

Decelerating Suspend and Resume in Operating Systems

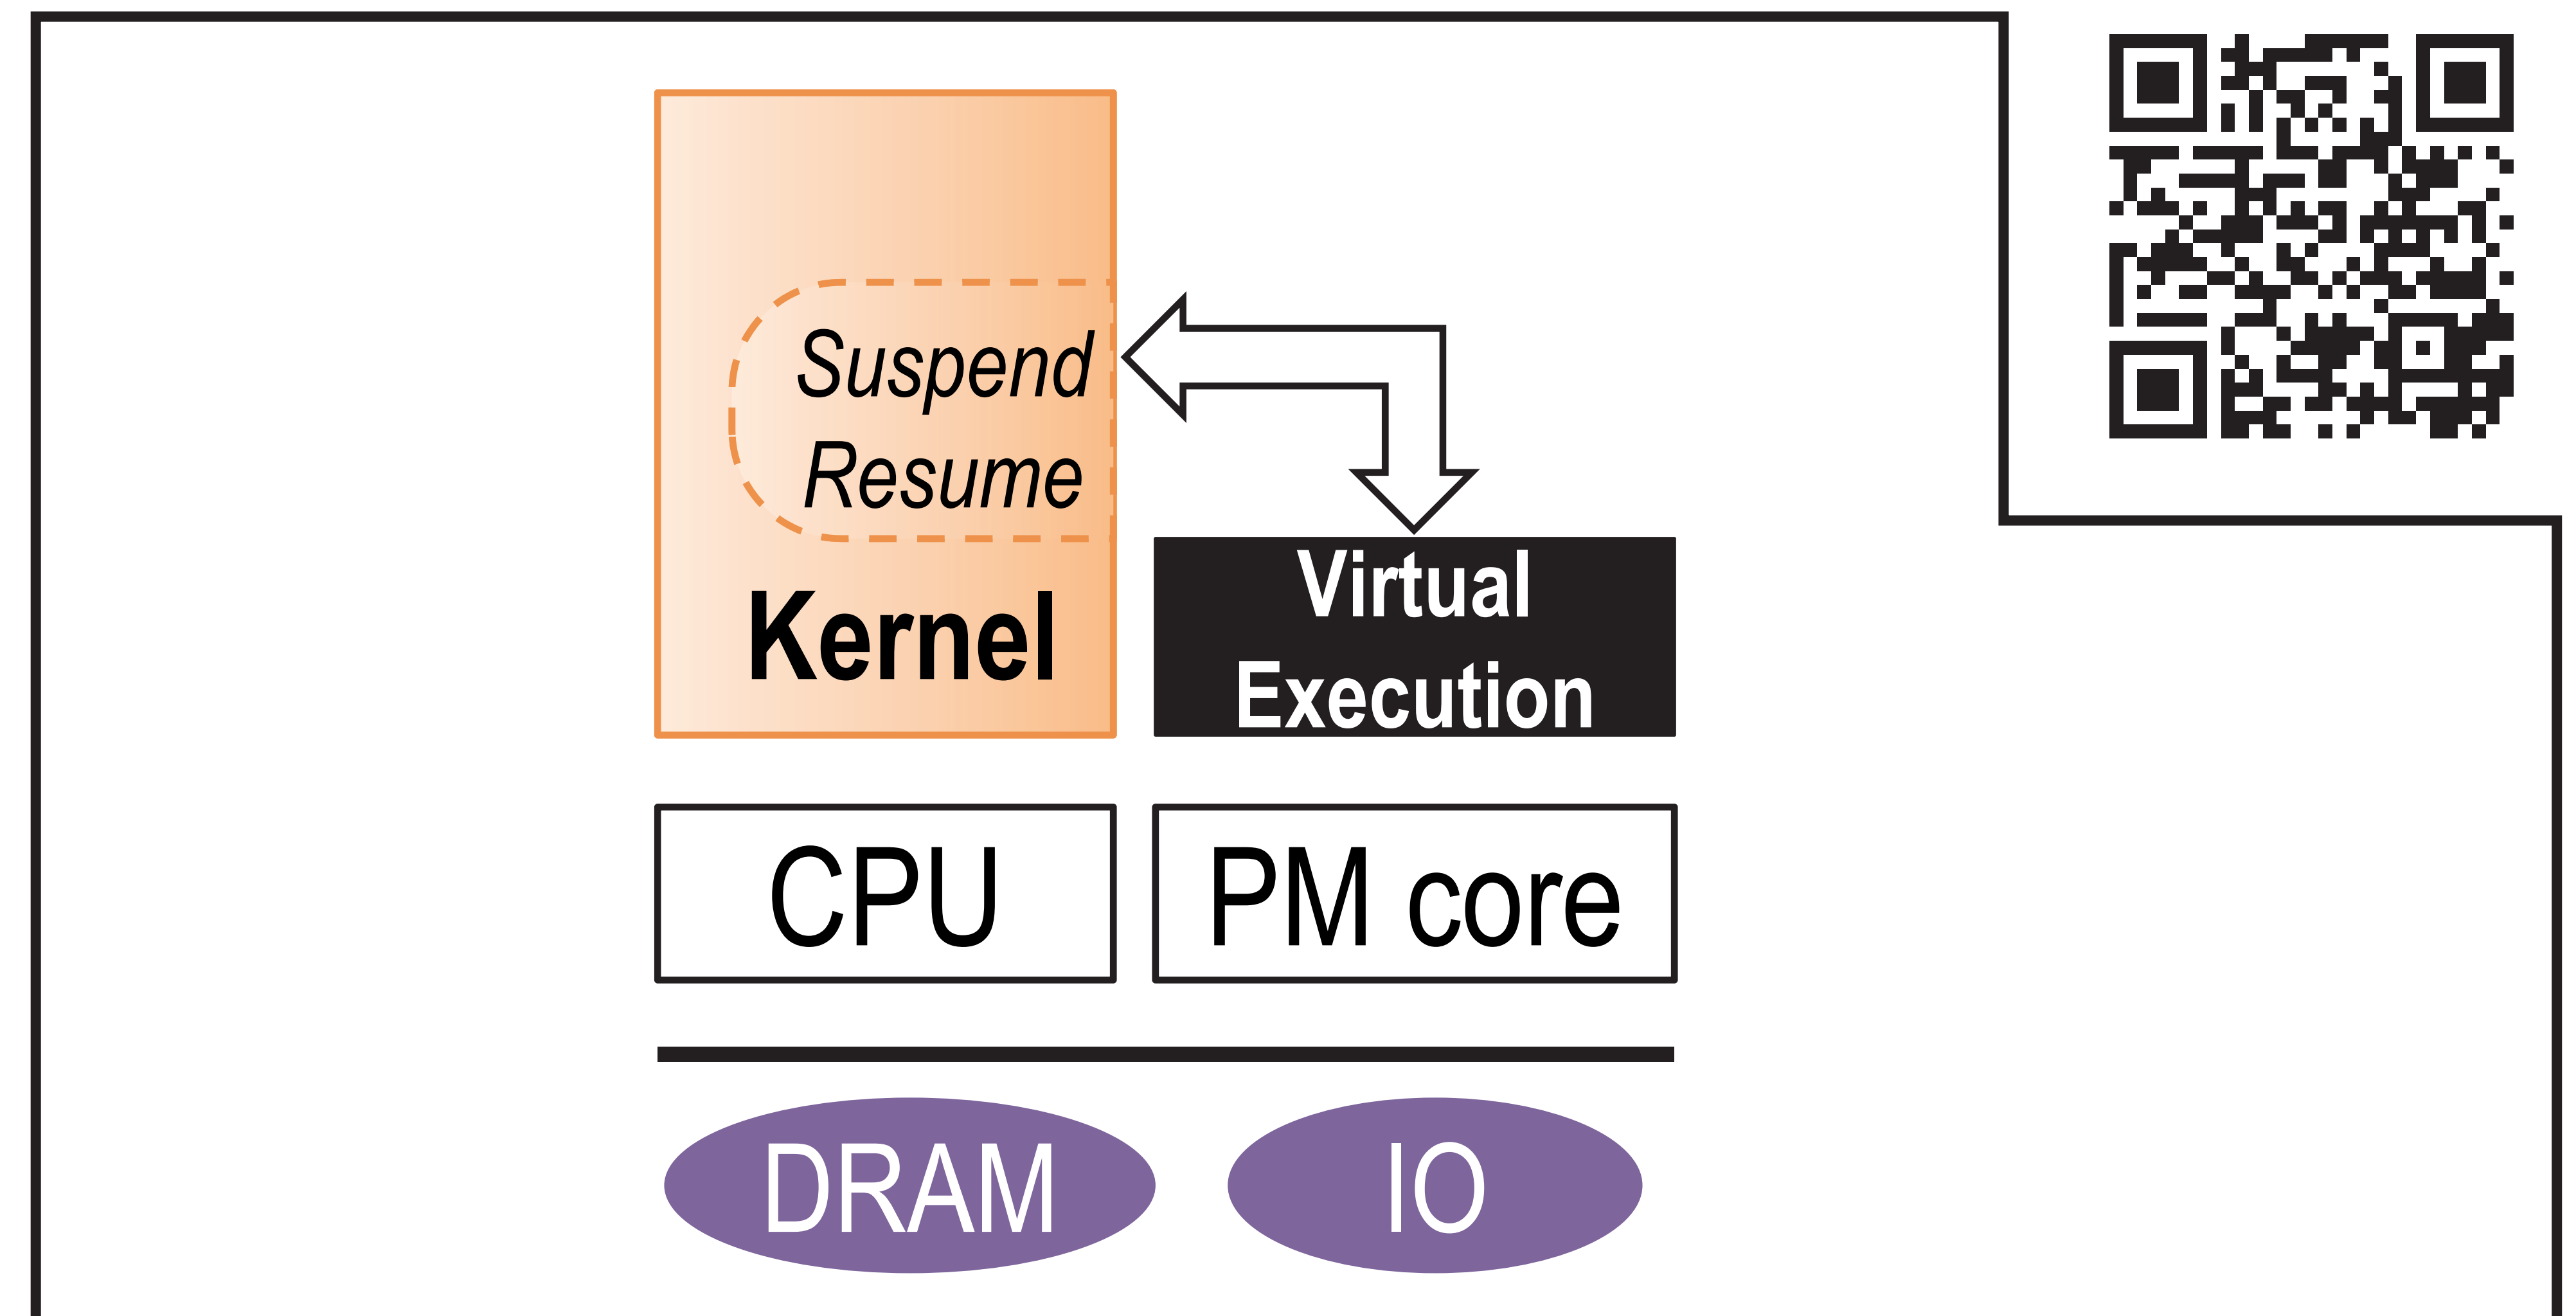
Shuang Zhai, Liwei Guo, Xiangyu Li, and Felix Xiaozhu Lin

{zhais, guo405, li1675, xzl}@purdue.edu

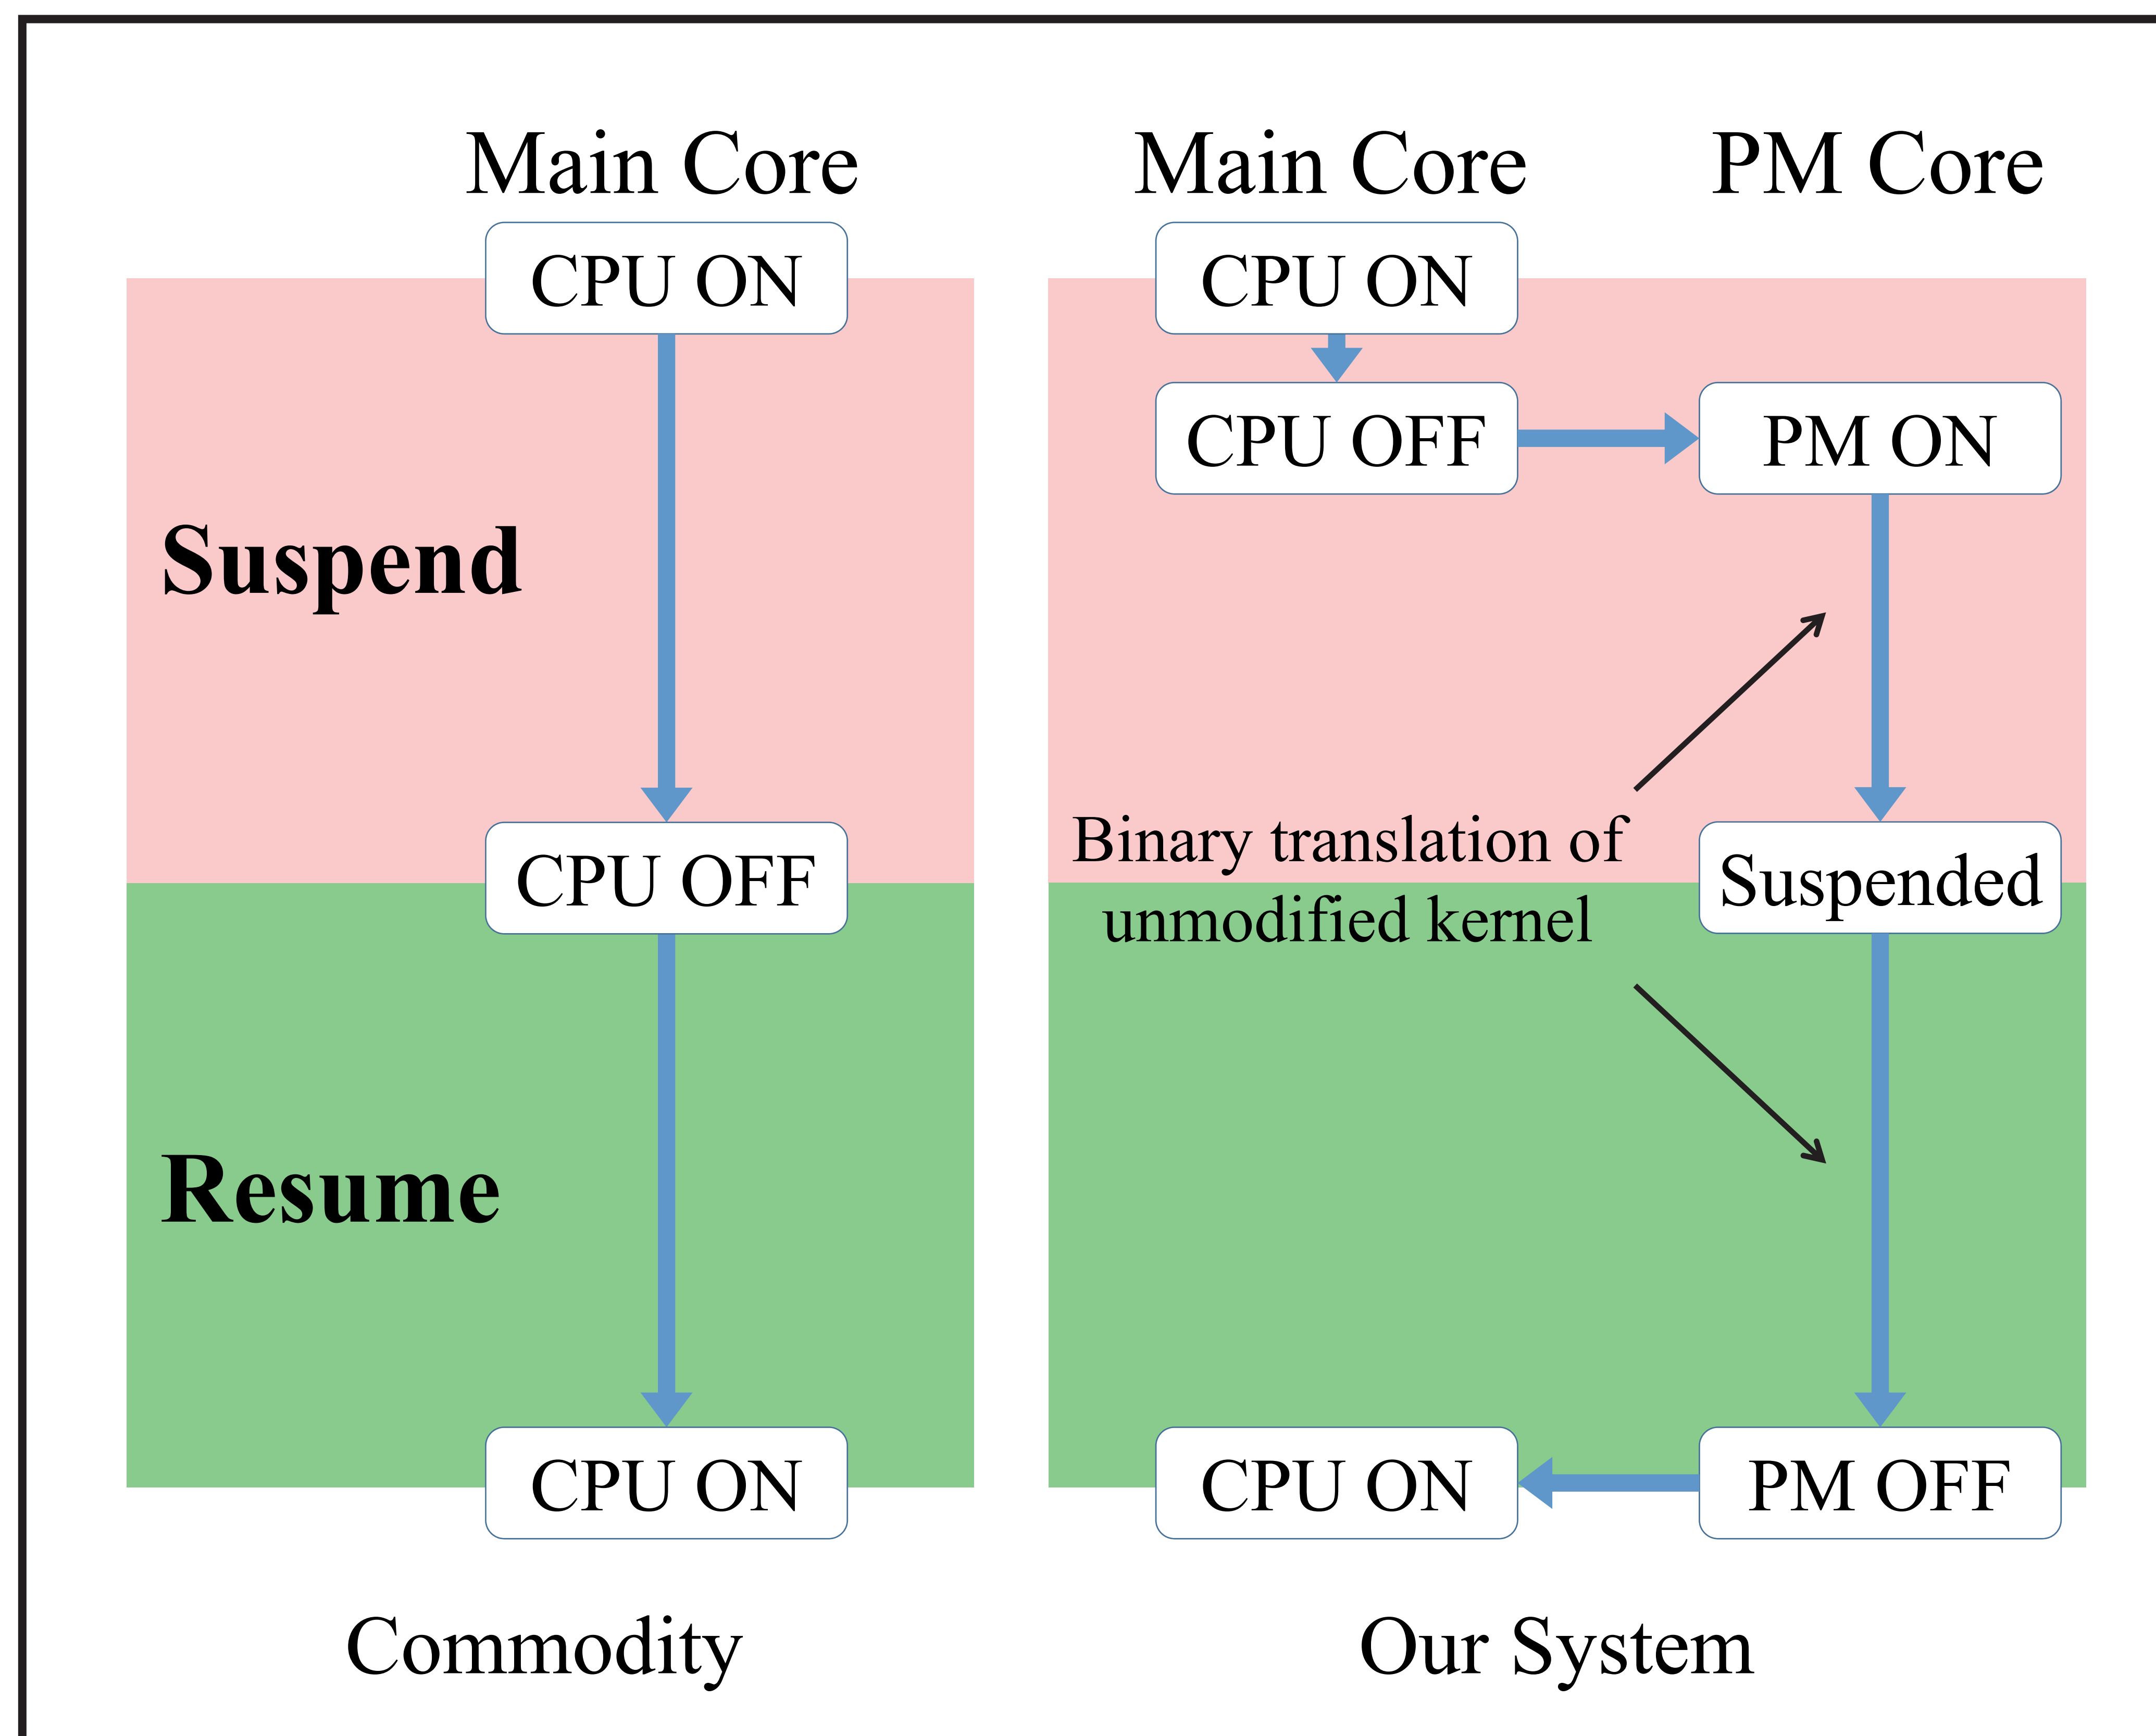
1. Suspend/Resume is inefficient

- Busy/idle wait is common in OS suspend/resume path.
- Time wasted on power state transitions of IO devices.
- Slow IO devices are diverse across different platforms.
- Proposal: offloading to a miniature core, dubbed PM core.

2. Solution



3. Workflow



4. Contributions

- Aggressively optimizing binary translation for ISA similarity: same number of registers, similar instruction format, etc.
- Key Optimizations
 - Direct register mapping
 - Baremetal stacks
 - Relaxed handling of interrupts and exceptions
- 4.5K SLoC new code atop 50.5K SLoC QEMU dynamic binary translator.
- Results: translation overhead reduced by 5x; energy of suspend/resume reduced by 70%.

A full implementation is in progress!

01/17/2017