Karlsruhe Institute of Technology

coreboot and Dasharo in Systems Research at KIT

Yussuf Khalil, Fabian Meyer

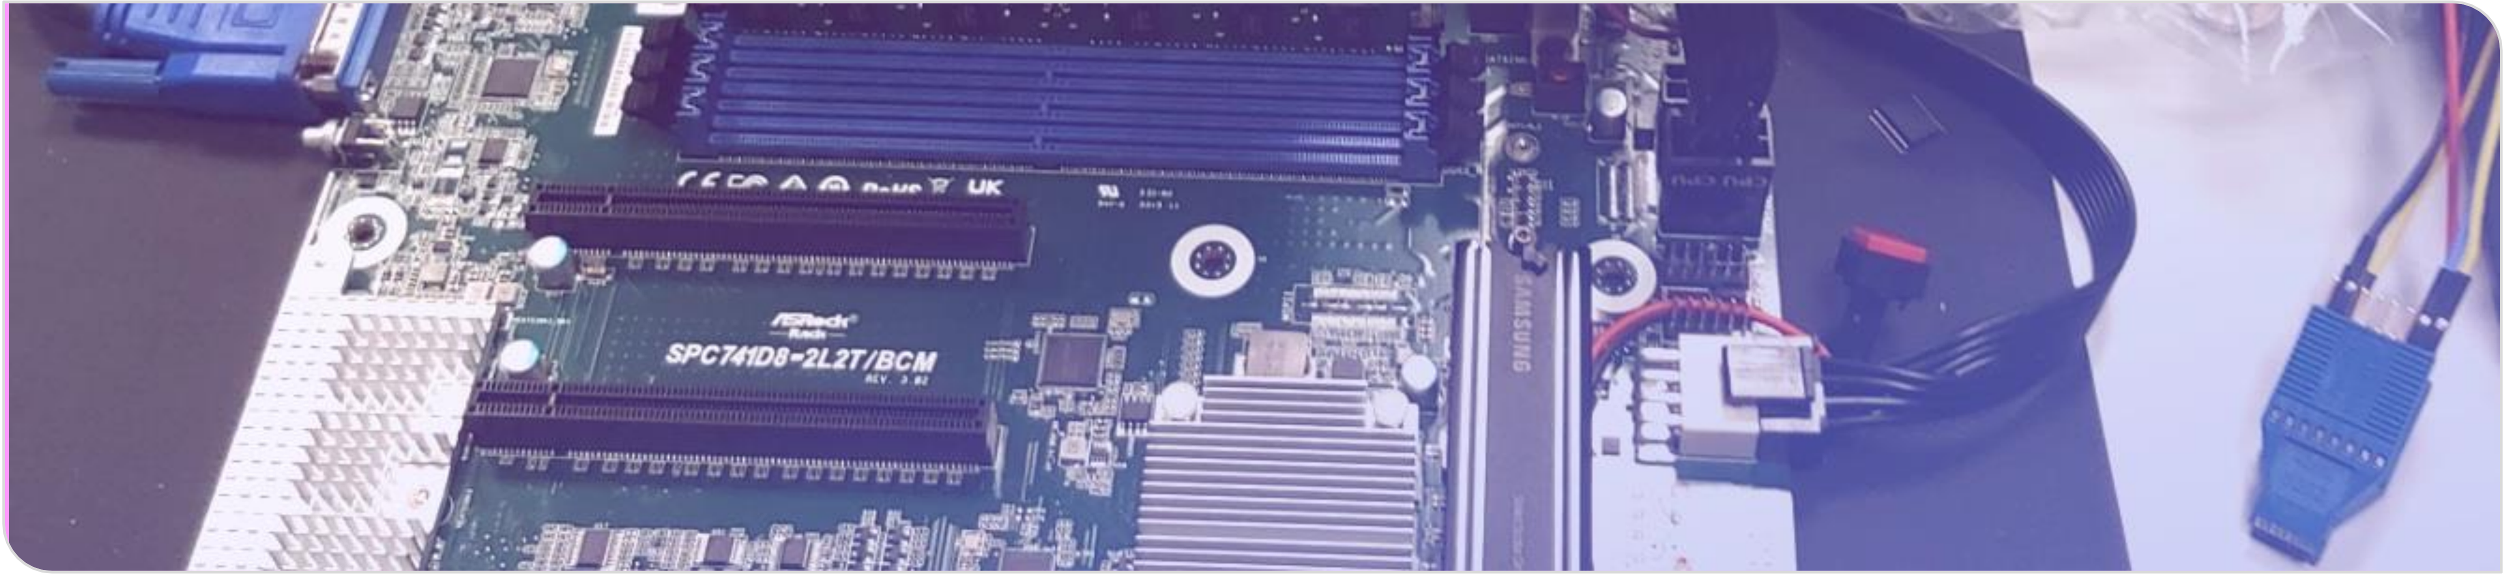

Who We Are

Yussuf Khalil

- 2nd year PhD student
- Working on power management and storage topics

Fabian Meyer

- Bachelor's thesis on firmware-based system suspend using non-volatile memory
- Porting coreboot to a Xeon SP mainboard
- Currently on the job market

Firmware in Research

- Firmware-related research typically focuses on security
  - Supply-chain attacks pose a huge risk for businesses and users
  - Firmware rootkits make all other security measures useless
  - LogoFAIL [1]
  - TPM GPIO fail [2]

[1] https://www.binarly.io/blog/the-far-reaching-consequences-of-logofail

[2] https://mkukri.xyz/2024/06/01/tpm-gpio-fail.html

Firmware in Research

- Firmware-related research typically focuses on security
  - Supply-chain attacks pose a huge risk for businesses and users
  - Firmware rootkits make all other security measures useless
  - LogoFAIL [1]
  - TPM GPIO fail [2]

But, what else can we do with firmware?

[1] https://www.binarly.io/blog/the-far-reaching-consequences-of-logofail

[2] https://mkukri.xyz/2024/06/01/tpm-gpio-fail.html

Firmware in Research

- Firmware also plays a vital role in power management
  - Important energy management information shared via ACPI
  - System suspend modes partially implemented in firmware
- Persistent caches require firmware support
  - eADR with Intel Optane
  - CXL Global Persistent Flush (GPF)

Firmware in Research

- Firmware also plays a vital role in power management
  - Important energy management information shared via ACPI
  - System suspend modes partially implemented in firmware
- Persistent caches require firmware support
  - eADR with Intel Optane
  - CXL Global Persistent Flush (GPF)

Let’s get creative!

Thesis: Suspending to Persistent Memory

- Aim to reduce:
  - Power consumption of a suspended device
  - Wake latency
- Implementation in coreboot (Dasharo) on real hardware

Thesis: Suspending to Persistent Memory

- Aim to reduce:
  - Power consumption of a suspended device
  - Wake latency
- Implementation in coreboot (Dasharo) on real hardware

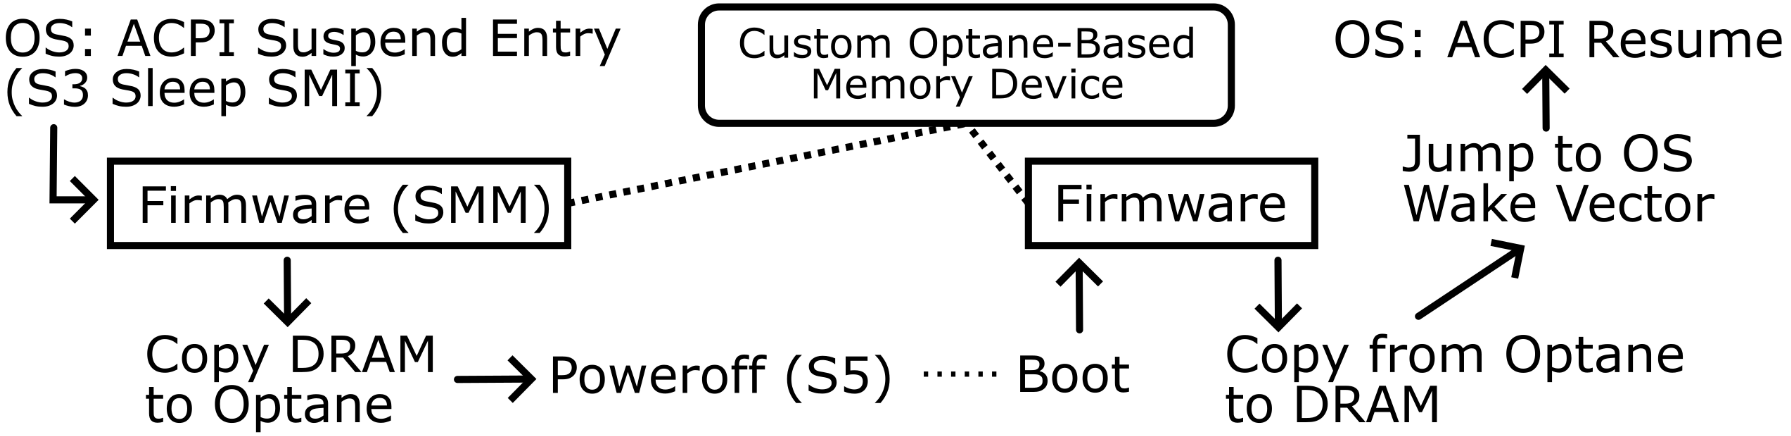

Thesis: Suspending to Persistent Memory

- Evaluation results:
  - Better energy efficiency than S3 after 1h of suspend time
  - Wake latency < 20s
  - Performance bottleneck: memory training

- Benefits when implemented in firmware vs. OS kernel:
  - Transparency to the OS, which performs a regular S3 entry/resume
  - Less overhead on resume: skip parts of firmware and bootloader
  - Otherwise, not competitive with S3

Thesis: https://os.itec.kit.edu/downloads/2023_BA_Meyer_Efficient_PMem_Suspend.pdf

Porting coreboot

- ASRock Rack SPC-741D8
  - Not restricted by Intel BootGuard, Platform Firmware Resiliency
  - Sapphire Rapids CPU
  - CXL support
  - Socketed flash IC

- First port of coreboot to off-the-shelf Eagle Stream board
  - Upstream patch: https://review.coreboot.org/c/coreboot/+/82203
  - 64-bit for all stages and SMM

Porting coreboot

- Most datasheets unavailable/confidential

- Debug methods severely limited
  - POST card: reliable, but not much bandwidth
  - USB/network: require arcane devices
  - cbmem: requires successful boot
  - Serial port/speaker modem: requires an NDA...

Future Research Plans

- So far: feasibility of custom suspend implementations demonstrated in Fabian’s thesis
  - But very similar to existing suspend-to-disk (S4) approaches
- Next up: implement our research ideas for a *new* suspend mode
  - Exploit hybrid CXL storage devices for fast wake-up with minimal energy consumption
    - Combine advantages of suspend-to-memory (S3) and suspend-to-disk (S4)
  - Can’t share too many details yet ☺

Future Research Plans

- CXL GPF flushes CPU caches to persistent memory on power loss
  - Caches effectively made persistent
- Flushing handled in firmware via System Management Mode

Future Research Plans

- CXL GPF flushes CPU caches to persistent memory on power loss
  - Caches effectively made persistent
- Flushing handled in firmware via System Management Mode

Can we integrate that with the OS in a smart manner?

- Don't want to put all process pages onto persistent memory
  - Partial persistence at runtime
  - Some more SMM logic to flush parts of DRAM on power loss?

Conclusion

- Previous firmware-related research focuses on security
- Firmware is interesting for other systems research areas as well
- First coreboot port to off-the-shelf Sapphire Rapids board
- Currently working on suspend and persistent memory topics