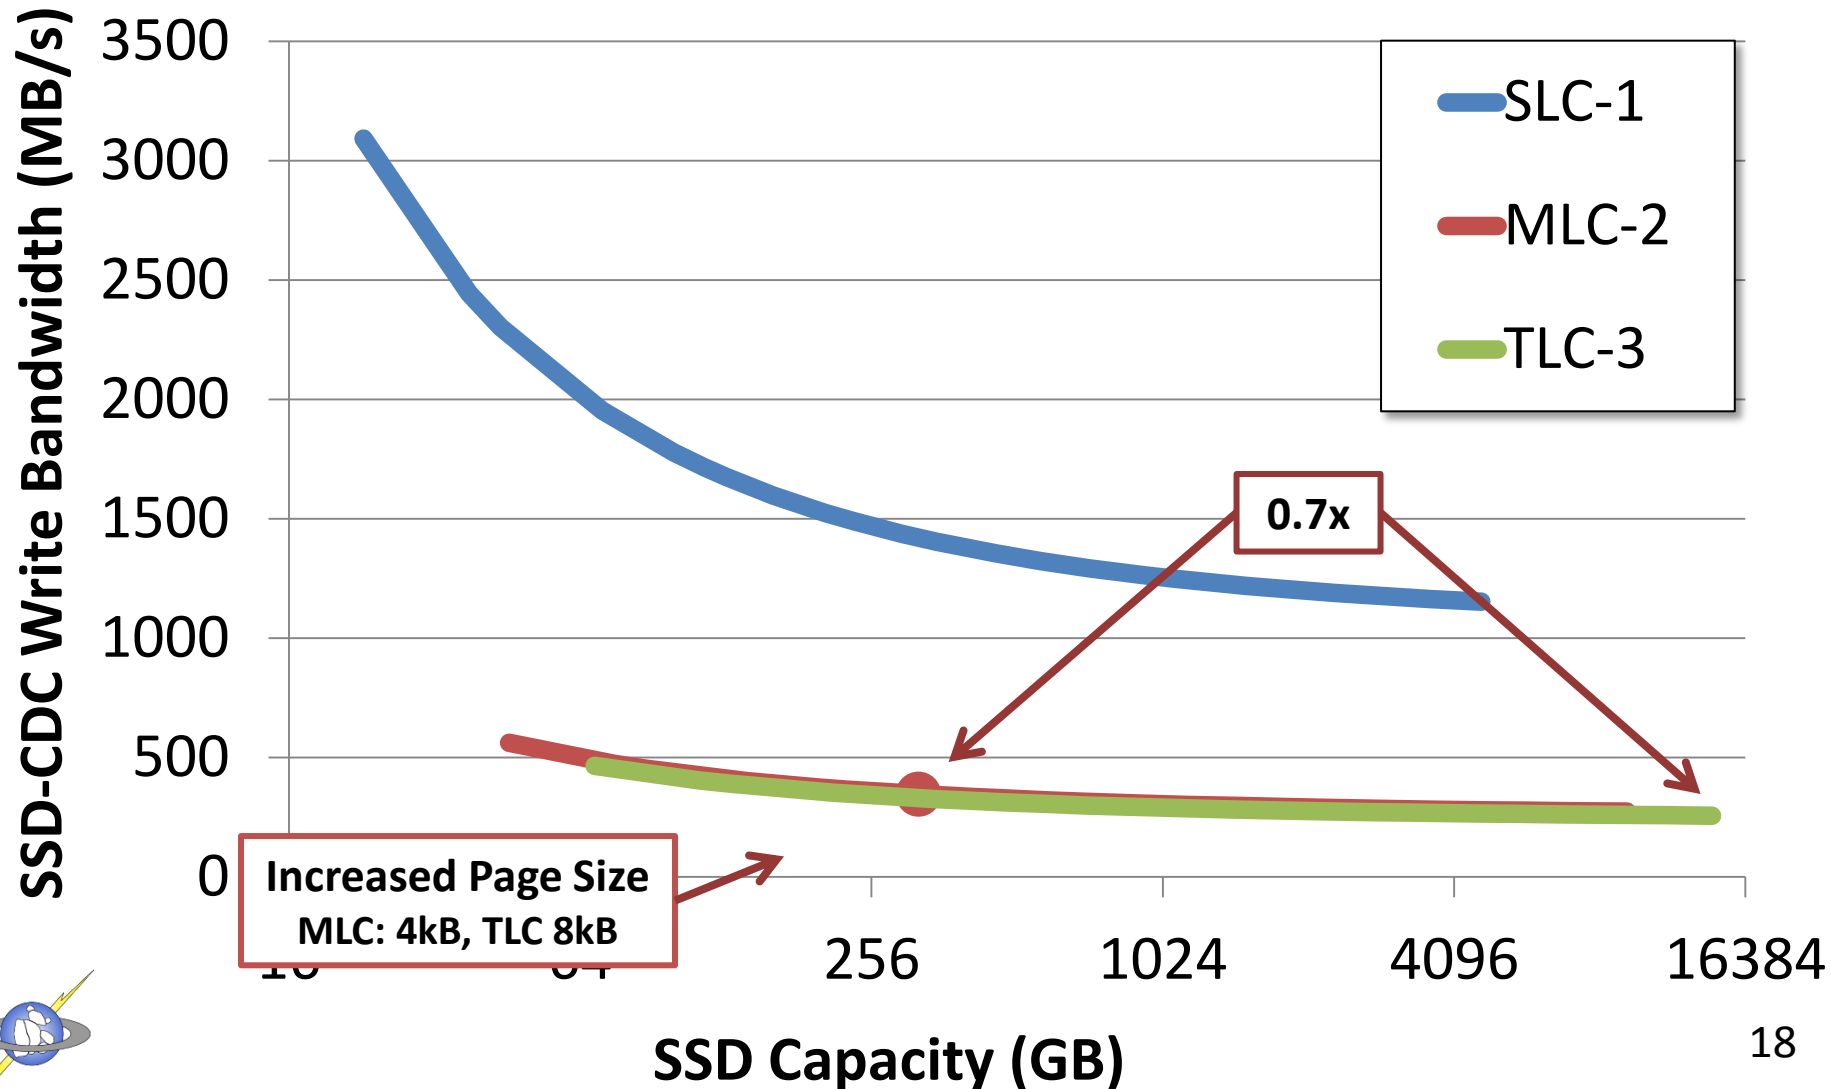
# The Metrics

---


- Capacity: 43x
- Latency: 2.6x
- Throughput



# Reduced Bandwidth




# IOPs - 512B Random Accesses



# The Metrics

---

- Capacity: 43x



- Latency: 2.6x



- Throughput: 0.7x, 0.4x



# Conclusion

---

- Chip Scaling: A Mixed Bag
  - Improved: Density and Cost
  - In Decline: Performance and Reliability
- SSDs: Not always a perfect replacement for disks
  - Do Get: High Capacity & High IOPs
  - Don't Get: Low Cost & Low Latency




# Questions?

---

## The Bleak Future of NAND Flash Memory

Laura M. Grupp, John D. Davis, Steven Swanson



Microsoft®

# NVSL

Non-volatile Systems Laboratory

# Research