FUJITSU

shaping tomorrow with you

# The Analysis of Inter-Process Interference on a Hybrid Memory System

**Satoshi Imamura** and Eiji Yoshida
Fujitsu Laboratories Ltd., Japan

IXPUG HPC Asia 2020, January 17, 2020

# Intel® Optane™ DC Persistent Memory (DCPM) FUJITSU

- First ***persistent memory*** product
  - Connected to DIMM slots
  - Byte-addressable
  - Non-volatile

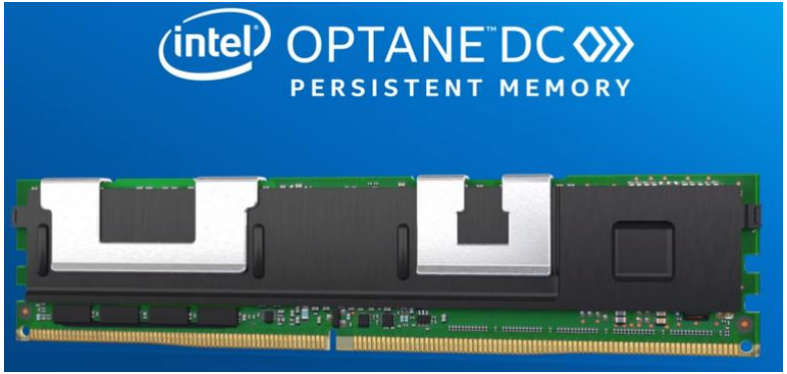


https://www.intel.co.jp/content/www/jp/ja/architecture-and-technology/optane-dc-persistent-memory.html

- Larger capacity than DRAM
  - 128, 256, 512 GB per DIMM

- Lower cost per GB than DRAM (**about half** [1, 2])

[1] https://www.tomshardware.com/news/intel-optane-dimm-pricing-performance,39007.html, [2] https://memory.net/memory-prices/

# Hybrid Memory Systems (HMS)

FUJITSU

- Combine DRAM and DCPM to take their advantages
  - DRAM is **faster**, but smaller and more expensive
  - DCPM is **larger and cheaper**, but slower [Izraelevitz+, arXiv 2019]

**Configuration of our experimental server**

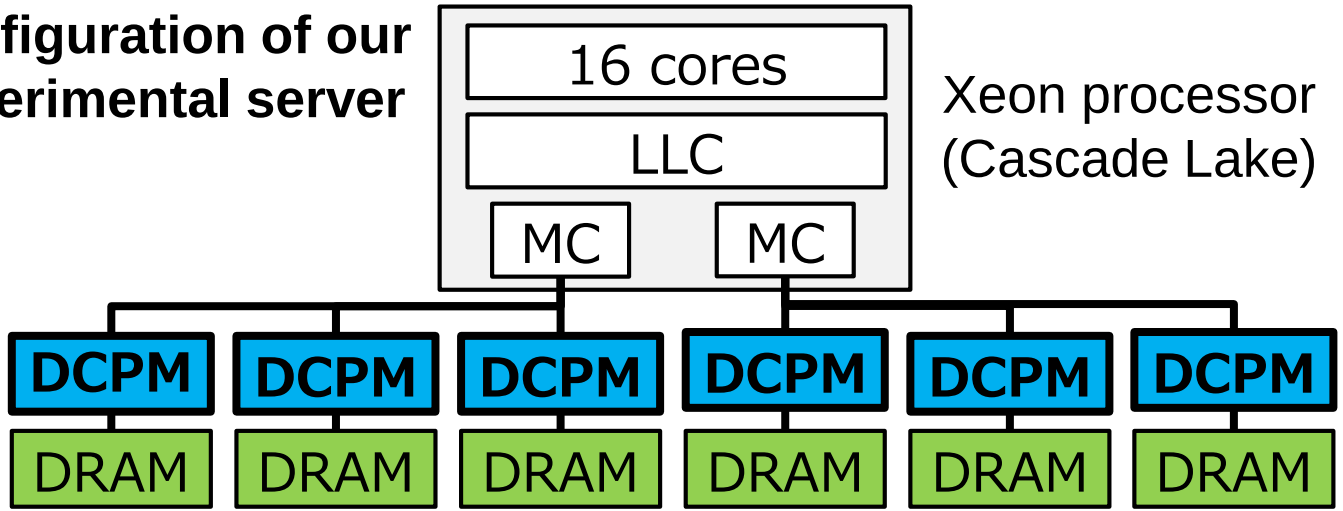

# A Use Case of HMS

FUJITSU

- **Virtualization** for cloud services such as IaaS
  - HMS will bring benefits to both service providers and customers

**For service providers**

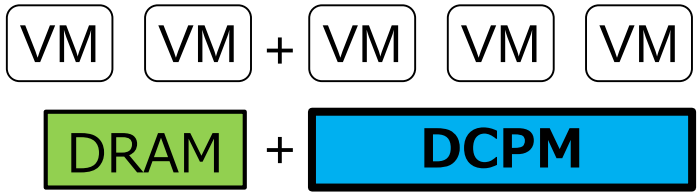


**Larger memory for more VMs**

**For customers**

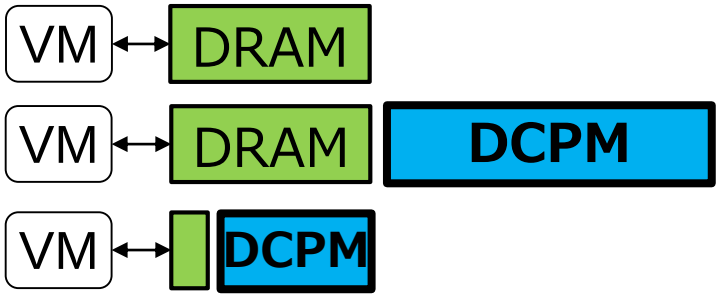


**More various VM instances**

**Multiple VMs are consolidated onto a single CPU**

# Executive Summary

FUJITSU

- Objective of this work
  - To analyze performance interference between two processes on a HMS (in a non-virtualized environment)
- Experimental methodology
  - We co-execute ***target*** and ***competing*** processes on a single CPU
  - We measure the slowdown (normalized throughput) of a target process compared to its solo execution
- Our finding
  - **Interference on a HMS is much different from that on a conventional DRAM-only memory system**

# Experimental Setup

FUJITSU

**8 types** of synthetic processes (Intel® Memory Latency Checker)

| | |
|---|---|
| **Access pattern** | Sequential (Sq) or Random (Rn) |
| **Read-write pattern** | Read-only (R) or Read-write mixed (W2) |
| **Access target** | DRAM (D) or DCPM (P) |

Delay = 0 cycle

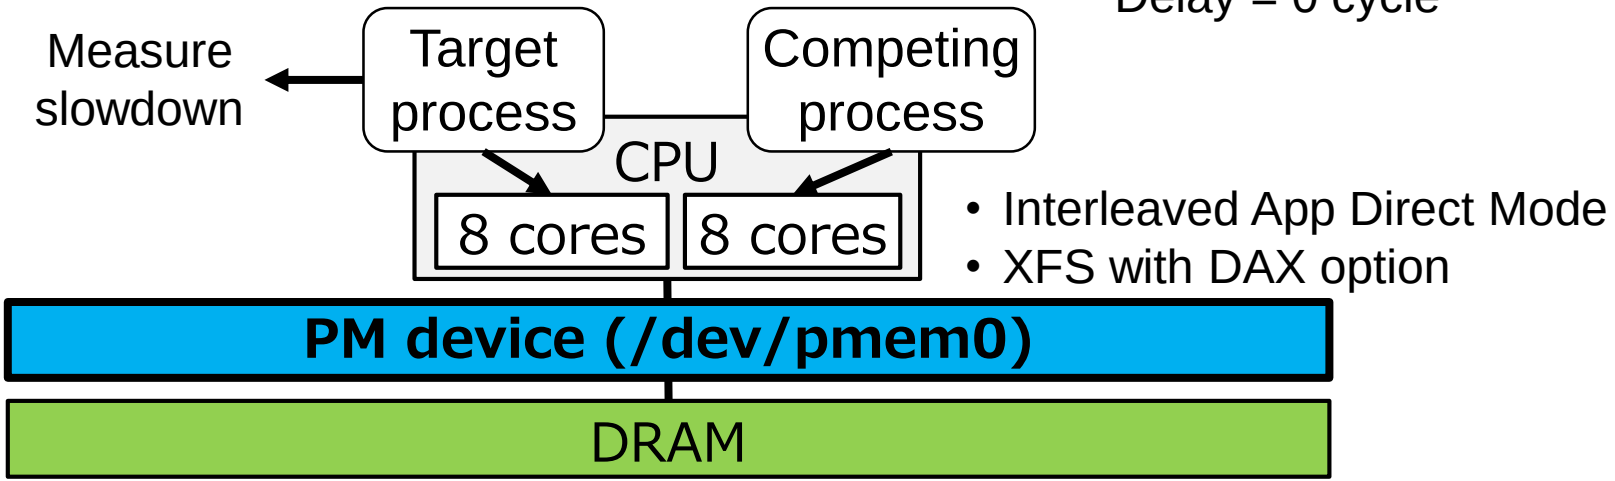


- Interleaved App Direct Mode
- XFS with DAX option

# Slowdown of Target Process

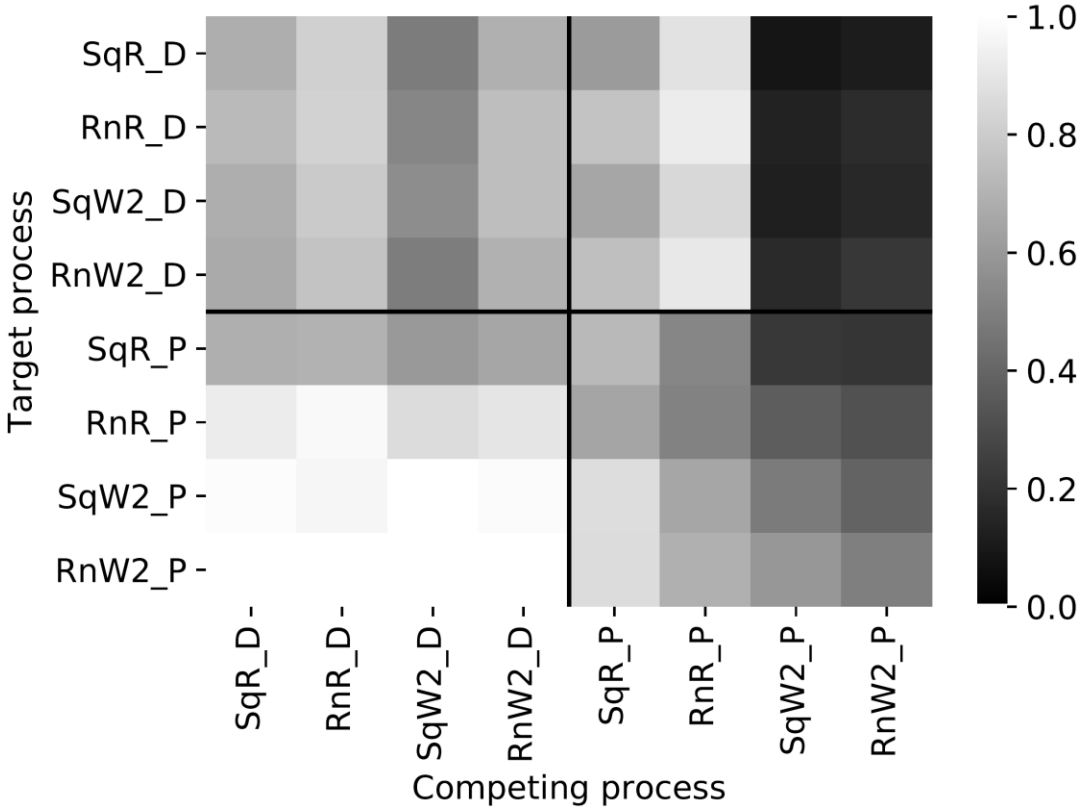


- 64 (= 8x8) combinations
- Darker color ⇒ larger slowdown compared to solo execution

# Interference b/w Processes Accessing DRAM

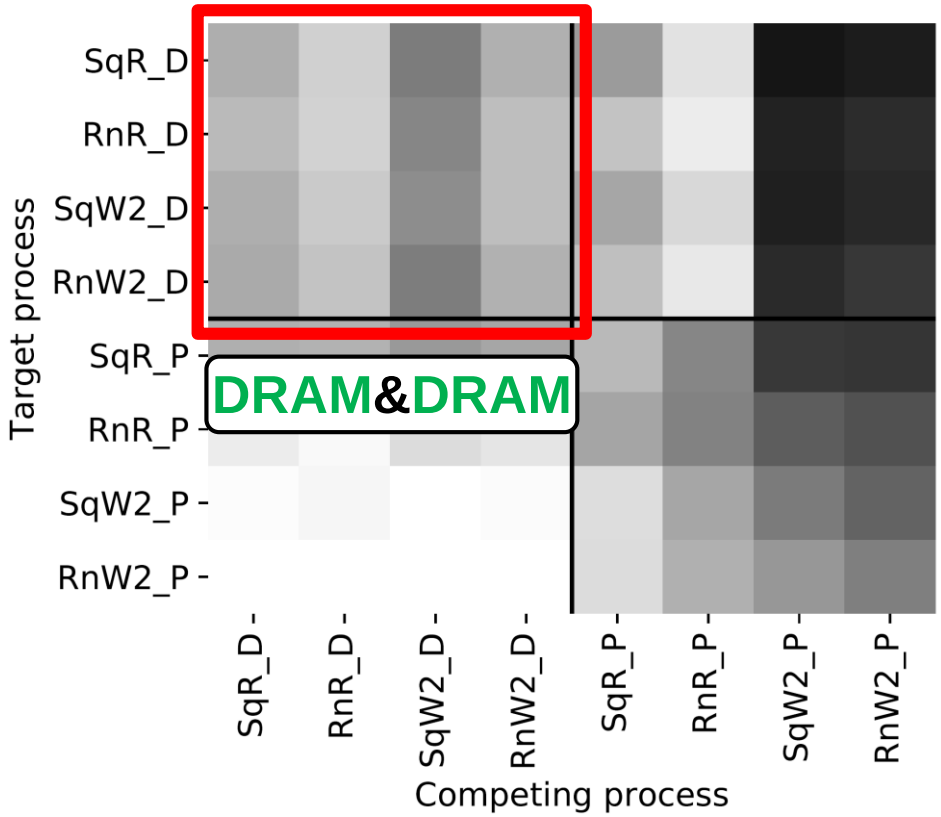


- Interference on a **DRAM-only system**
- **Small differences between different combinations**
  - Entirely grayed

# Interference b/w Processes Accessing DCPM

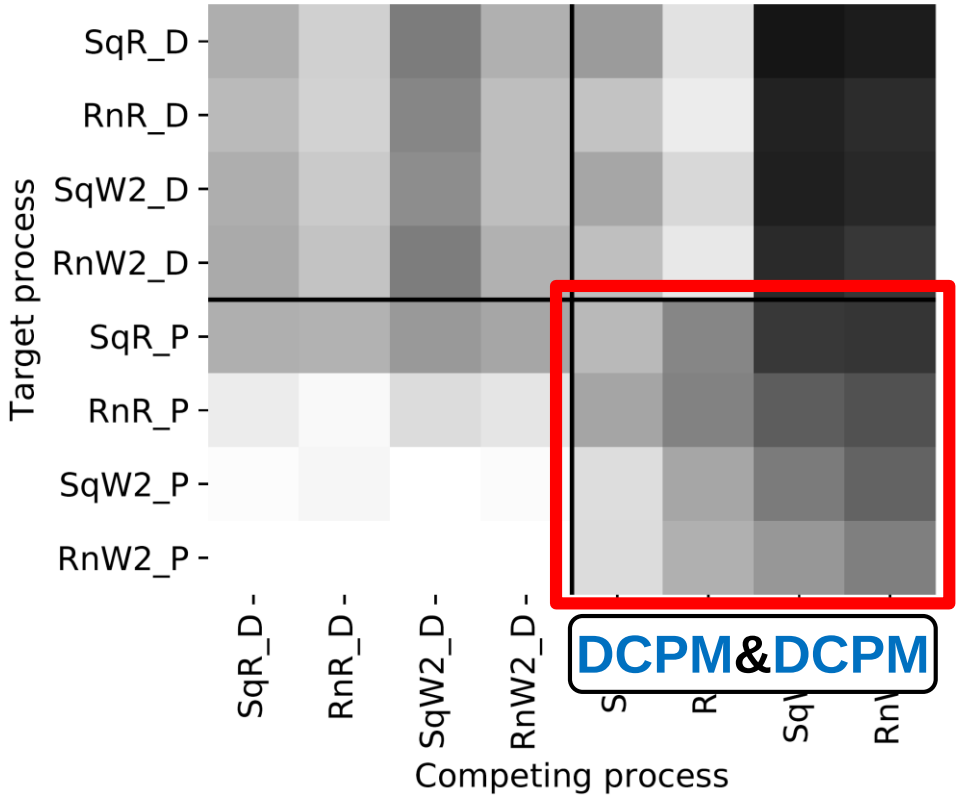


- Similar results to "DRAM&DRAM" case
  - Entirely grayed
- Some exceptions
  - Please refer to our paper for the details

# Interference between DRAM/DCPM Accesses

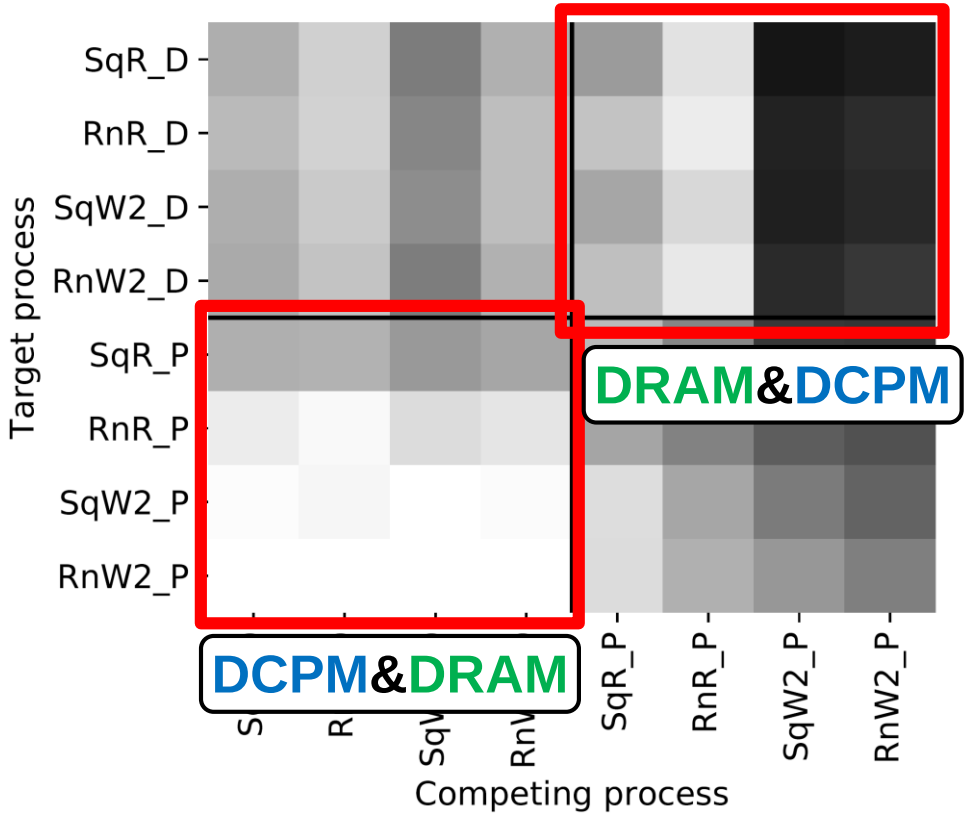


- Much different results from "DRAM&DRAM" and "DCPM&DCPM" cases
- **Significant differences between different combinations**
  - White ⇒ small slowdown
  - Black ⇒ large slowdown

# Observation 1: Small Interference

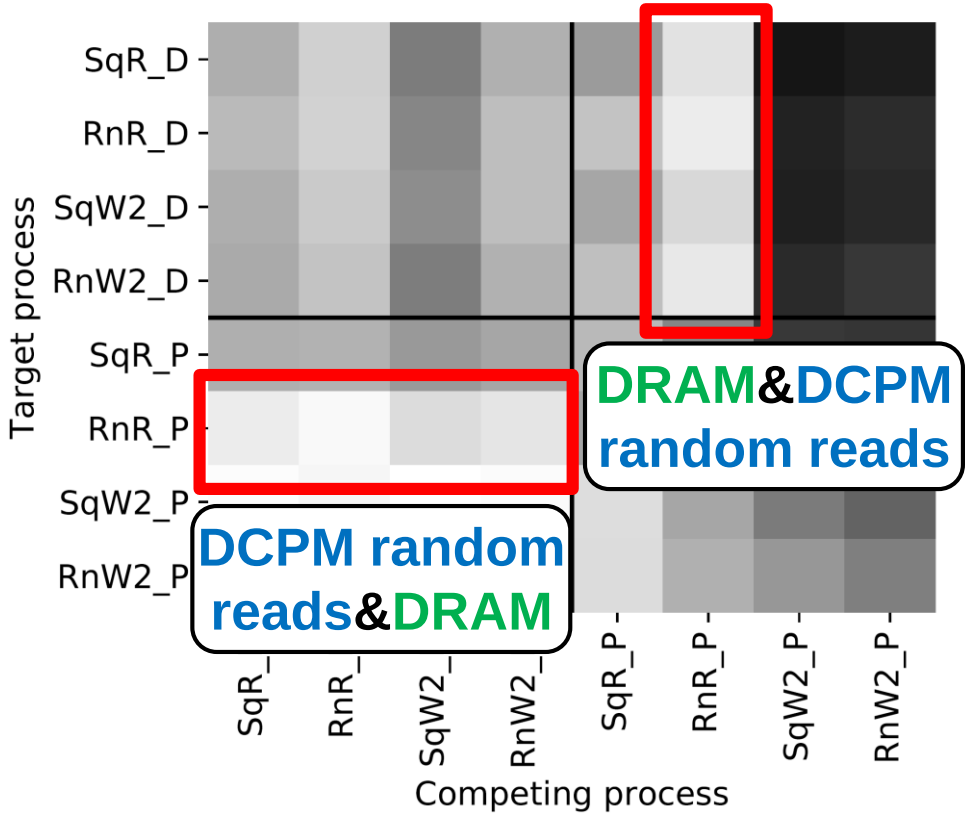


- **Small interference between DRAM accesses and DCPM random reads**
- Our analysis shows
  - Degree of interference: sequential > random
  - Low contention of row buffers with separate accesses to DRAM and DCPM

# Observation 2: Significant Interference

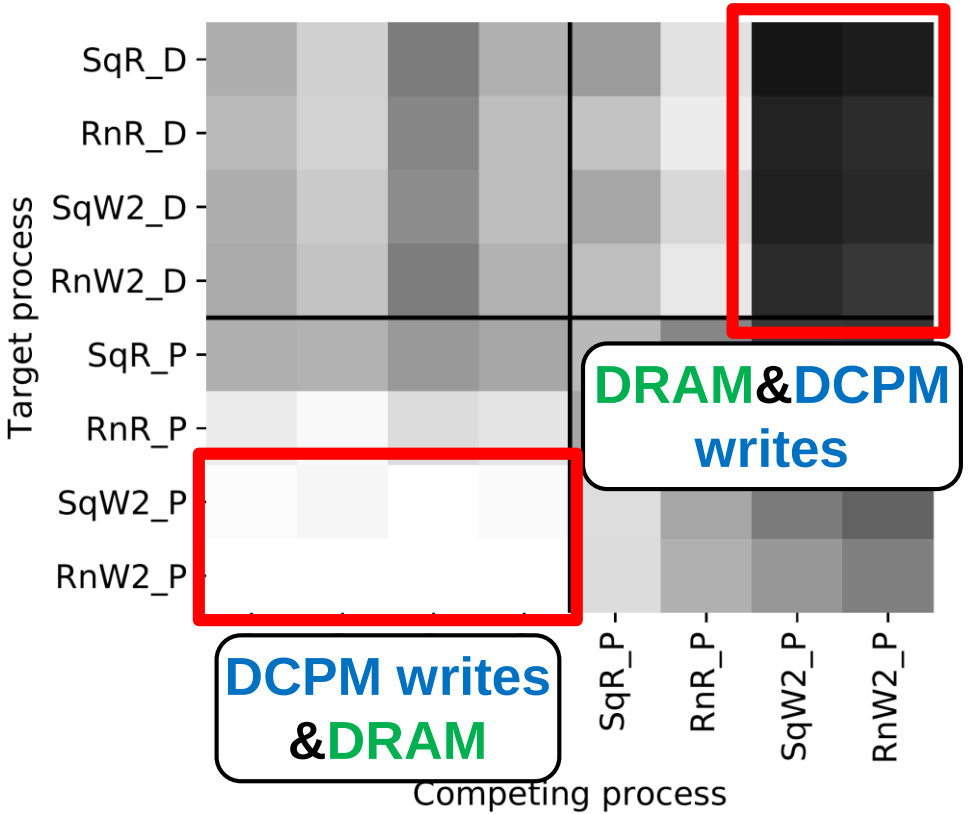


- **DRAM accesses are heavily disturbed by DCPM writes**
- Our analysis shows
  - Frequent DCPM writes fill up DCPM write queues
    ⇒ DRAM throughput drops significantly

# Conclusion

FUJITSU

- HMS will bring benefits to could services such as IaaS
  - Multiple VMs are consolidated onto a single CPU
- We evaluate and analyze the performance interference between two processes on a HMS
- Two observations:
  - Small interference between DRAM accesses and DCPM random reads
  - DRAM accesses are heavily disturbed by frequent DCPM writes

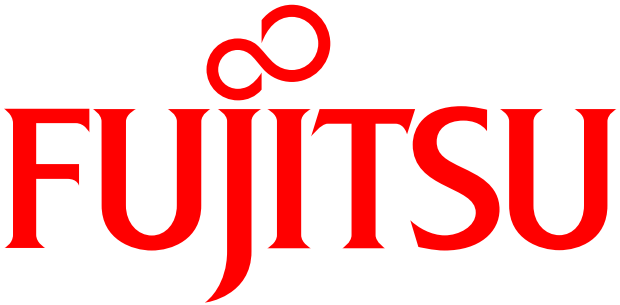


shaping tomorrow with you

# Analysis of Small Interference on HMS

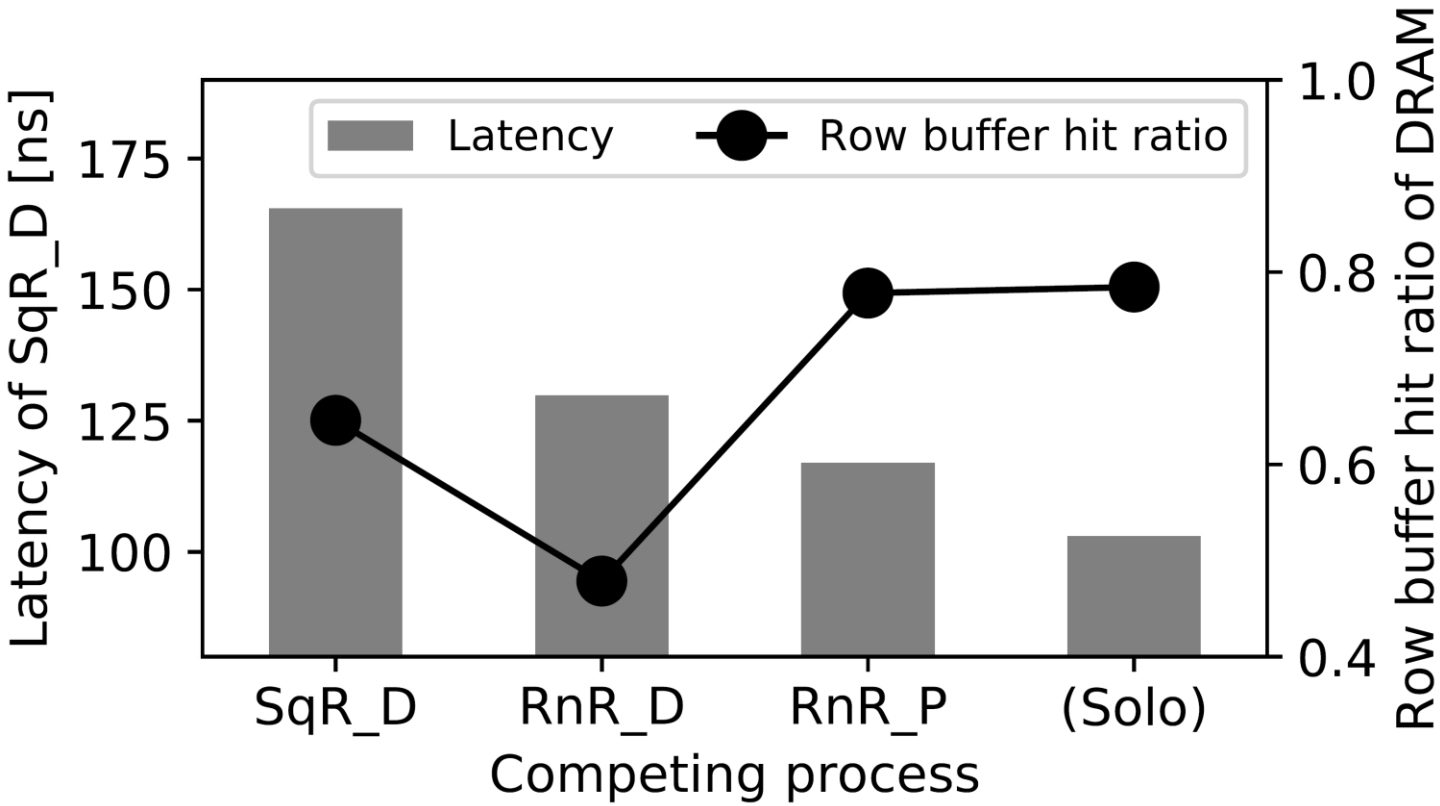

# Analysis of Significant Interference on HMS