SPP2377 Annual Meeting 2025

DRAMaOS: **DRAM-a**ware **OS**


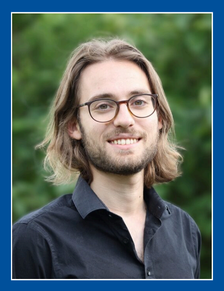

Illia Ostapyshyn Nils Wilbert Matthias Jung Daniel Lohmann

Breaking the Barrier Between OS and HW Memory Management

Problem: Modern DRAM is increasingly complex and gains new features, but the existing hardware and software lack support to unleash its full potential.

Solution:

- OS becomes aware of the DRAM's internal structure to manage it effectively.
- Smart memory controller exposes control and performs simple tasks autonomously.

Example: Power Management

- Allocated memory is *scattered over time*
- Existing features apply to *coarse granularities*

→ **Suitable compaction must be *cost-aware***

Available power management features

■ (◪) supported (optional) ■ with retention ■ accessible

Power mode	Granularity	DDR	LPDDR	HBM
		3 4 5 6	3 4 5 6	2 3 4
Full power off	Channel	■■■■	■■■■	□□□
Self refresh	Channel	■■■ ?	■■■■	■■ ?
MPSM	Channel	□■■ ?	□□□□	□□□
MPSM (DDR5)	Rank	□□■ ?	□□□□	□□□
Power-Down	Rank	■■■ ?	■■■■	■■ ?
PASR	1/8 Rank	◪□◪ ?	■■■■	□□□
PARC	1/8 Rank	□□□ ?	□□■■	□□□
Efficiency mode	Subchannel	□□□ ?	□□□■	□□□

Effect of defragmentation effort by granularity

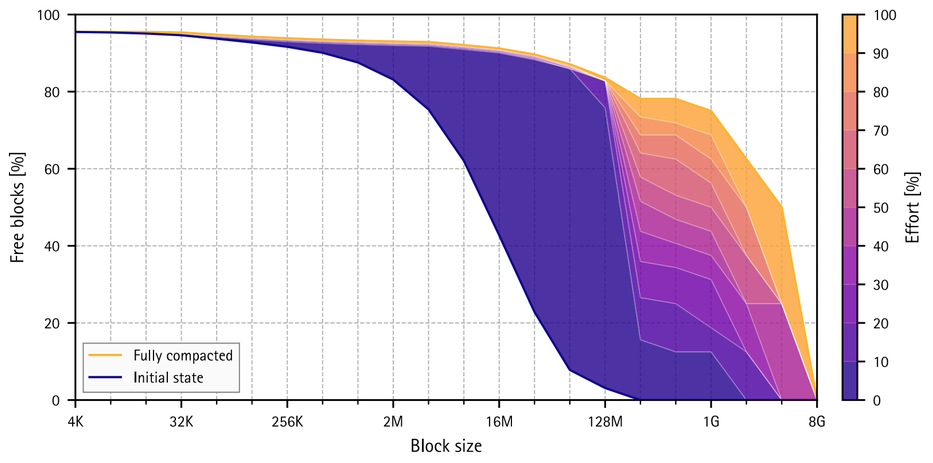

Improved OS Memory Management

Problem: Current memory management is based on outdated beliefs

1. Only used memory is good memory (cache 'em all)
2. All frames are equal (no external fragmentation)

Our Belief: OS must be aware and in control of DRAM's capabilities.

- Page cache actually costs energy and therefore money
- Targeted allocation can increase performance and security/safety

→ Reflect DRAM capabilities and *memory slices* they affect in the memory management structures and policies of the OS.

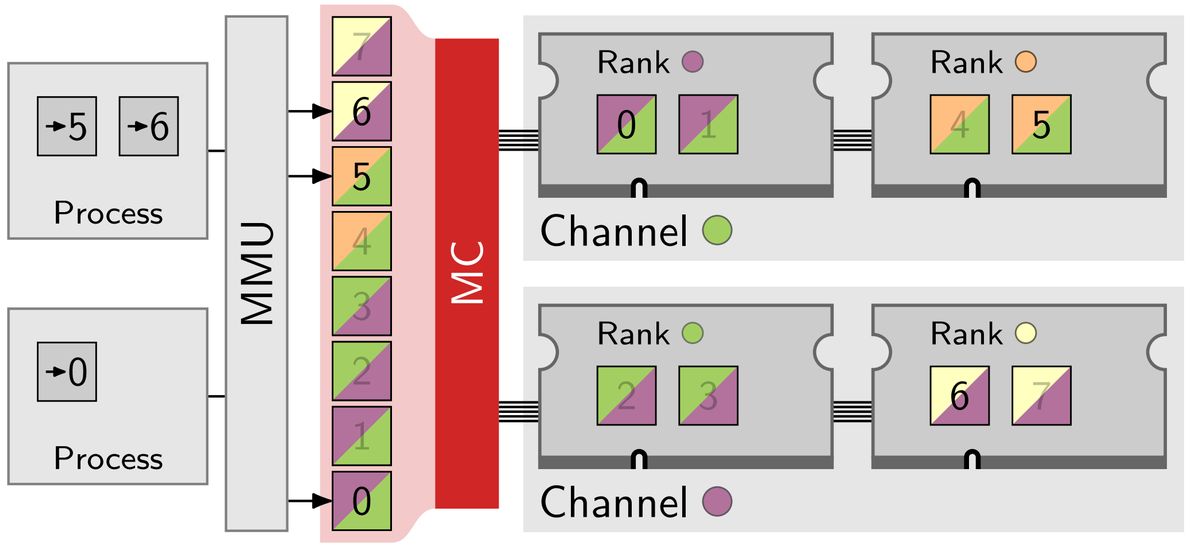

Smart Memory Controller

Problem: Traditional DRAM Memory Controllers (MC)

1. Are non-programmable
2. Can only use locally buffered requests to perform optimizations
3. Do not expose DRAM internals to OS/software stack

Our Proposal: Add programmable OS machines to the MC.

- OS can offload tasks to the MC
- Explore different levels of OS-machine complexity

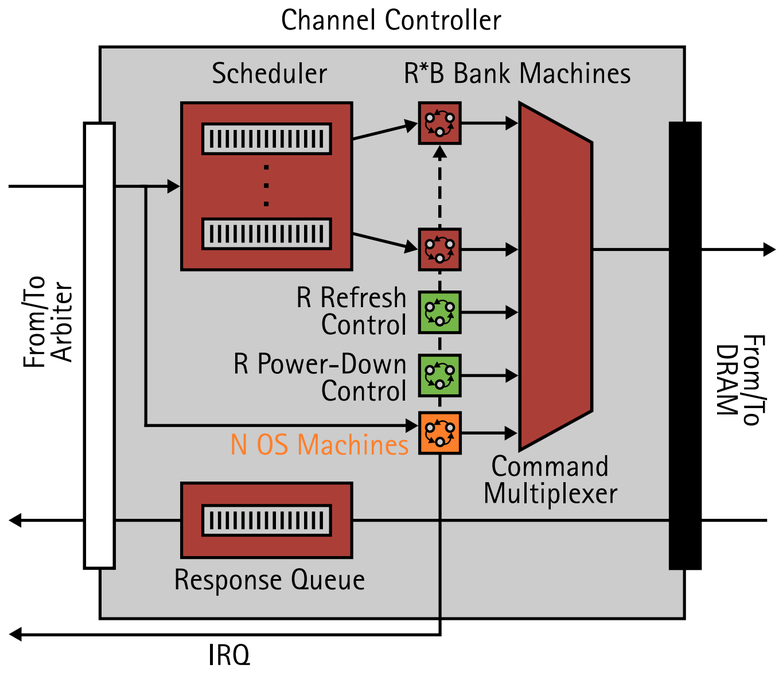

Virtual Platform

DRAMSys 5.0

- DRAMSys
  - Cycle-accurate DRAM simulator for all JEDEC standards
  - Fast TLM-level simulation
  - Includes power model and trace analyzer
  - Available on https://github.com/tukl-msd/DRAMSys
- gem5
  - Full-system simulator capable of booting Linux

Our Framework: Couple DRAMSys with gem5 to create a fast and flexible simulation platform supporting our proposed OS-MC codesign.

Working Plan

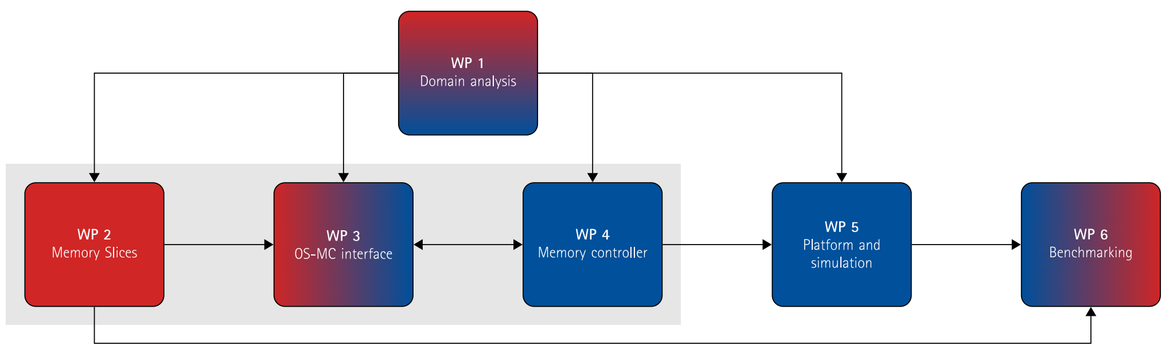

Role within the SPP: Cross-domain project combining HW simulation and OS areas, allowing for great cooperation potential.

Systems Research and Architecture
https://sra.uni-hannover.de

Leibniz Universität Hannover

Computer Engineering
https://www.informatik.uni-wuerzburg.de/ce

Julius-Maximilians-Universität Würzburg

567410490
JU 3158/1-1, LO 1719/10-1

DFG