

Empirical FTL Evaluation and Modeling

Thomas W. McCormick & David R. Kaeli
Northeastern University (NUCAR)





Background - Segments

- Flash Memory Markets
 - Client – Cost (\$/GB)
 - Enterprise – Speed/Endurance (IOPS, DWPD)
 - Embedded – Reliability (BER)

Background - Embedded Systems


- “Fixed function system”
 - Telecom, automotive, industrial control systems, medical equipment ...
 - 2009: 10B embedded processors (EETimes)
 - 2009: 300M PCs (Gartner), 3%
- Commonality: Flash Storage
 - Code & data



(Wikipedia – Creative Commons)

Motivation

- ~\$30B market (2014, iSupply) dominated by client (\$/GB) needs (Apple ~30%, iSupply)
- Smaller process, more bits per cells
 - Endurance & ECC
 - Retention
- Challenge: Maintain acceptable service-life for embedded systems





Lifetime & Write Amplification

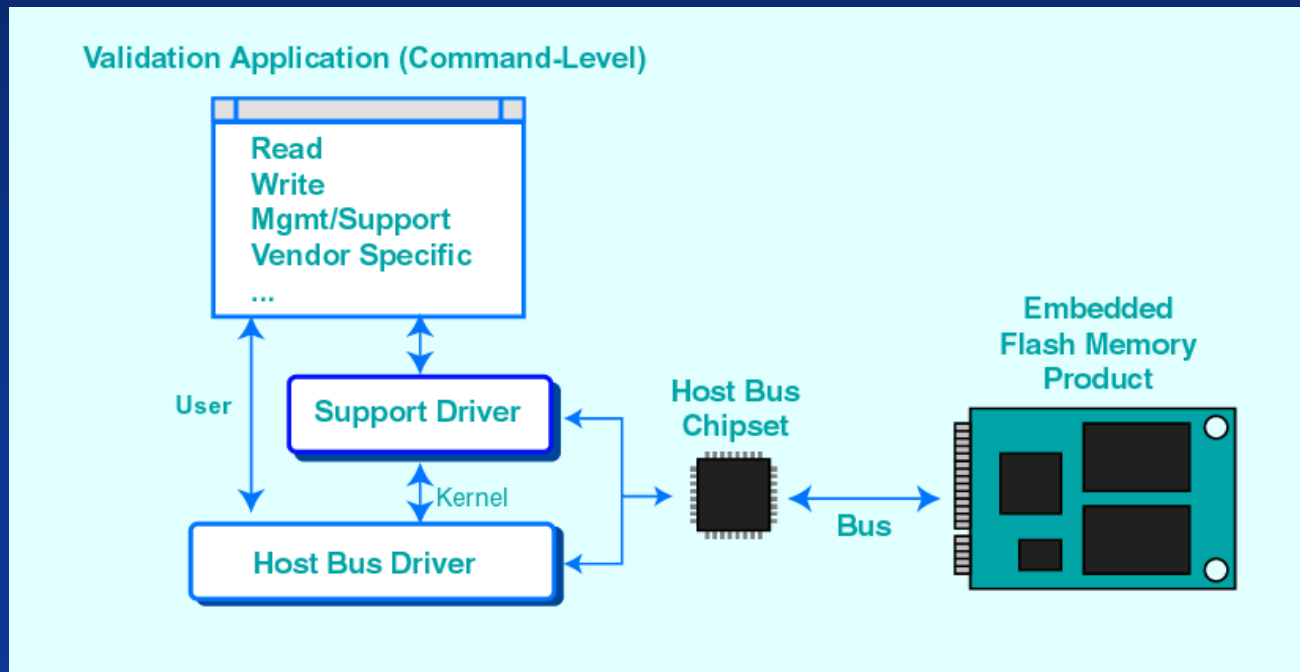
$$\text{Lifetime} \propto \frac{(\text{Capacity})(\text{Endurance})}{(\text{Write Amplification})}$$

$$\text{Write Amplification} = \frac{\text{Data Written to Flash}}{\text{Data Written by Host}}$$

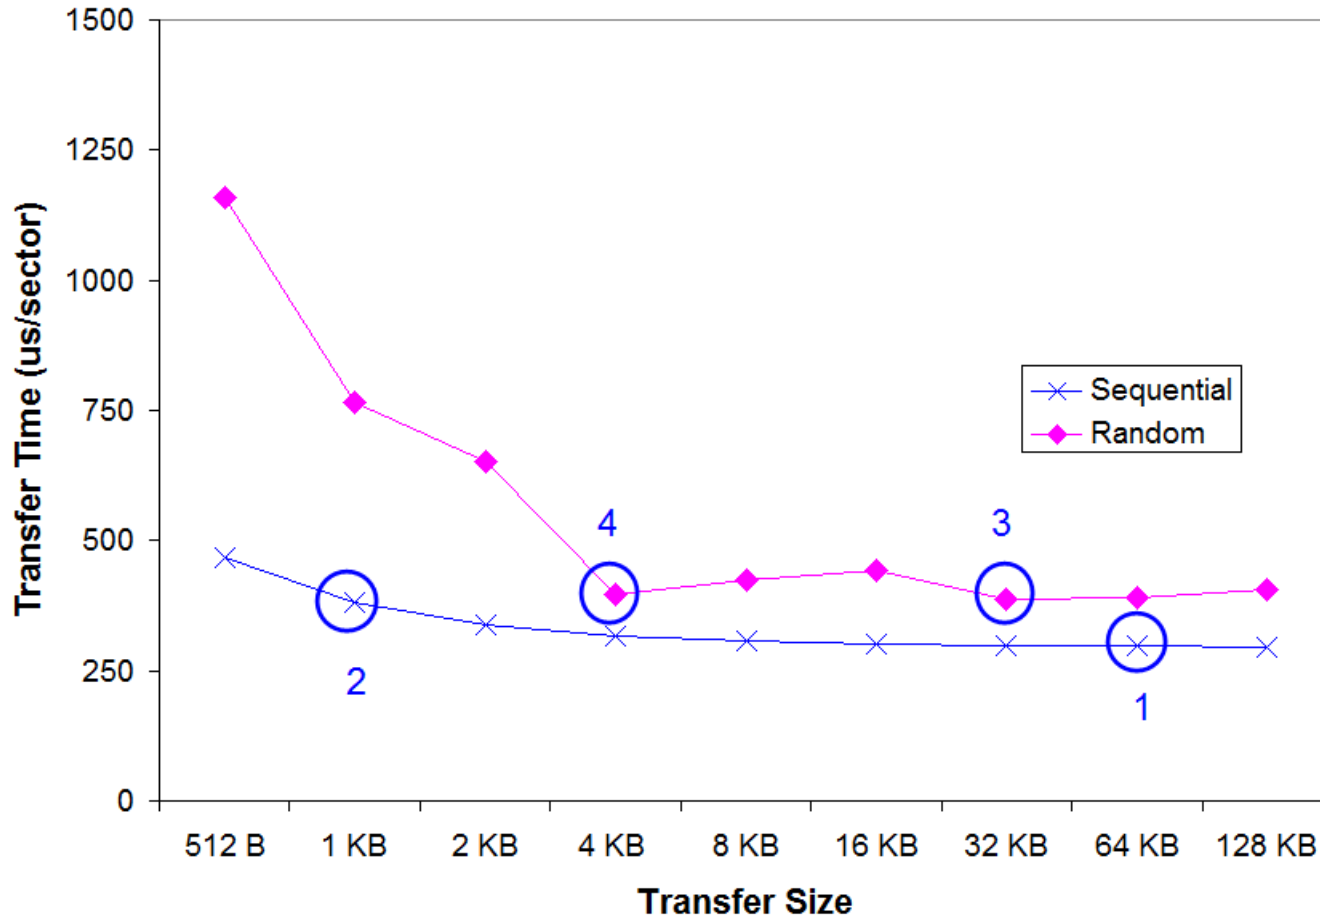
- Write amplification factor (WAF) is a function of Flash Translation Layer (FTL)
- FTLs are trade secrets. WAF measurements are unpublished

Technique

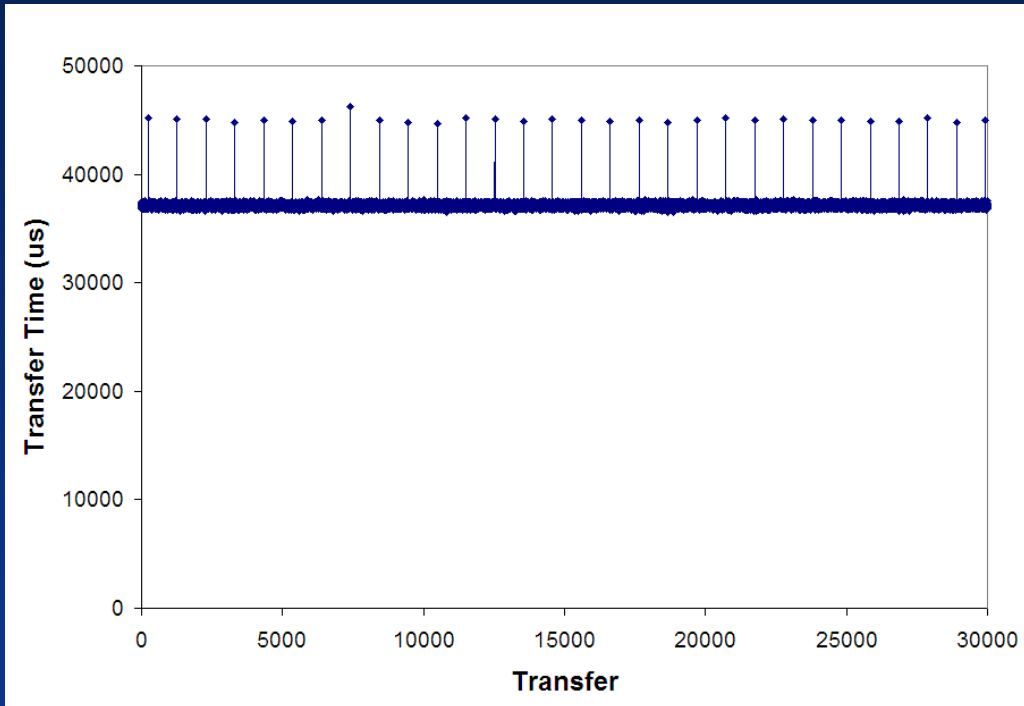
- Software application to performed fine-grained write performance measurements and characterize FTL and provide WAF measurements



Measurements



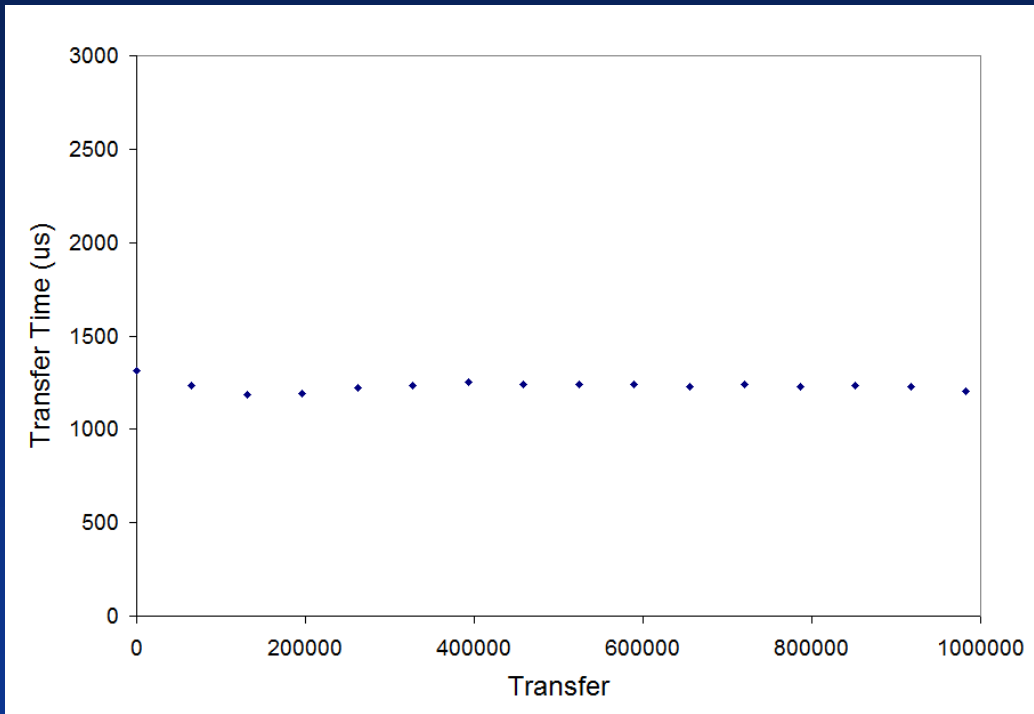
Measurements – Sequential, 64 KB



- Frequency: 64 MB
- Time: ~10 ms
- Block Erase: 5 ms
- Block Size: 2 MB
- Conclusion: 64 MB superblocs with 2 groups of 16 parallel erase blocks

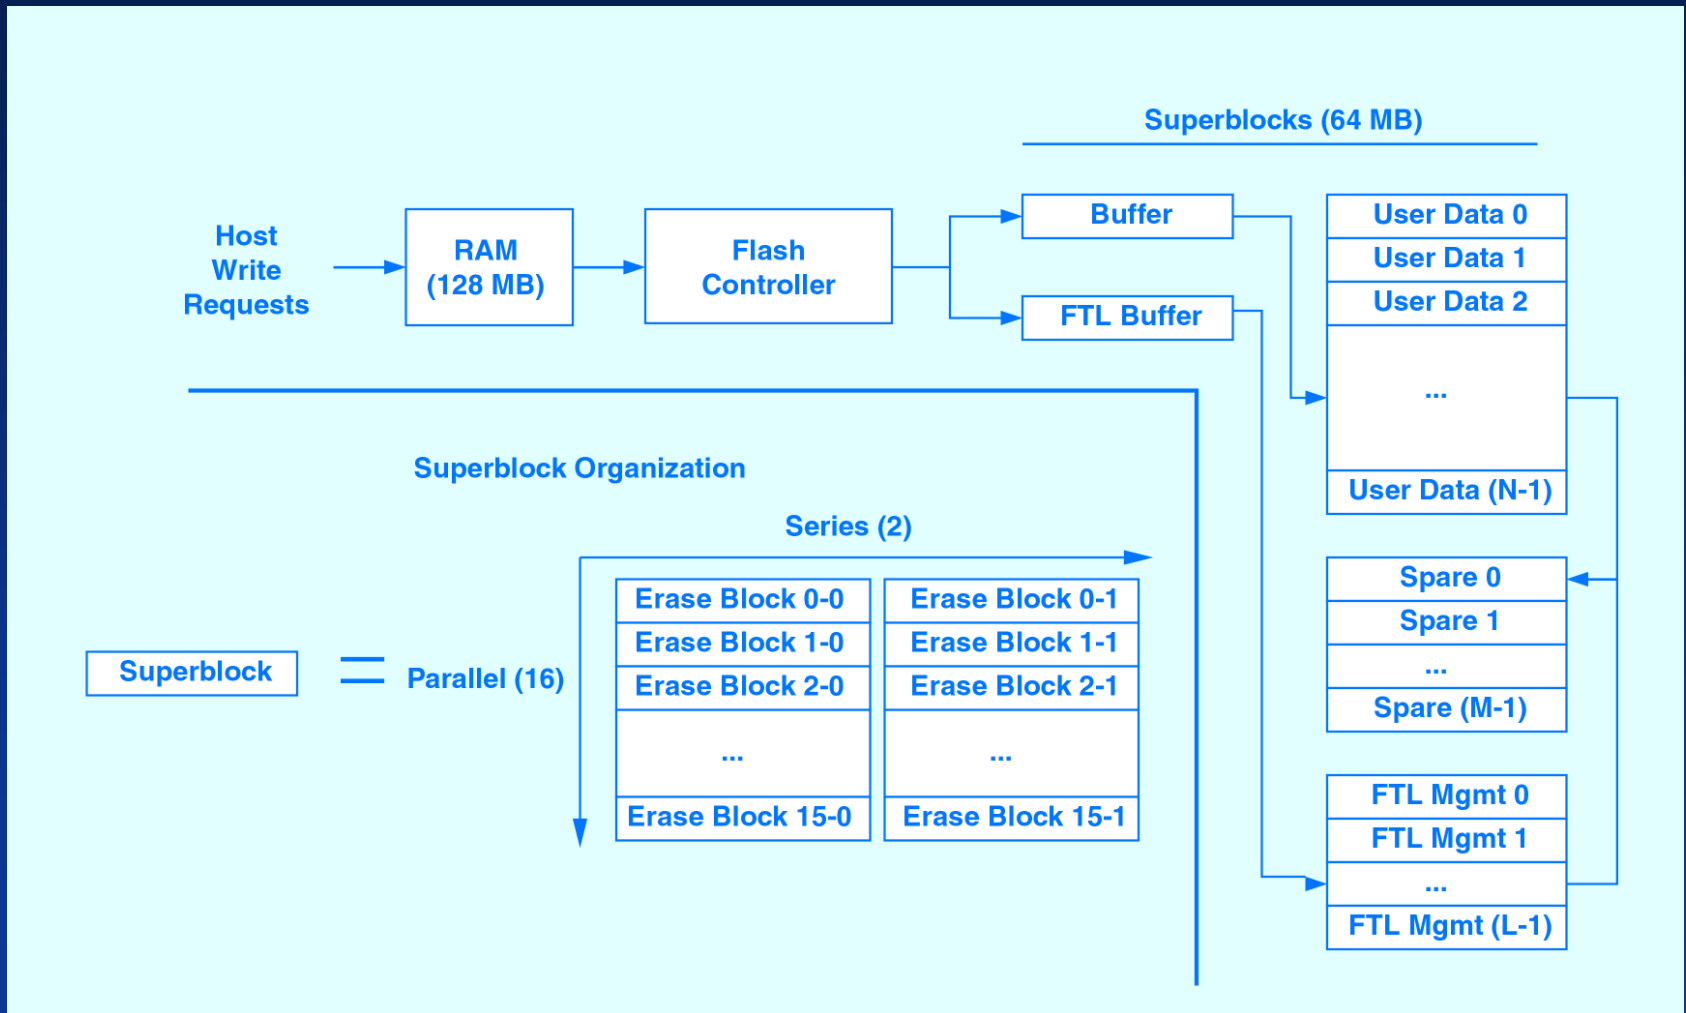


Measurements – Sequential, 1 KB



- Frequency: 64 MB
- Time: ~1200 us
- Page Program: 1400 us
- Page Size: 8 KB
- Conclusion: FTL overhead (256 KB/64 MB data)

FTL Model (Sequential)



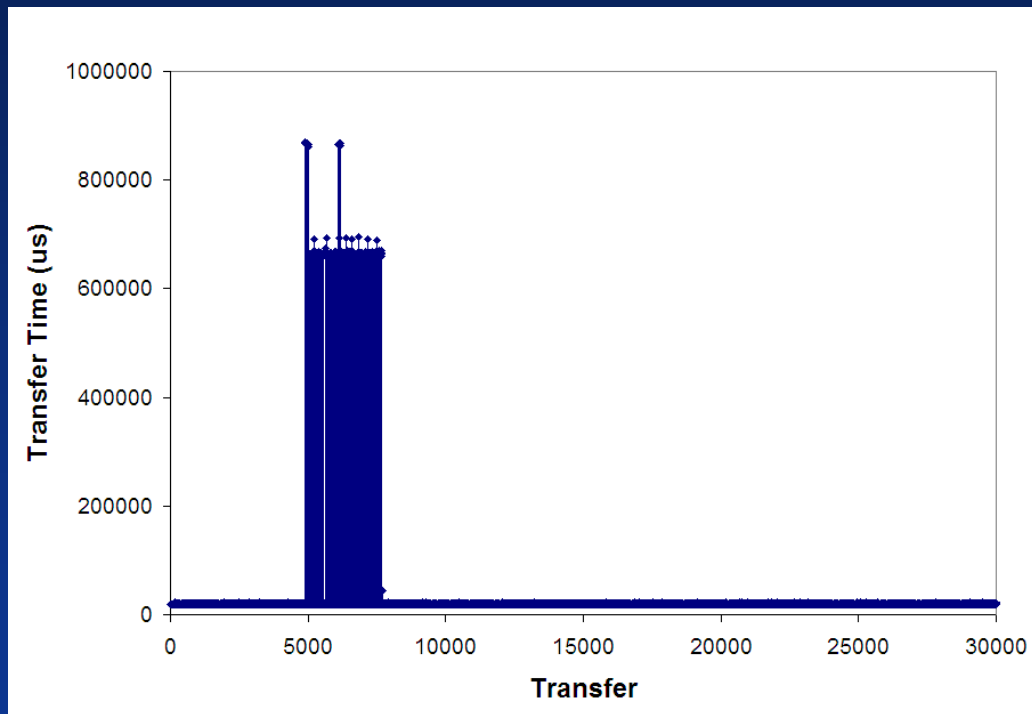
Write Amplification (Sequential)

$$WAF_{\text{Sequential}} = \frac{\text{Superblock Size} + \text{Management Data}}{\text{Superblock Size}}$$

$$WAF_{\text{Sequential}} = \frac{65536 \text{ KB} + 256 \text{ KB}}{65536 \text{ KB}} = 1.004$$

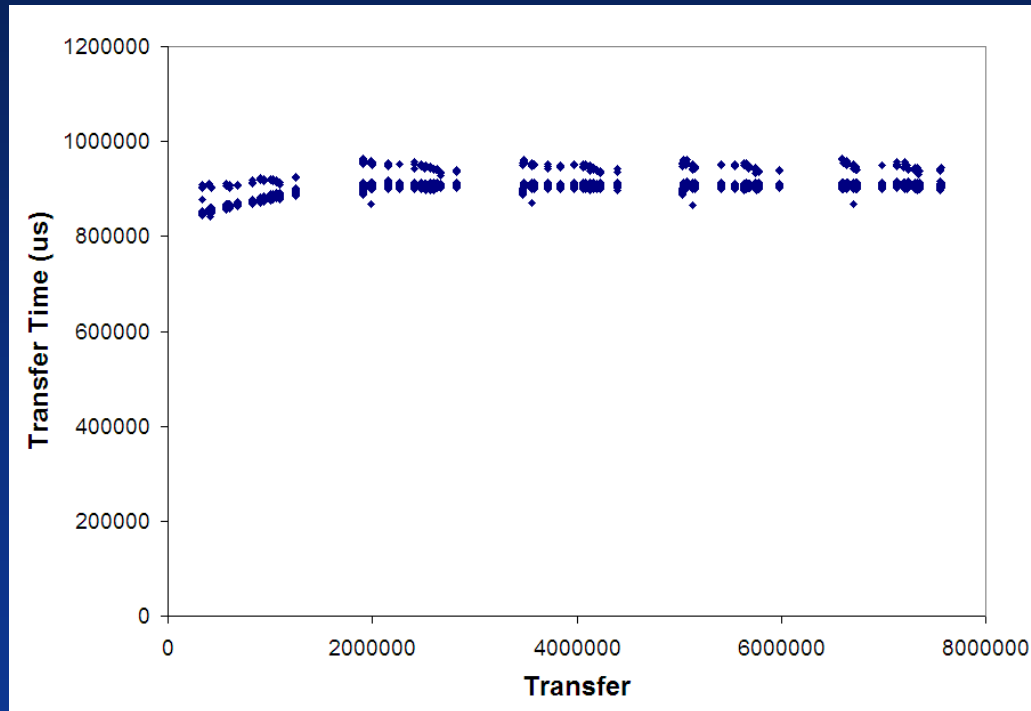


Measurements – Random, 32 KB





Measurements – Random, 4 KB

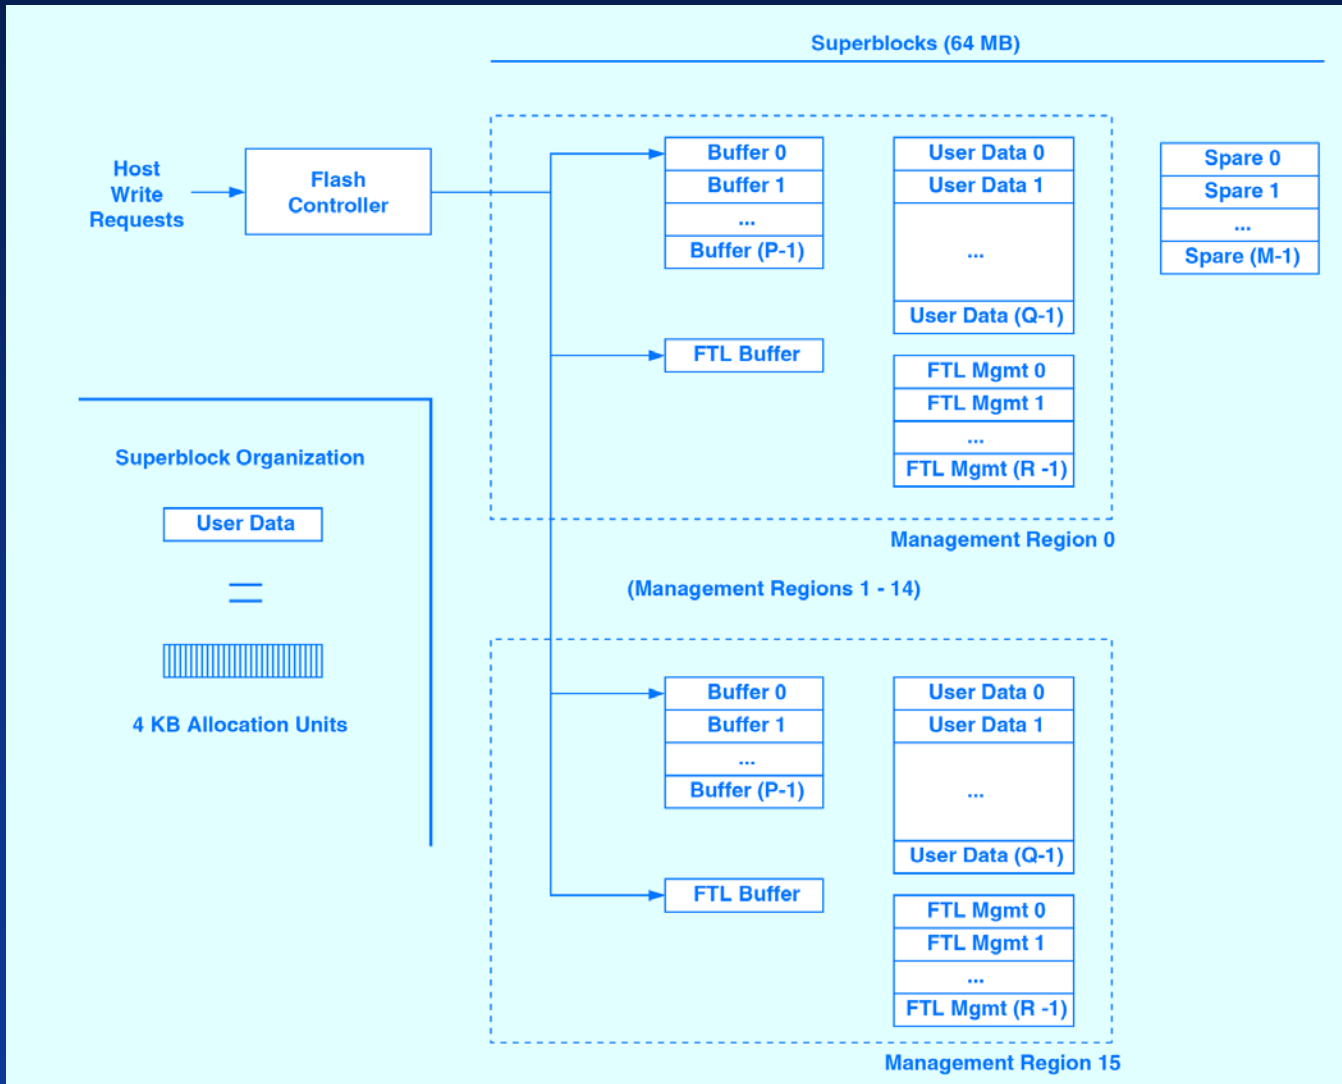




Modeling (Random)

- Frequency (Overall): 5.742 GB
- Events (Overall): 1908
- Events (Region): 119
- Regions: 16 (1908/119)
- Time: ~725 us
- Block Copy: 362.4 us
- Conclusion: 120 GB Drive partitioned into 16 independent management regions (7.5 GB) each with cache blocks (365 MB). Allocation units are 4 KB

FTL Model (Random)



Write Amplification (Sequential)

$$WAF_{\text{Random}} = \frac{\text{User Data} + \text{Buffer Data} + \text{Management Data}}{\text{Buffer Data}}$$

$$WAF_{\text{Random}} = \frac{(119 \times 64 \text{ MB}) + (5.7 \times 64 \text{ MB}) + (11.4 \text{ MB})}{5.7 \times 64 \text{ MB}} = 21.91$$



Summary & Conclusions

1. Technique for empirical FTL characterization
 - a. Application for fine-grained control and measurement
 - b. Bus specific functionality
2. Empirical block-mode FTL models for sequential and random modes of operations
 - a. First order (data) and second-order (FTL)
3. WAF Measurements
 - a. Sequential: 1.004
 - b. Random 21.91



More Embedded Sessions

- Flash Memory System Embedded Events:
 - Beer, Pizza, and Chat with the Experts
 - Tues 7:00 – 8:30 PM
 - Embedded SSD Applications (203-B)
 - Wed 3:20 – 4:25 PM
 - Embedded Flash Applications (204-B)
 - Wed 4:40 – 5:45 PM



Questions?

Tom McCormick - Chief Engineer/Technologist
Swissbit

tom.mccormick@swissbit.com

