OS Dependencies on CPU SEFIs

NASA Goddard Space Flight Center

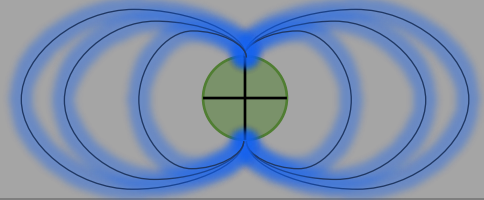

Radiation Effects & Analysis Group (REAG)

Dr. Seth Roffe
NASA GSFC
Code 561

Edward Wyrwas
SSAI

5/13/2024 - 5/17/2024

To be presented by Seth Roffe at SEE/MAPLD in San Diego, CA on May 16, 2024

Acronym List

OS: Operating System

SoC: System-on-Chip

SoCs: Systems-on-Chip

TFTP: Trivial File Transfer Protocol

MM: Matrix Multiplication

SEFI: Single-Event Functional Interrupt

CPU: Central Processing Unit

RAM: Random-Access Memory

TLYF: Test Like You Fly

TFTP: Trivial File Transfer Protocol

Key Points

- Linux operating system configuration variables
  - Which one has most affect?
  - How is this information useful to reliability/availability?
- Experimental controls
  - How was the test designed?
  - How do we control the experiment to know what we were finding?
- Where do we go from here?
  - How does this help us with standardization?
  - Future work

Problem Statement

- SoCs are complex
  - Hard to test
  - Hard to isolate complications and subsystems
- Current Fix: “Test as you Fly” or “Test Like You Fly” (TLYF)
  - Test the **exact** configuration and software as is going to be on mission
  - *Any* change, even compilation/optimization flags, could invalidate tests
- Need to find and standardize experimental controls
  - Control and isolate experimental variables to observe effects
  - Find correlations between isolated subsystems and reliability

SoCs: Systems-on-Chip

What do we already know?

- Linux system memory in two key parts
  - User space
  - Kernel space
- Hardware drivers exist in kernel space
  - Necessary to bridge the gap between hardware and OS
  - Huge part of the kernel are drivers
- Driver crash or failure can lead to kernel panic
  - Hangs entire system
  - ⇒ **Drivers leading cause of significant kernel failures**

	Address	Data
KERNEL SPACE	0000 0000	YYYY YYYY
	0000 0001	YYYY YYYY
	…	…
	XXXX XXXX	YYYY YYYY
USER SPACE	XXXX XXX(X+1)	YYYY YYYY
	…	…
	…	…
	…	…
	…	…
	…	…
	…	…
	FFFF FFFF	YYYY YYYY

*simplified for illustration

Linux Setup

- PYNQ-Z2 SoC
  - Artix-7 FPGA fabric
  - **Dual-core ARM Cortex-A9 processor**
- Generate Buildroot Linux images for **16 configurations**
  - Ensure minimalist Linux image with as many parts known as possible
  - Use same first-stage bootloader (boot.bin) for all configurations
- Initramfs
  - Load entire image in RAM via TFTP boot
  - Refresh each reboot to fresh image.
  - Operations and driver setup run with initscripts
- Expect drivers to have impact on system reliability

SoC: System-on-Chip
TFTP: Trivial File Transfer Protocol

OS Variables

L2 Cache	Operation	Num. Drivers	Loaded State
ON	MM	LARGE	LOADED
ON	MM	LARGE	UNLOADED
ON	MM	SMALL	LOADED
ON	MM	SMALL	UNLOADED
ON	IDLE	LARGE	LOADED
ON	IDLE	LARGE	UNLOADED
ON	IDLE	SMALL	LOADED
ON	IDLE	SMALL	UNLOADED
OFF	MM	LARGE	LOADED
OFF	MM	LARGE	UNLOADED
OFF	MM	SMALL	LOADED
OFF	MM	SMALL	UNLOADED
OFF	IDLE	LARGE	LOADED
OFF	IDLE	LARGE	UNLOADED
OFF	IDLE	SMALL	LOADED
OFF	IDLE	SMALL	UNLOADED

500×500

Sleep 10 sec

256 Drivers Built-In

170 Drivers Loaded

126 Drivers Built-In

10 Drivers Loaded

MM: Matrix Multiplication

Methodology

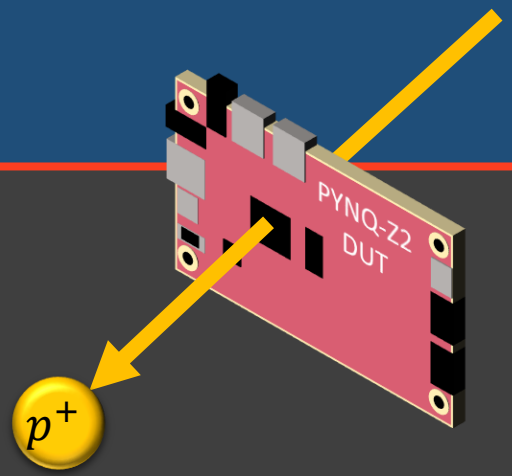

- MGH Proton Therapy Clinic
  - 200 MeV p^+
  - $\sim 10^8 \frac{p^+}{cm^2 s}$ flux
- System booted and stable between each run
  - Power cycle via web-powered power switch
  - Image pulled from host computer via TFTP boot in u-boot
  - Beam only on when operations began for fluence tracking
- ~15 runs for each configuration

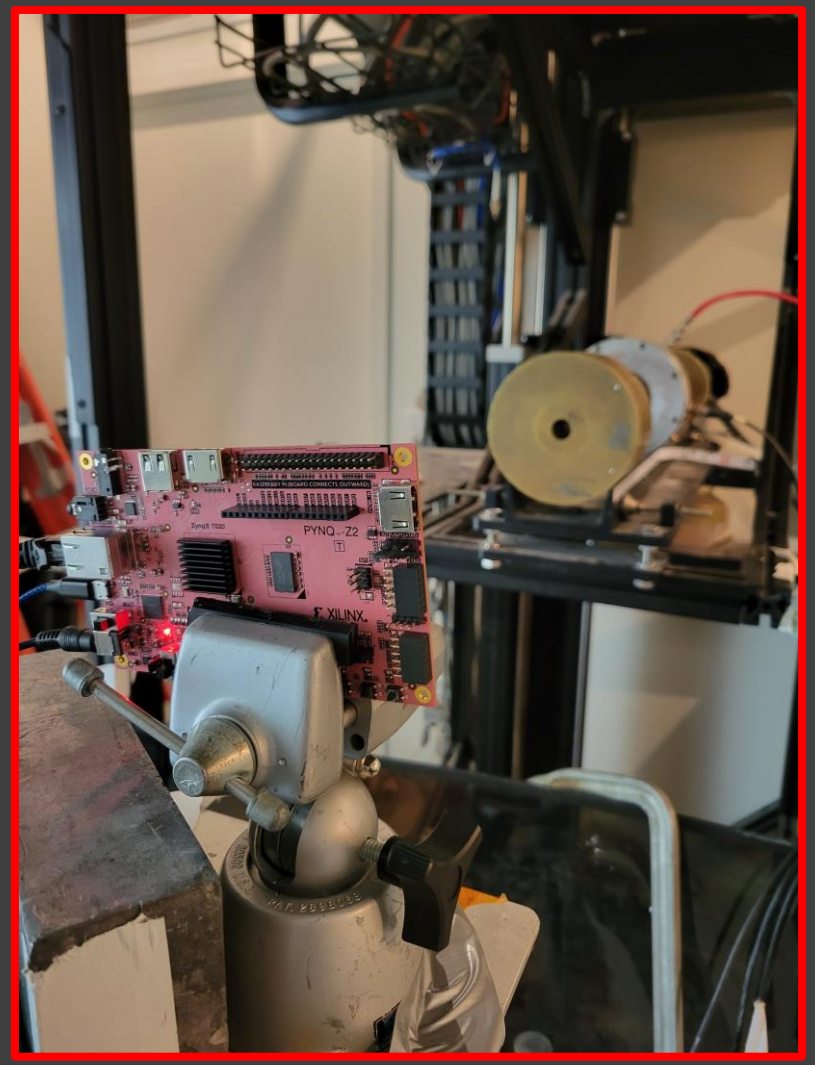

TFTP: Trivial File Transfer Protocol

Experimental Setup

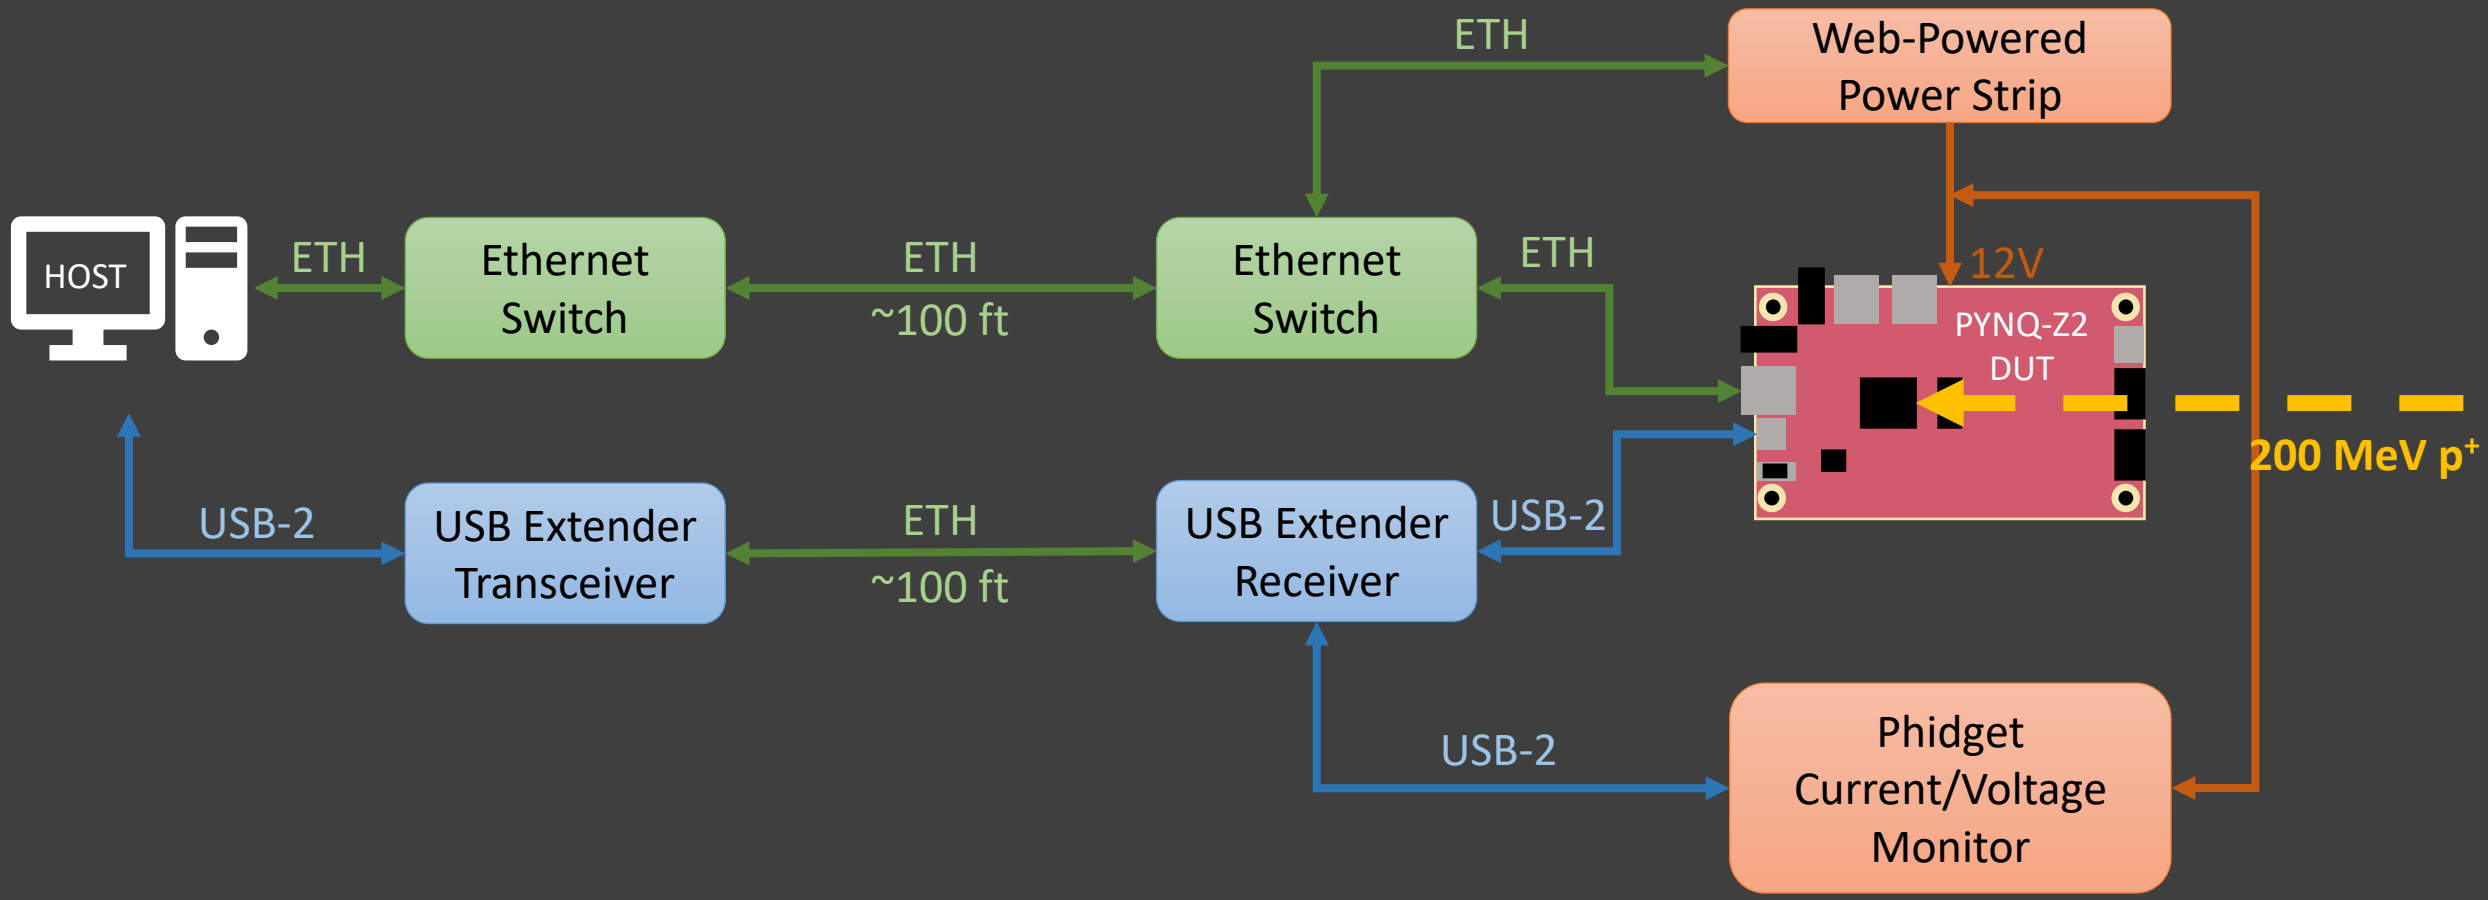

Mean Fluence to Failure (Cache ON)

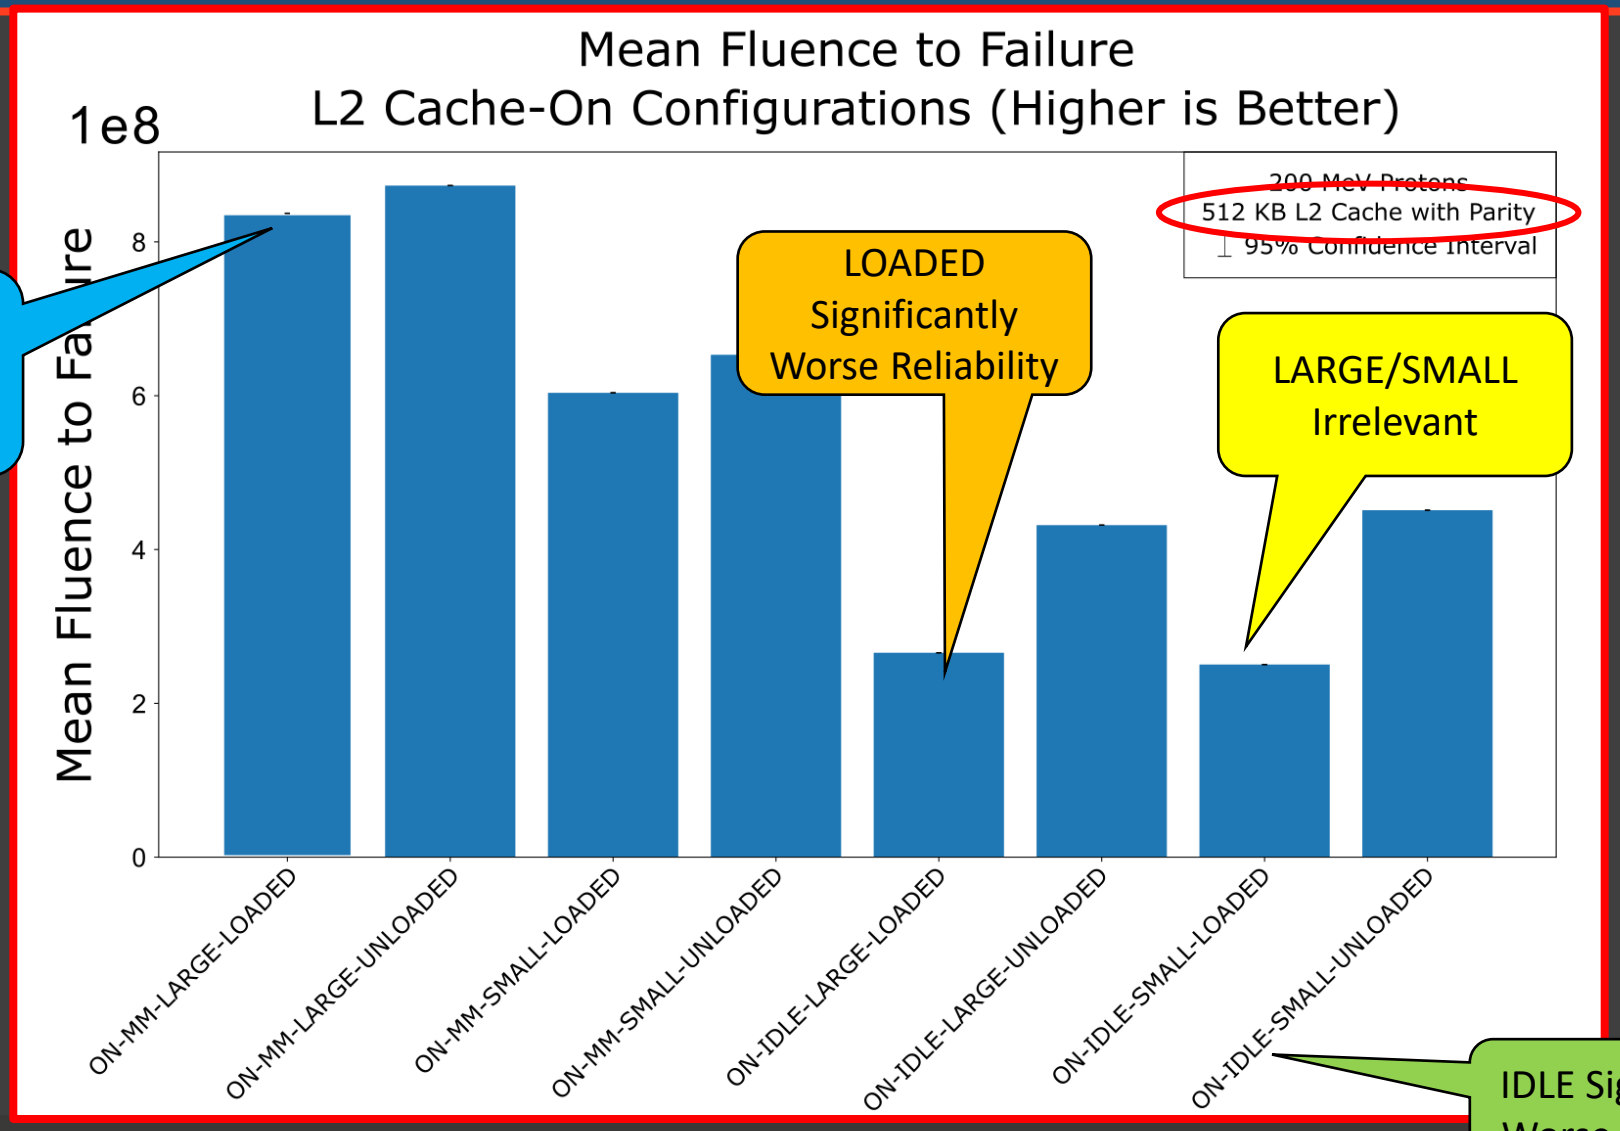

95% CI constrained by event number of 1

LOADED Significantly Worse Reliability

LARGE/SMALL Irrelevant

IDLE Significantly Worse Reliability

Mean Fluence to Failure (Cache OFF)

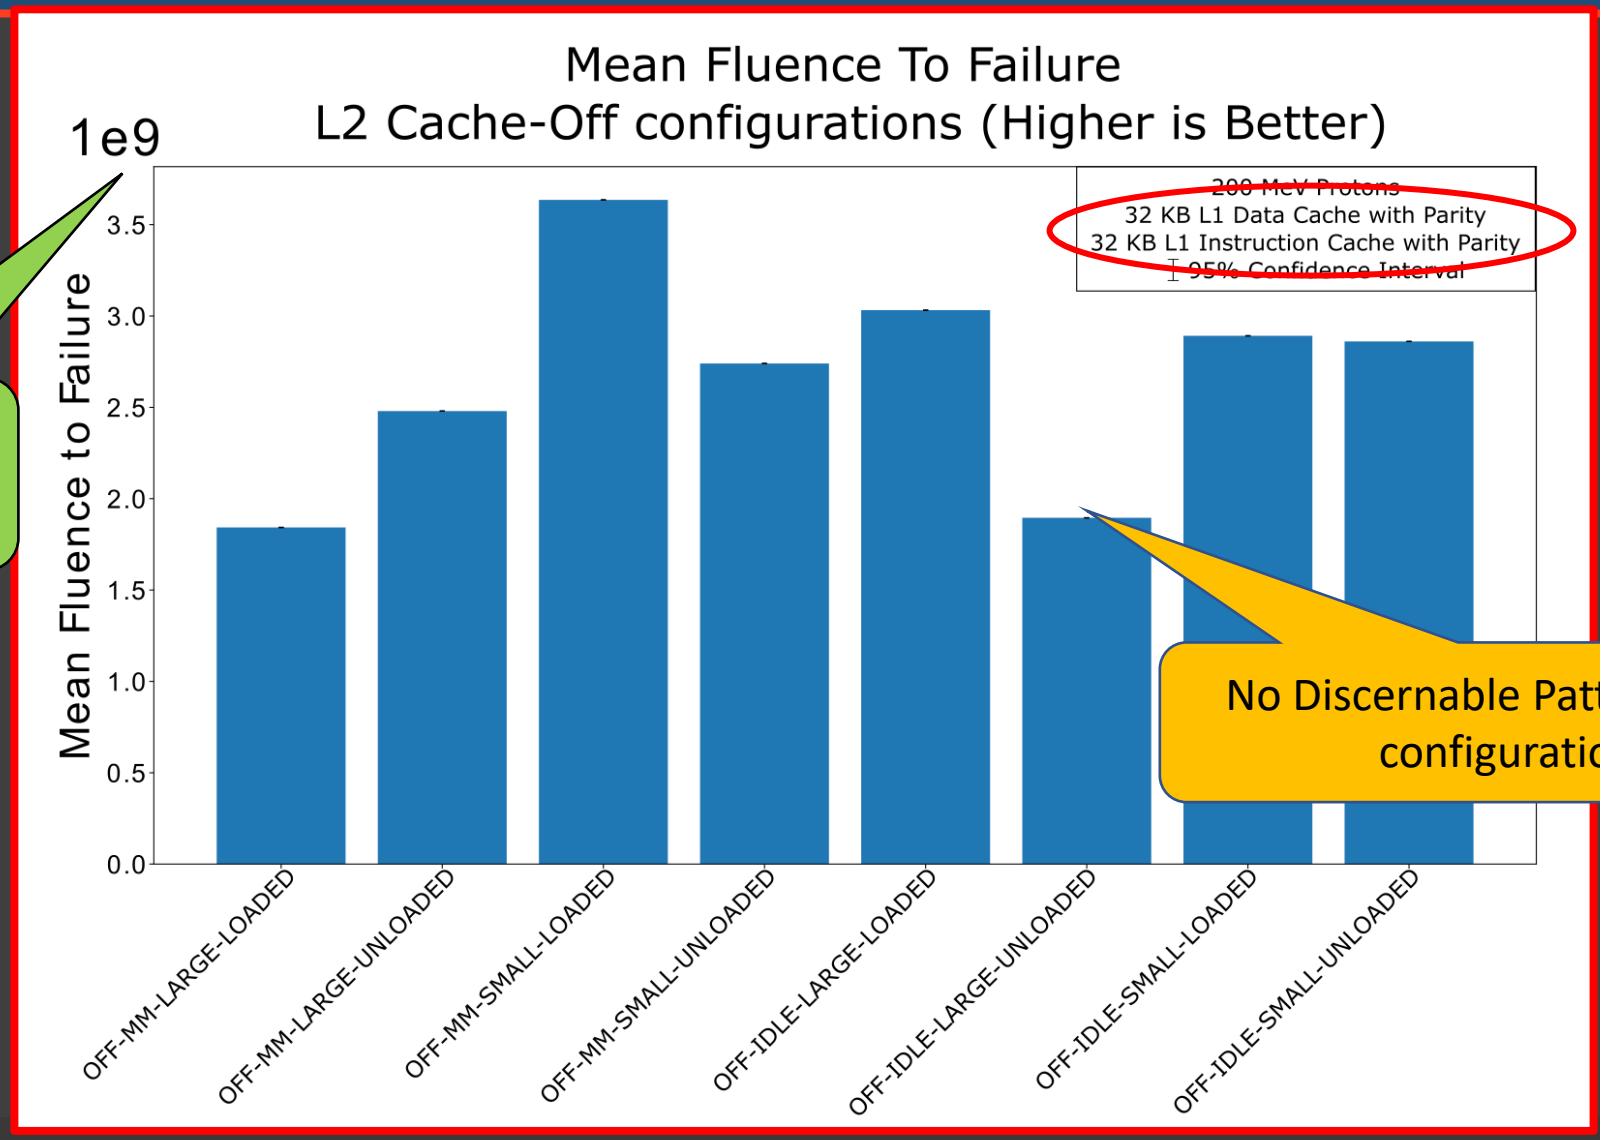

Significantly more reliable than ON

No Discernable Pattern in OFF configurations

Extrapolations and Implications

- Reliability/Available more **easily understood** when L2-Cache is ON
- LARGE/SMALL show **no significant difference**
  - May be a sample size problem
- Drivers **do have an effect** via LOADED with insmod
  - Unload unused drivers when not needed
- IDLE **significantly more vulnerable** to SEFI than MM
  - Critical data flushed in L2 when math is being performed
  - **Current operations have an impact on the reliability**
- **Understanding the memory hierarchy is critical**

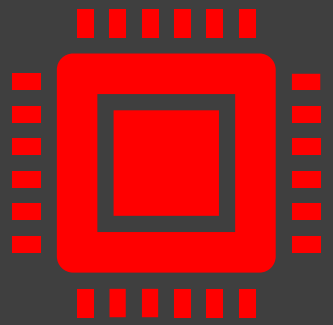

MM: Matrix Multiplication
SEFI: Single-Event Functional Interrupt

Some Extra Interesting Results

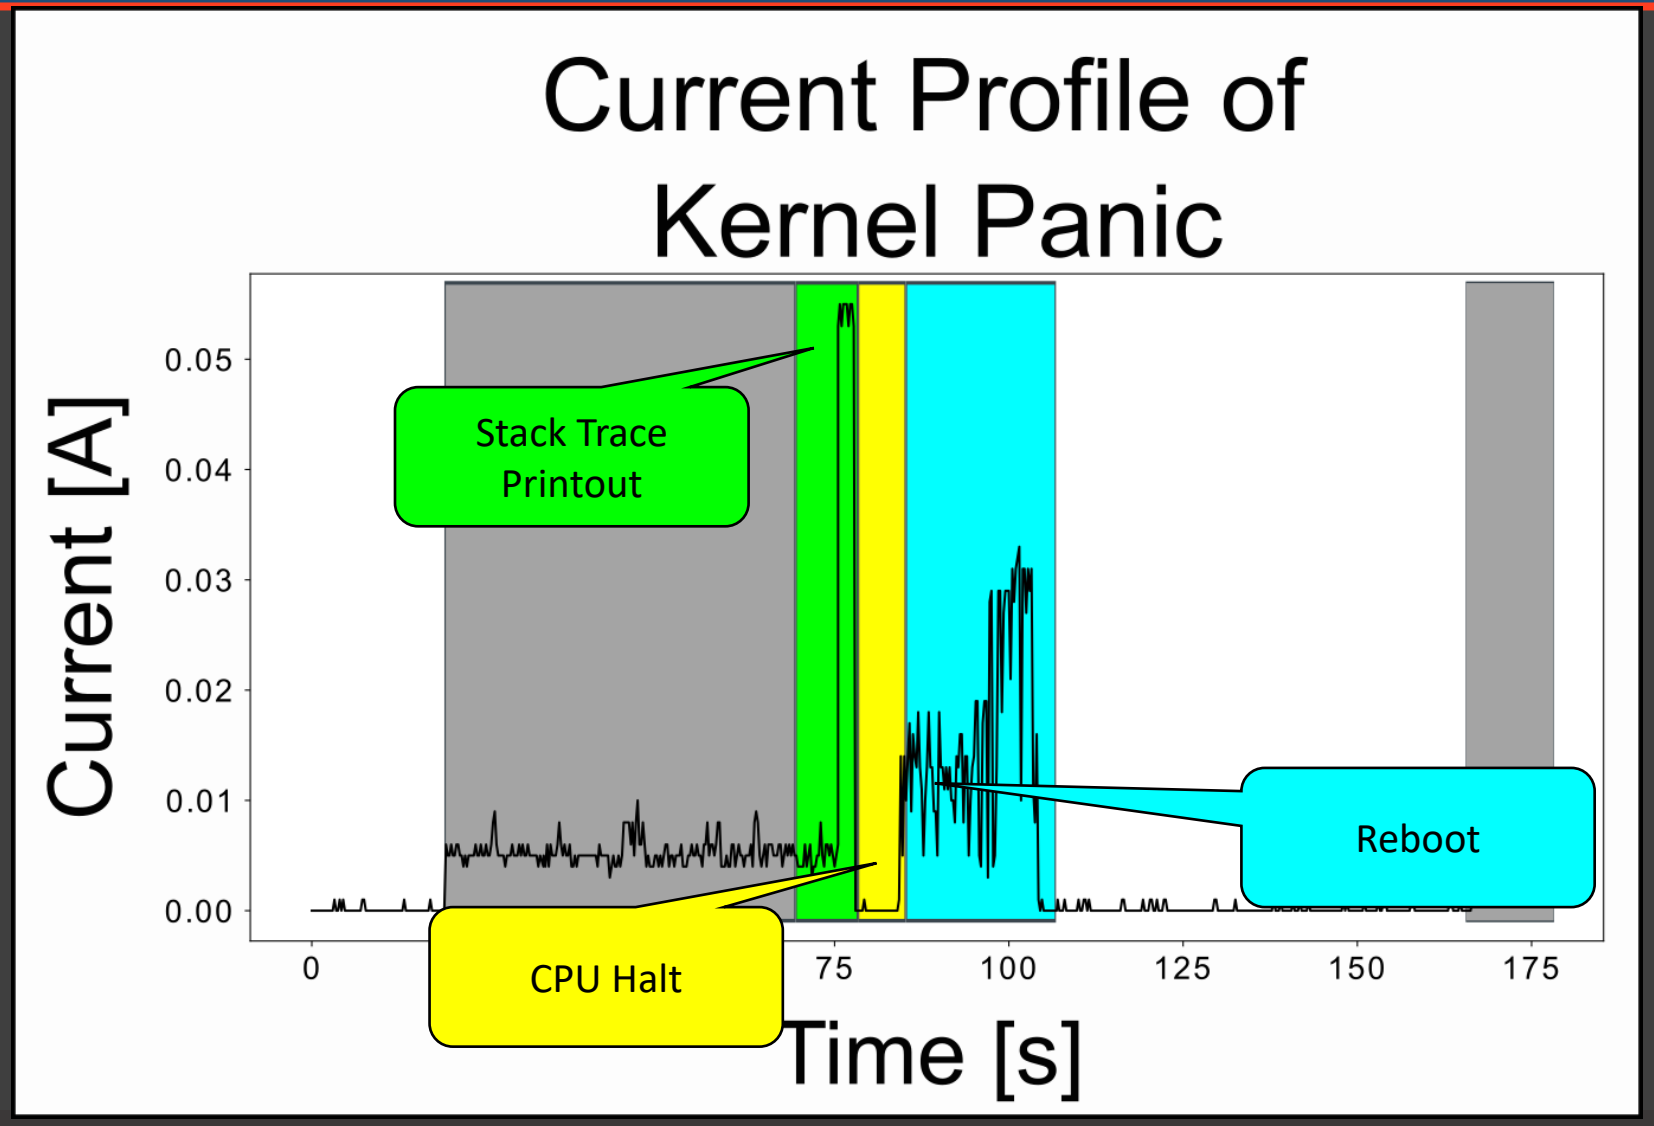

Microlatchup

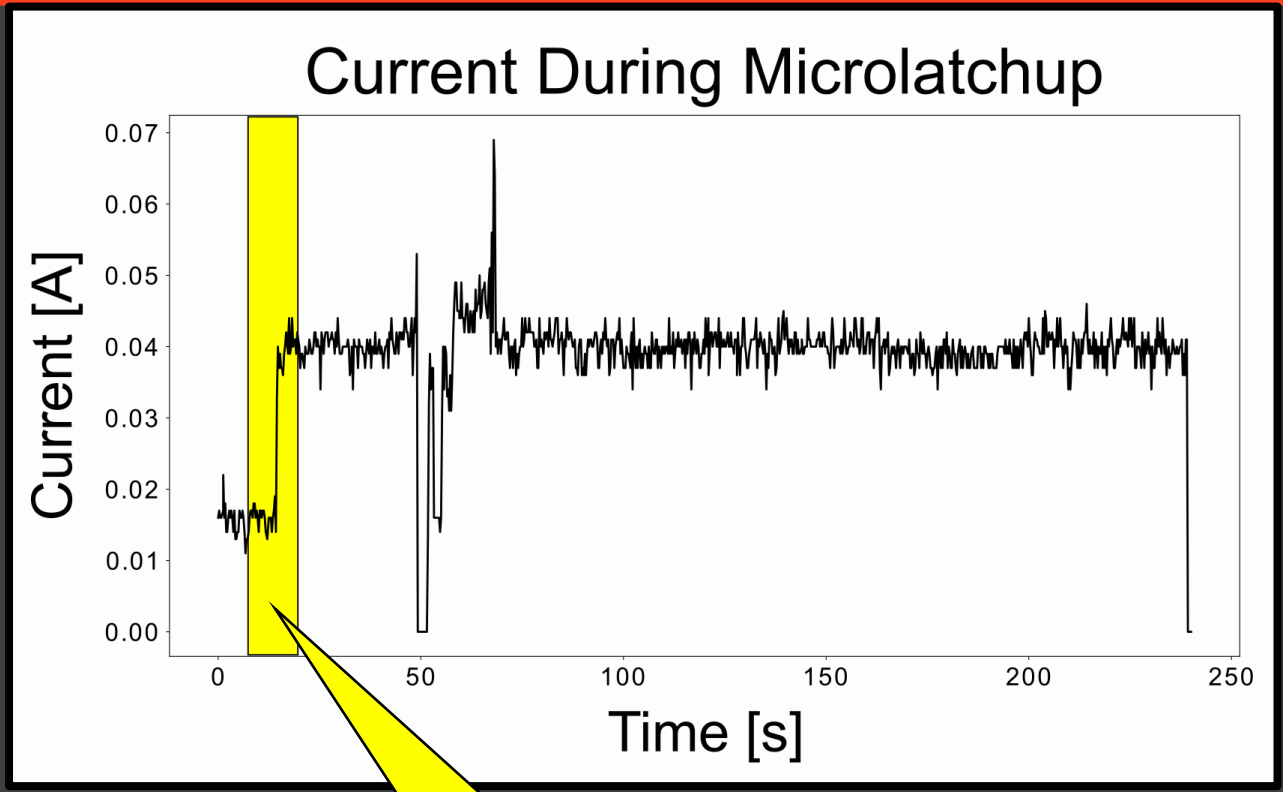

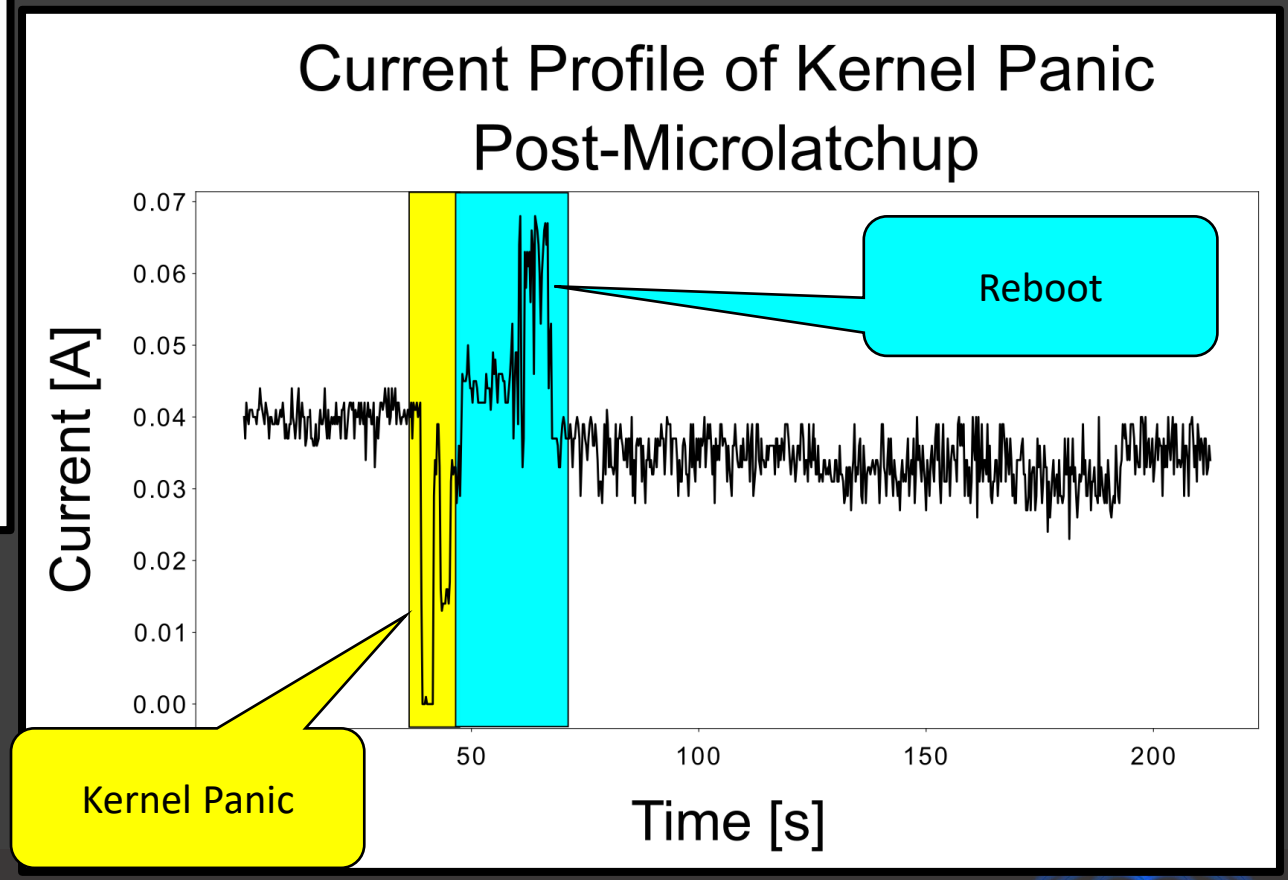

Conclusions and Future Work

- Things are hard
  - Many variables to keep track of
  - TLYF is not as effective as we may think
- Replication for better statistics
- Extend to other Architectures
  - RISC-V
  - GPUs and other Accelerators
- Optimize test procedures
  - Design specific TLYF to ensure coverage
  - Compile dependencies of a system

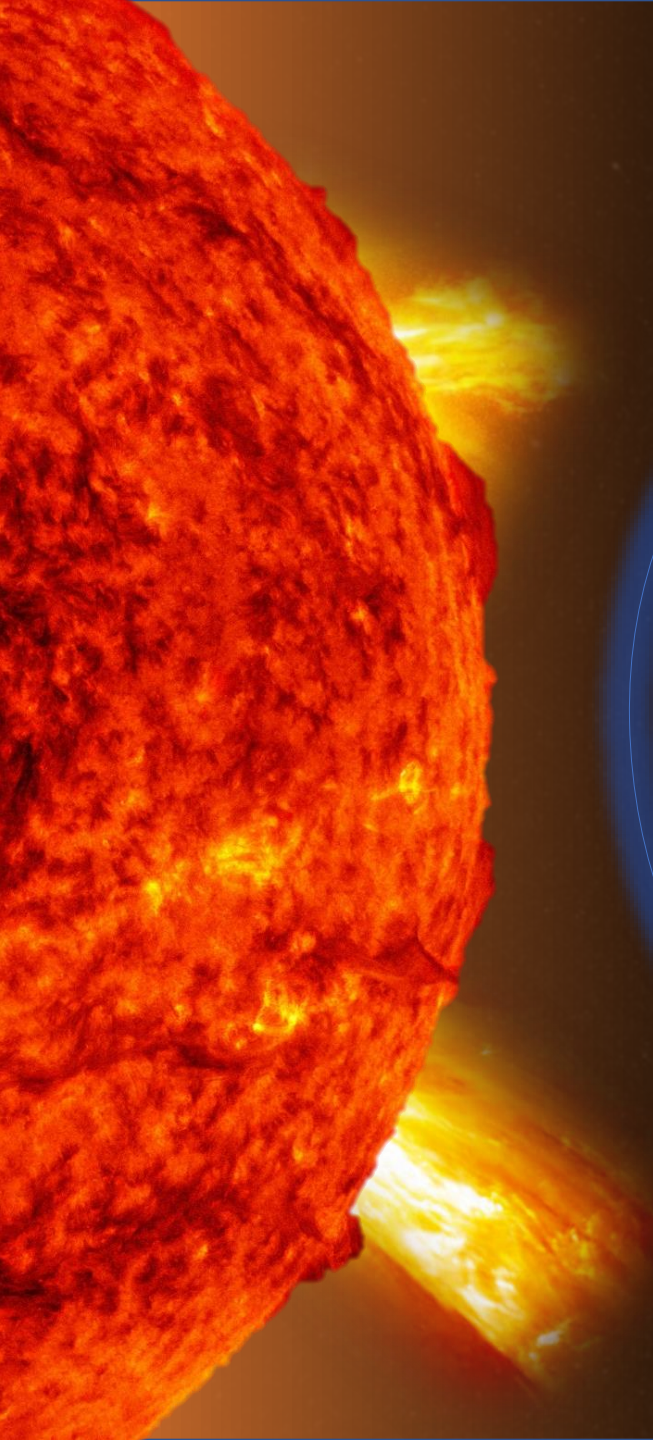

Thank you for your time!

Questions?

Comments?

Philosophies?

Single Event Effects (SEE) Symposium
Military & Aerospace Programmable Logic Devices (MAPLD)
Combined Workshop

SEE MAPLD

NASA Goddard Space Flight Center
Radiation Effects & Analysis Group (REAG)

To be presented by Seth Roffe at SEE/MAPLD in San Diego, CA on May 16, 2024