

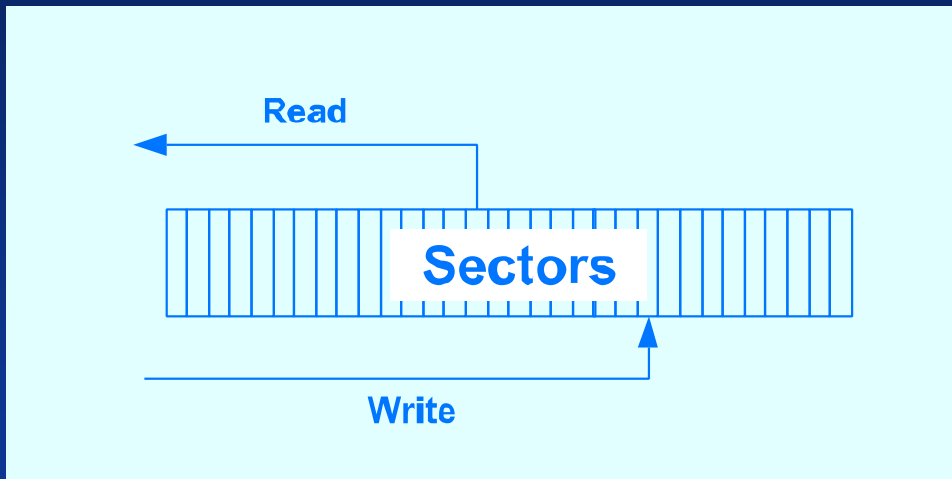
Modeling Flash Translation Layers to Enhance System Lifetime

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



Background – Flash Memory System

- Modern Flash Memory System (HDD -> SSD)

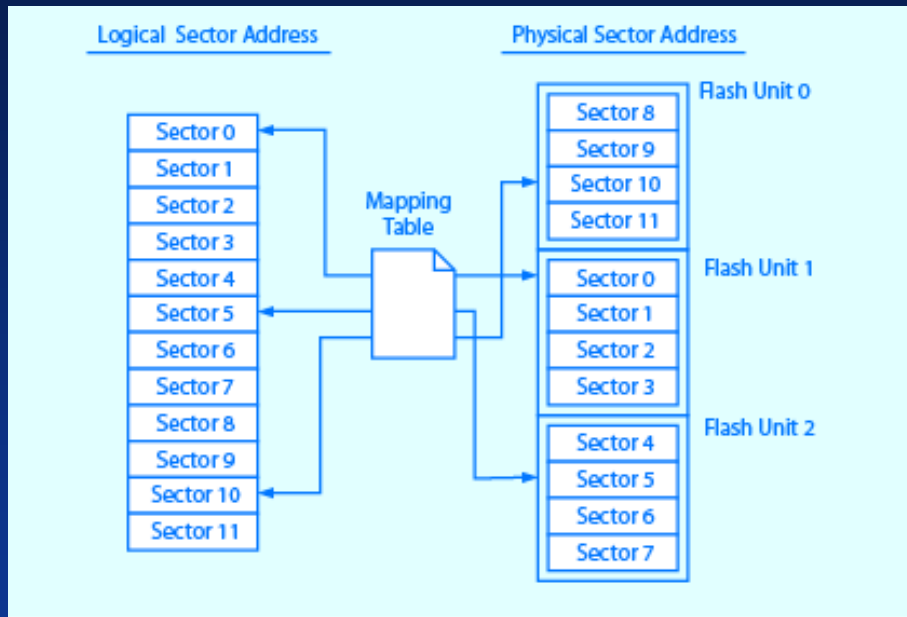




Background – Flash Challenges

1. Program-Then-Erase
2. Course-Grained Structures (Block & Page)
3. Finite Endurance

Background - FTL



- FTL Types: Block-Mode and Page-Mode



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



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

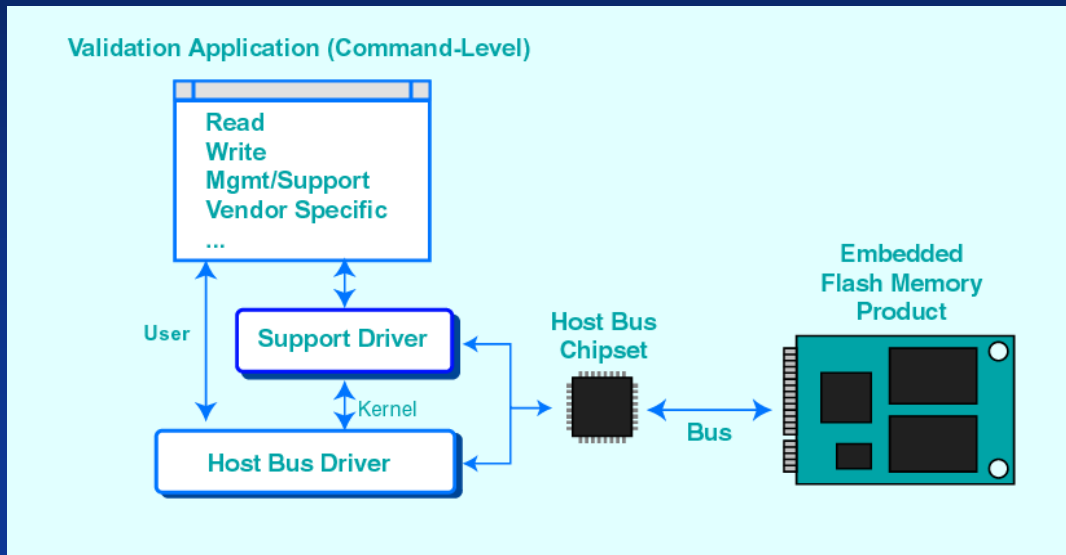
$$\text{Lifetime} \propto \frac{(\text{Capacity})(\text{Endurance})}{(\text{Write Amplification})}$$

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

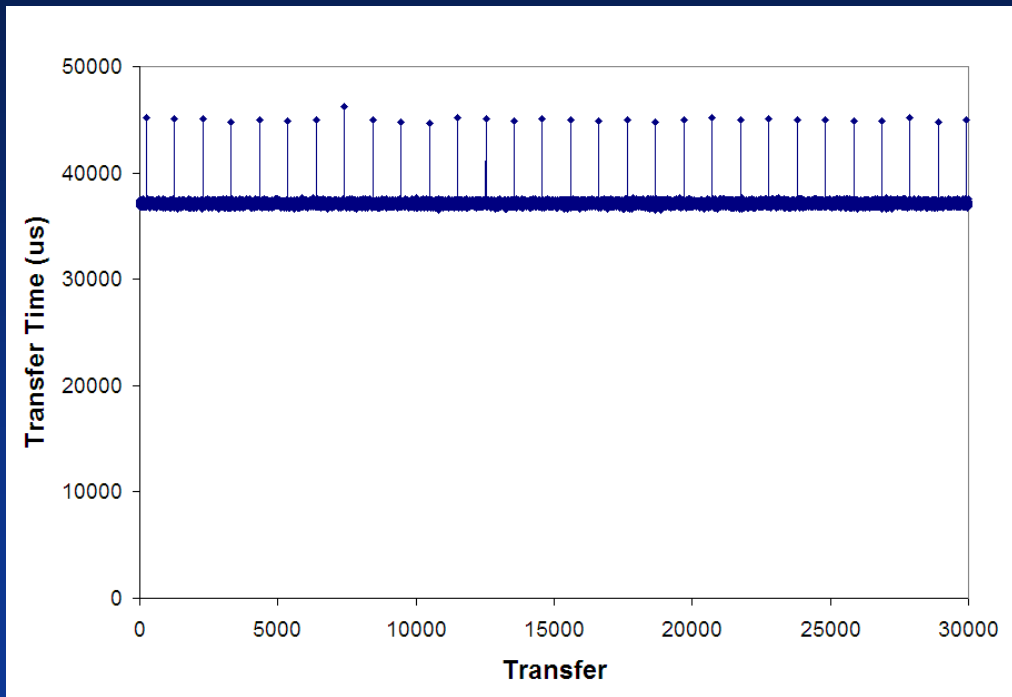
- Write amplification factor (WAF) is a function of Flash Translation Layer (FTL)
- FTLs are trade secrets. WAF measurements are unpublished

Technique

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



Measurements – Sequential, 64 KB



WAF Equation (Block-Mode FTL)

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

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

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

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

WAF Equation (Page-Mode FTL)

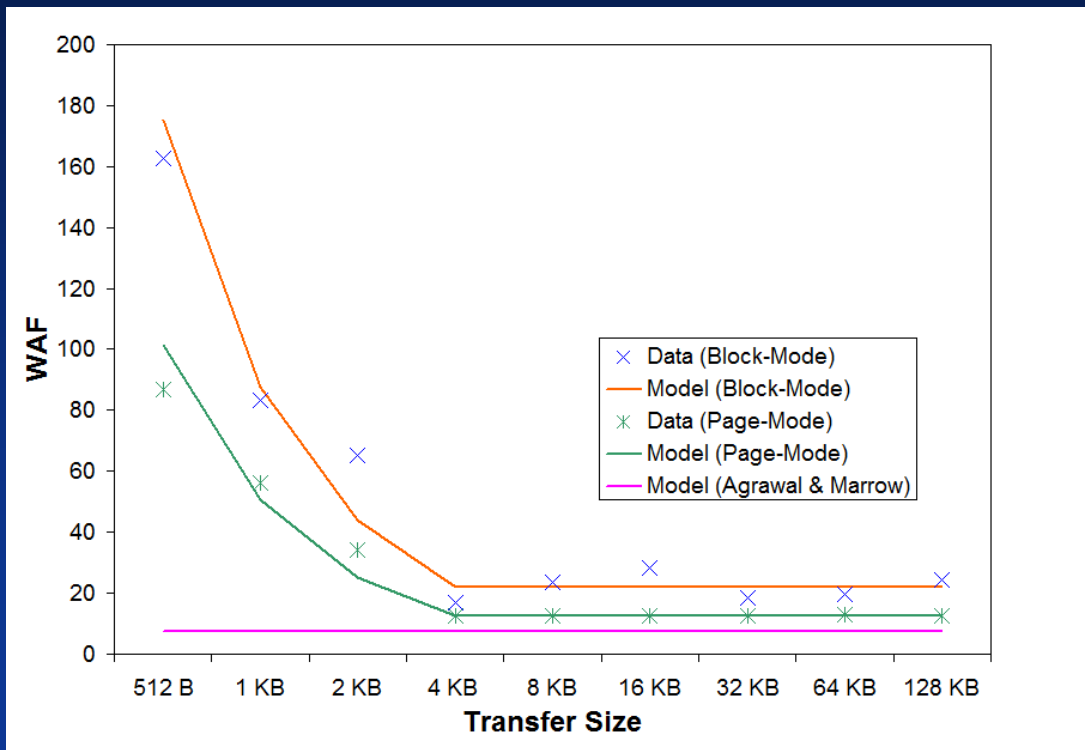
$$\text{WAF}_{\text{Sequential}} = \frac{\text{Superblock Size} + \text{Management Data}}{\text{Superblock Size}}$$

$$\text{WAF}_{\text{Sequential}} = \frac{65536 \text{ KB} + 128 \text{ KB}}{65536 \text{ KB}} = 1.002$$

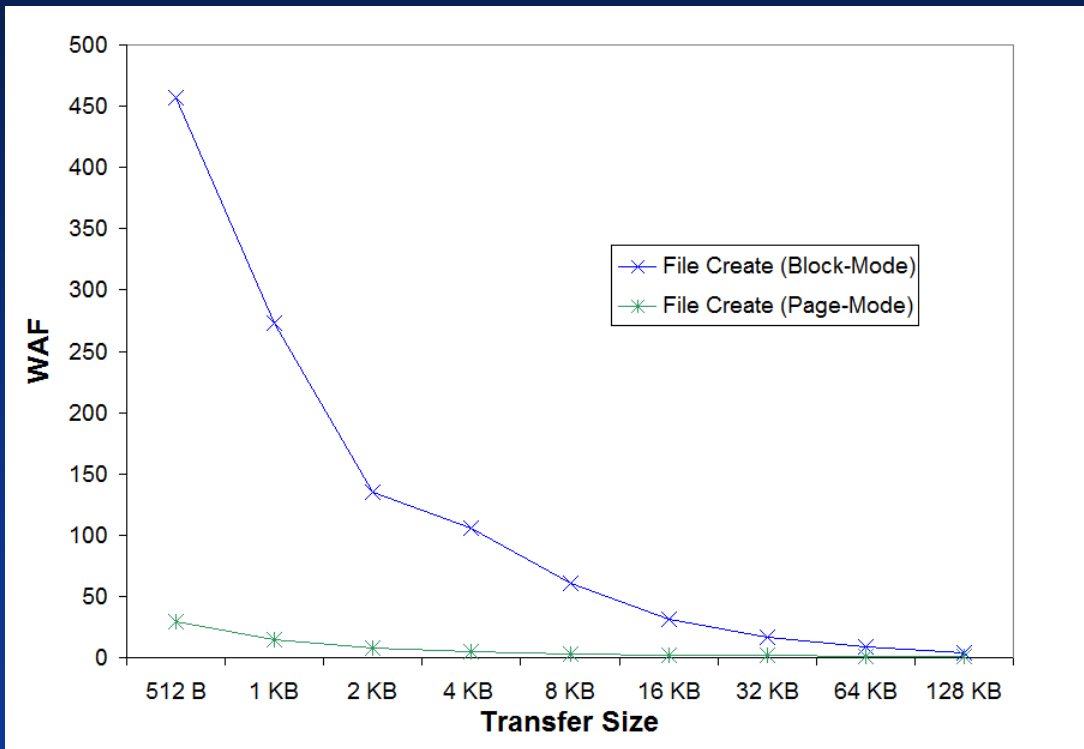
$$\text{WAF}_{\text{Random}} = \frac{\text{Superblock Size} + \text{Management Data}}{\text{Buffer Data}}$$

$$\text{WAF}_{\text{Random}} = \frac{(64 \text{ MB}) + (2.0 \text{ MB})}{5.206 \text{ MB}} = 12.66$$

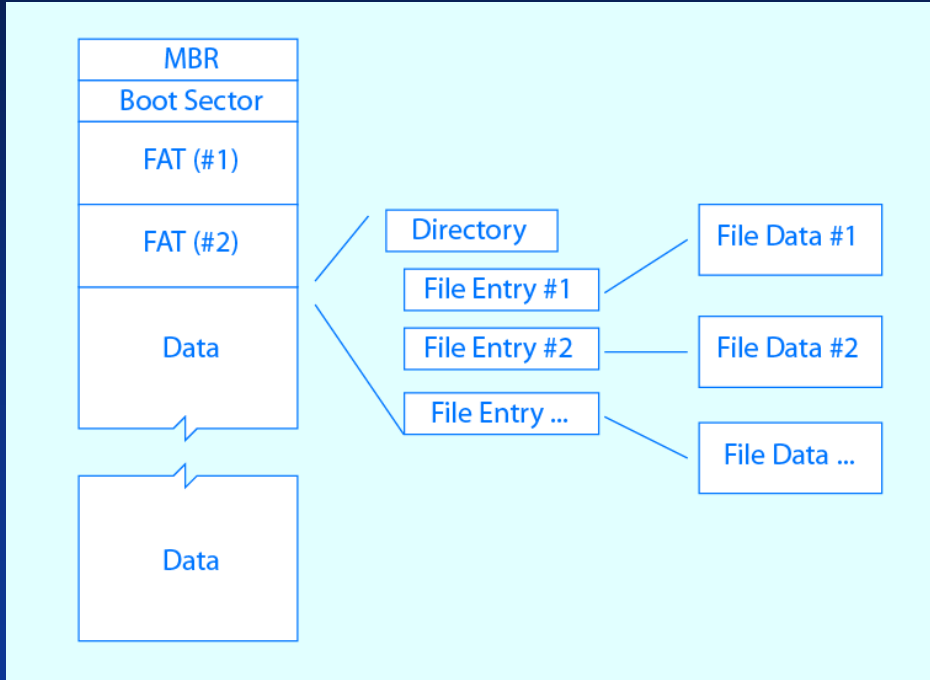
WAF Modeling (Random)



WAF (File System)



File System Operations (FAT)



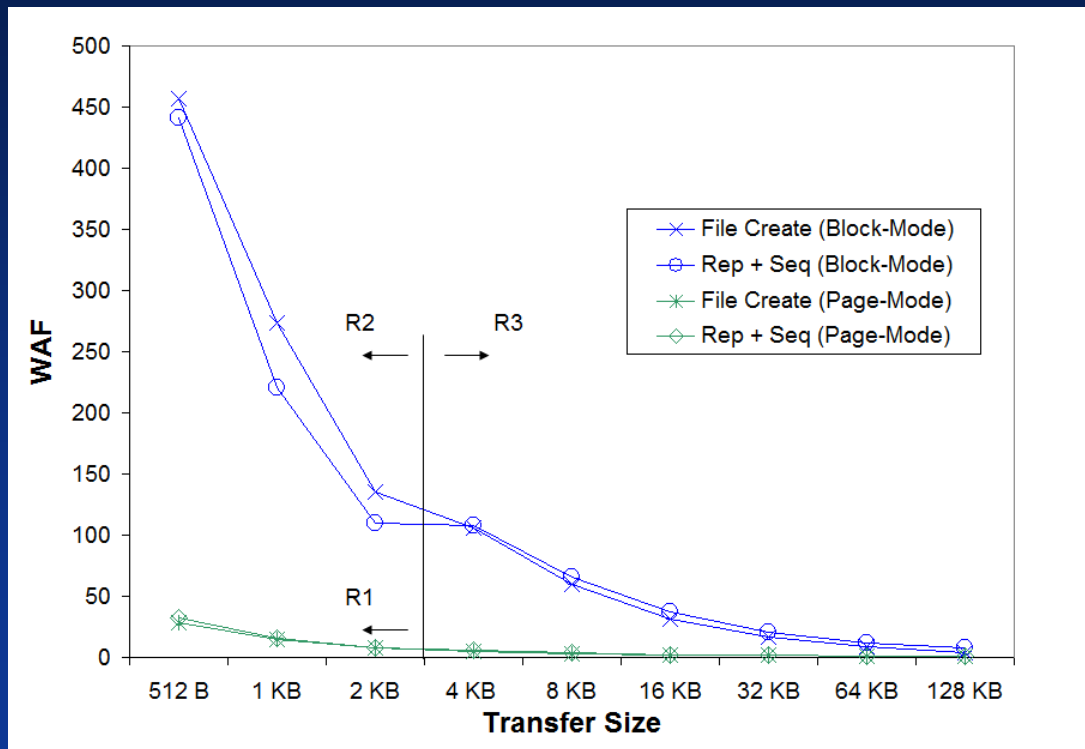
File Creation:

1. Directory (Repeated)
2. FAT #1 (Repeated)
3. FAT #2 (Repeated)
4. Data (Sequential)

File Deletion:

1. Directory (Repeated)
2. FAT #1 (Repeated)
3. FAT #2 (Repeated)

WAF (File System - Simulated)



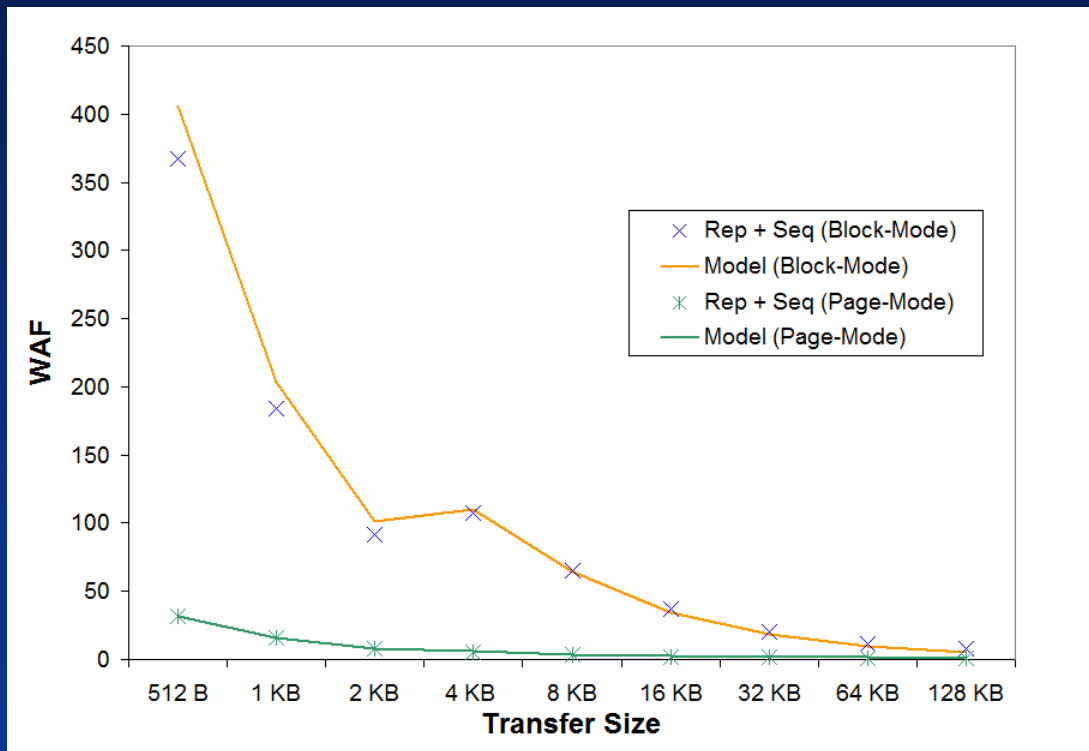
WAF Equation (Repeated)

$$WAF_{\text{Repeated}} = \frac{N_{\text{Repeated}} \cdot \text{Allocation Unit}_{\text{Repeated}} + \text{Allocation Unit}_{\text{Transfer}} + \text{Consolidation}}{N_{\text{Repeated}} \cdot \text{Size}_{\text{Repeated}} + \text{Size}_{\text{Transfer}}}$$

$$WAF_{\text{Block-Mode, 4KB}} = \frac{(3 \cdot 32 \text{ KB}) + (32 \text{ KB}) + (478.8 \text{ KB})}{3 \cdot 512 \text{ B} + 4 \text{ KB}} = 110.3$$

$$WAF_{\text{Page-Mode, 4KB}} = \frac{(3 \cdot 8 \text{ KB}) + (8 \text{ KB})}{3 \cdot 512 \text{ B} + 4 \text{ KB}} = 5.818$$

WAF Modeling (File Creation)



Summary & Conclusions

1. WAF Measurements (Block-Mode)
 - a. Sequential: 1.004
 - b. Random 21.91
 - c. Repeated (4K): 110.3
2. WAF Measurements (Page-Mode)
 - a. Sequential: 1.002
 - b. Random 12.66
 - c. Repeated (4K): 5.818
3. Page-Mode FTL incurs less WAF than Block-Mode FTL
4. Characterization of file-system operations needs to include repeated writes
5. FTL design should include consideration of repeated writes



More Embedded Sessions

- Flash Memory System Embedded Events:
 - Beer, Pizza, and Chat with the Experts
 - Tues 7:00 – 8:30 PM
 - Embedded Applications, Part 1 (203-B)
 - Wed 3:10 – 4:15 PM
 - Embedded Applications, Part 2 (204-B)
 - Wed 4:30 – 5:45 PM
 - Embedded Applications, Part 3 (301-B)
 - Thurs 8:30 – 9:35 AM



Questions?

Tom McCormick - Chief Engineer/Technologist
Swissbit

tom.mccormick@swissbit.com

