2nd Workshop on Memory-Centric Computing: Processing-Near-Memory - Part II

Dr. Geraldo F. Oliveira

https://geraldofojunior.github.io

ICS 2025

08 June 2025

SAFARI

ETH zürich

Possible PNM Designs

- General-purpose programmable cores
  - Wimpy cores (possibility of running any workload)
  - E.g. from academia: Tesseract PIM for Graph Processing
  - E.g. from industry: UPMEM PIM

- **Fixed-function units**
  - Hardware/software co-designed PIM for efficiency
  - E.g. from academia: Mensa for NN Edge Inference
  - E.g. from industry: Samsung HBM-PIM, SK hynix AiM

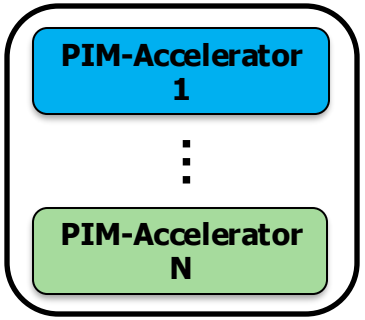

- Reconfigurable architectures
  - PNM cores coupled with FPGAs, CGRA
  - E.g. from academia: NERO for Weather Prediction
  - E.g. from industry: Samsung AxDIMM

Samsung Function-in-Memory DRAM (2021)

Samsung Newsroom

CORPORATE | PRODUCTS | PRESS RESOURCES | VIEWS | ABOUT US

Samsung Develops Industry's First High Bandwidth Memory with AI Processing Power

Korea on February 17, 2021 Audio Share

The new architecture will deliver over twice the system performance and reduce energy consumption by more than 70%

Samsung Electronics, the world leader in advanced memory technology, today announced that it has developed the industry's first High Bandwidth Memory (HBM) integrated with artificial intelligence (AI) processing power — the HBM-PIM. The new processing-in-memory (PIM) architecture brings powerful AI computing capabilities inside high-performance memory, to accelerate large-scale processing in data centers, high performance computing (HPC) systems and AI-enabled mobile applications.

Kwangil Park, senior vice president of Memory Product Planning at Samsung Electronics stated, "Our groundbreaking HBM-PIM is the industry's first programmable PIM solution tailored for diverse AI-driven workloads such as HPC, training and inference. We plan to build upon this breakthrough by further collaborating with AI solution providers for even more advanced PIM-powered applications."

https://news.samsung.com/global/samsung-develops-industrys-first-high-bandwidth-memory-with-ai-processing-power

Samsung Function-in-Memory DRAM (2021)

- FIMDRAM based on HBM2

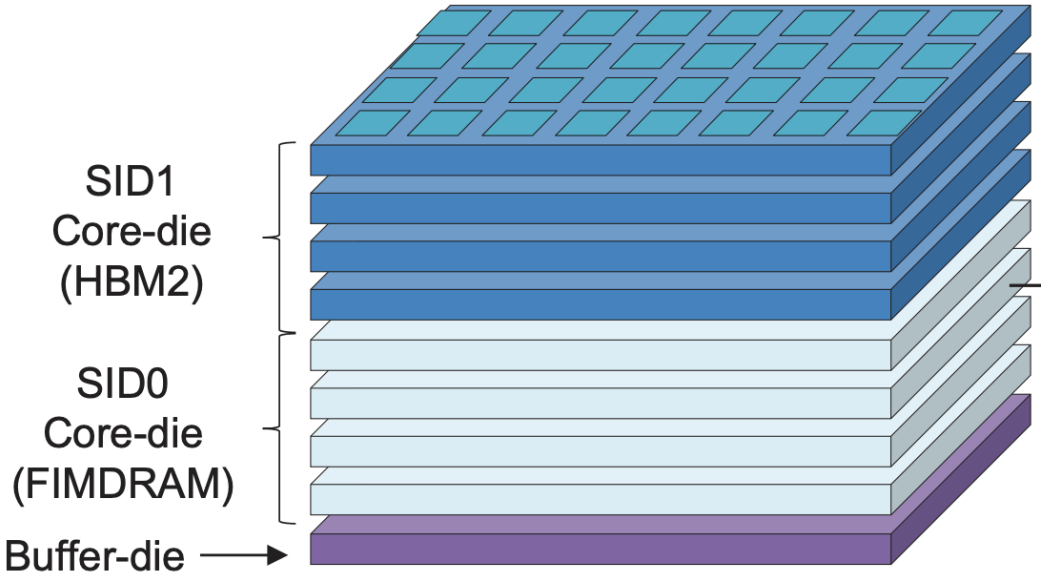

[3D Chip Structure of HBM with FIMDRAM]

Chip Specification
128DQ / 8CH / 16 banks / BL4
32 PCU blocks (1 FIM block/2 banks)
1.2 TFLOPS (4H)
FP16 ADD / Multiply (MUL) / Multiply-Accumulate (MAC) / Multiply-and- Add (MAD)

ISSCC 2021 / SESSION 25 / DRAM / 25.4

25.4 A 20nm 6GB Function-In-Memory DRAM, Based on HBM2 with a 1.2TFLOPS Programmable Computing Unit Using Bank-Level Parallelism, for Machine Learning Applications

Young-Cheon Kwon[1], Suk Han Lee[1], Jaehoon Lee[1], Sang-Hyuk Kwon[1], Je Min Ryu[1], Jong-Pil Son[1], Seongil O[1], Hak-Soo Yu[1], Haesuk Lee[1], Soo Young Kim[1], Youngmin Cho[1], Jin Guk Kim[1], Jongyoon Choi[1], Hyun-Sung Shin[1], Jin Kim[1], BengSeng Phuah[1], HyoungMin Kim[1], Myeong Jun Song[1], Ahn Choi[1], Daeho Kim[1], SooYoung Kim[1], Eun-Bong Kim[1], David Wang[2], Shinhaeng Kang[1], Yuhwan Ro[3], Seungwoo Seo[3], JoonHo Song[3], Jaeyoun Youn[1], Kyomin Sohn[1], Nam Sung Kim[1]

[1]Samsung Electronics, Hwaseong, Korea
[2]Samsung Electronics, San Jose, CA
[3]Samsung Electronics, Suwon, Korea

Samsung Function-in-Memory DRAM (2021)

Programmable Computing Unit

- Configuration of PCU block
  - Interface unit to control data flow
  - Execution unit to perform operations
  - Register group
    - 32 entries of CRF for instruction memory
    - 16 GRF for weight and accumulation
    - 16 SRF to store constants for MAC operations

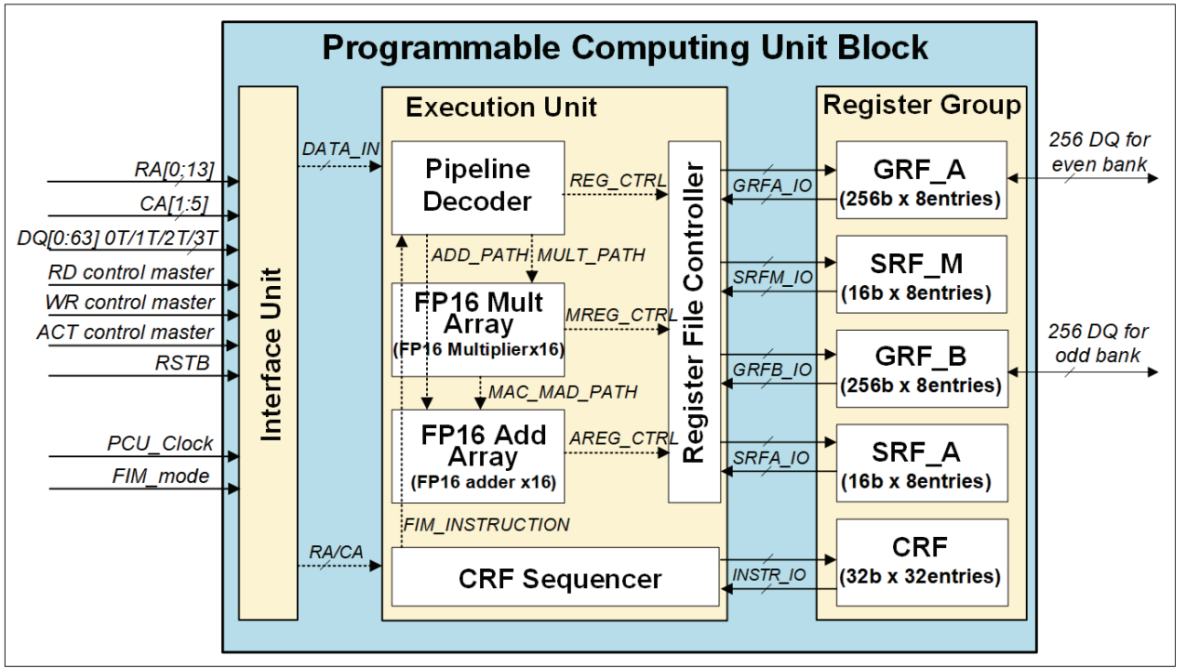

[Block diagram of PCU in FIMDRAM]

ISSCC 2021 / SESSION 25 / DRAM / 25.4

25.4 A 20nm 6GB Function-In-Memory DRAM, Based on HBM2 with a 1.2TFLOPS Programmable Computing Unit Using Bank-Level Parallelism, for Machine Learning Applications

Young-Cheon Kwon[1], Suk Han Lee[1], Jaehoon Lee[1], Sang-Hyuk Kwon[1], Je Min Ryu[1], Jong-Pil Son[1], Seongil O[1], Hak-Soo Yu[1], Haesuk Lee[1], Soo Young Kim[1], Youngmin Cho[1], Jin Guk Kim[1], Jongyoon Choi[1], Hyun-Sung Shin[1], Jin Kim[1], BengSeng Phuah[1], HyoungMin Kim[1], Myeong Jun Song[1], Ahn Choi[1], Daeho Kim[1], SooYoung Kim[1], Eun-Bong Kim[1], David Wang[2], Shinhaeng Kang[1], Yuhwan Ro[3], Seungwoo Seo[2], JoonHo Song[2], Jaeyoun Youn[1], Kyomin Sohn[1], Nam Sung Kim[1]

[1]Samsung Electronics, Hwaseong, Korea
[2]Samsung Electronics, San Jose, CA
[3]Samsung Electronics, Suwon, Korea

Samsung Function-in-Memory DRAM (2021)

[Available instruction list for FIM operation]

Type	CMD	Description
Floating Point	ADD	FP16 addition
	MUL	FP16 multiplication
	MAC	FP16 multiply-accumulate
	MAD	FP16 multiply and add
Data Path	MOVE	Load or store data
	FILL	Copy data from bank to GRFs
Control Path	NOP	Do nothing
	JUMP	Jump instruction
	EXIT	Exit instruction

ISSCC 2021 / SESSION 25 / DRAM / 25.4

25.4 A 20nm 6GB Function-In-Memory DRAM, Based on HBM2 with a 1.2TFLOPS Programmable Computing Unit Using Bank-Level Parallelism, for Machine Learning Applications

Young-Cheon Kwon[1], Suk Han Lee[1], Jaehoon Lee[1], Sang-Hyuk Kwon[1], Je Min Ryu[1], Jong-Pil Son[1], Seongil O[1], Hak-Soo Yu[1], Haesuk Lee[1], Soo Young Kim[1], Youngmin Cho[1], Jin Guk Kim[1], Jongyoon Choi[1], Hyun-Sung Shin[1], Jin Kim[1], BengSeng Phuah[1], HyoungMin Kim[1], Myeong Jun Song[1], Ahn Choi[1], Daeho Kim[1], SooYoung Kim[1], Eun-Bong Kim[1], David Wang[2], Shinhaeng Kang[1], Yuhwan Ro[3], Seungwoo Seo[2], JoonHo Song[2], Jaeyoun Youn[1], Kyomin Sohn[1], Nam Sung Kim[1]

[1]Samsung Electronics, Hwaseong, Korea
[2]Samsung Electronics, San Jose, CA
[3]Samsung Electronics, Suwon, Korea

Samsung Function-in-Memory DRAM (2021)

Chip Implementation

- Mixed design methodology to implement FIMDRAM
  - Full-custom + Digital RTL

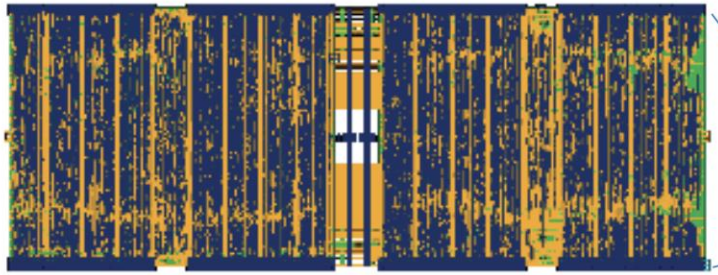

[Digital RTL design for PCU block]

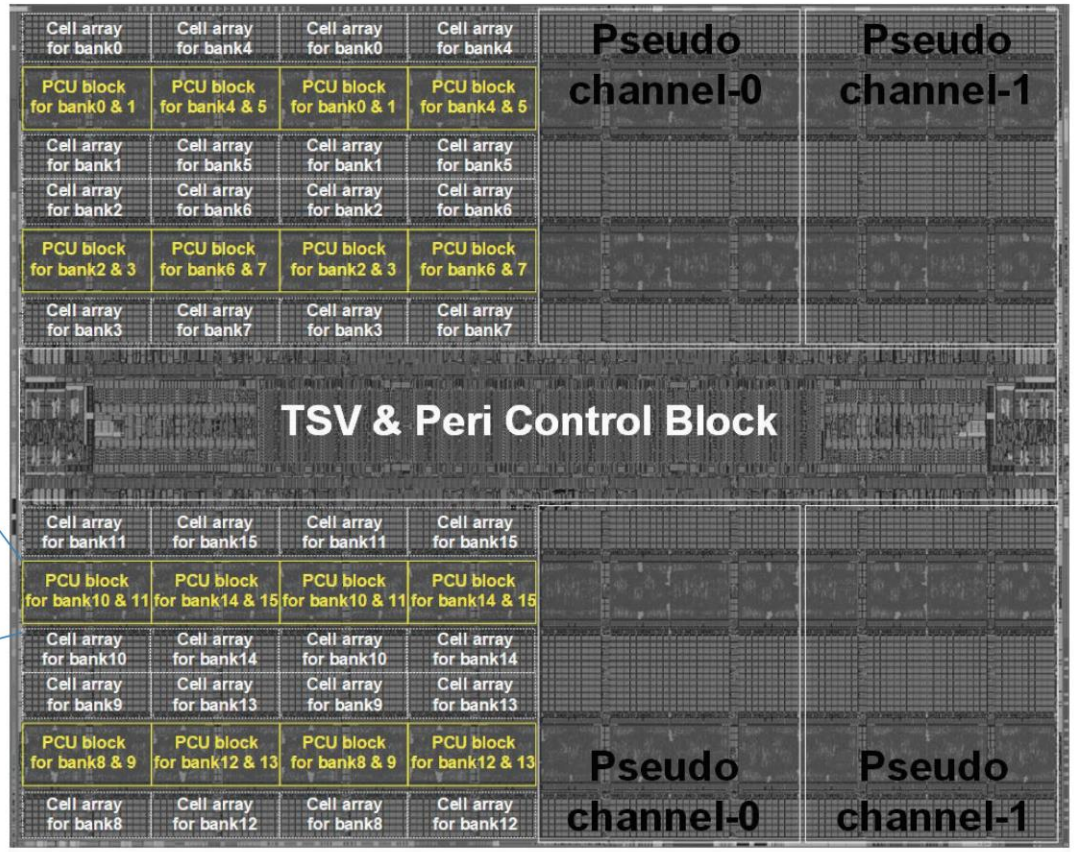

Samsung PNM Solutions for Generative AI (2023)

- Main target: transformer decoders used in ChatGPT, GPT-3
  - Compute-bound step: Summarization
  - Memory-bound step: Generation
    - Most of the execution time is spent on the memory copy from the host CPU memory to the CPU memory

- GEMV portion can be 60%-80% of total generation latency, which is the target of PIM/PNM

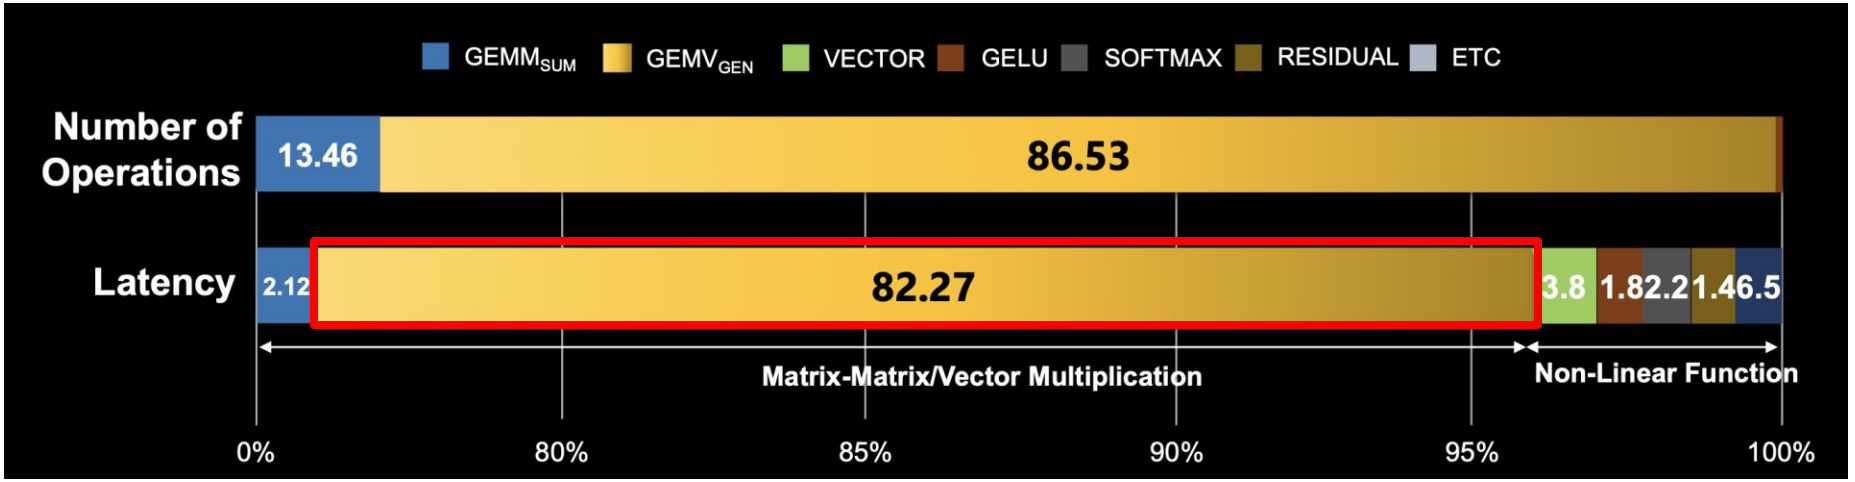

[J. H. Kim+ HC, 2023]

Solution I: Samsung's HBM-PIM (2023)

- AMD MI100 GPUs fabricated with HBM-PIM
- Experimental setup: GPT-J (6B, 32 input tokes), single AMD MI100-PIM GPU

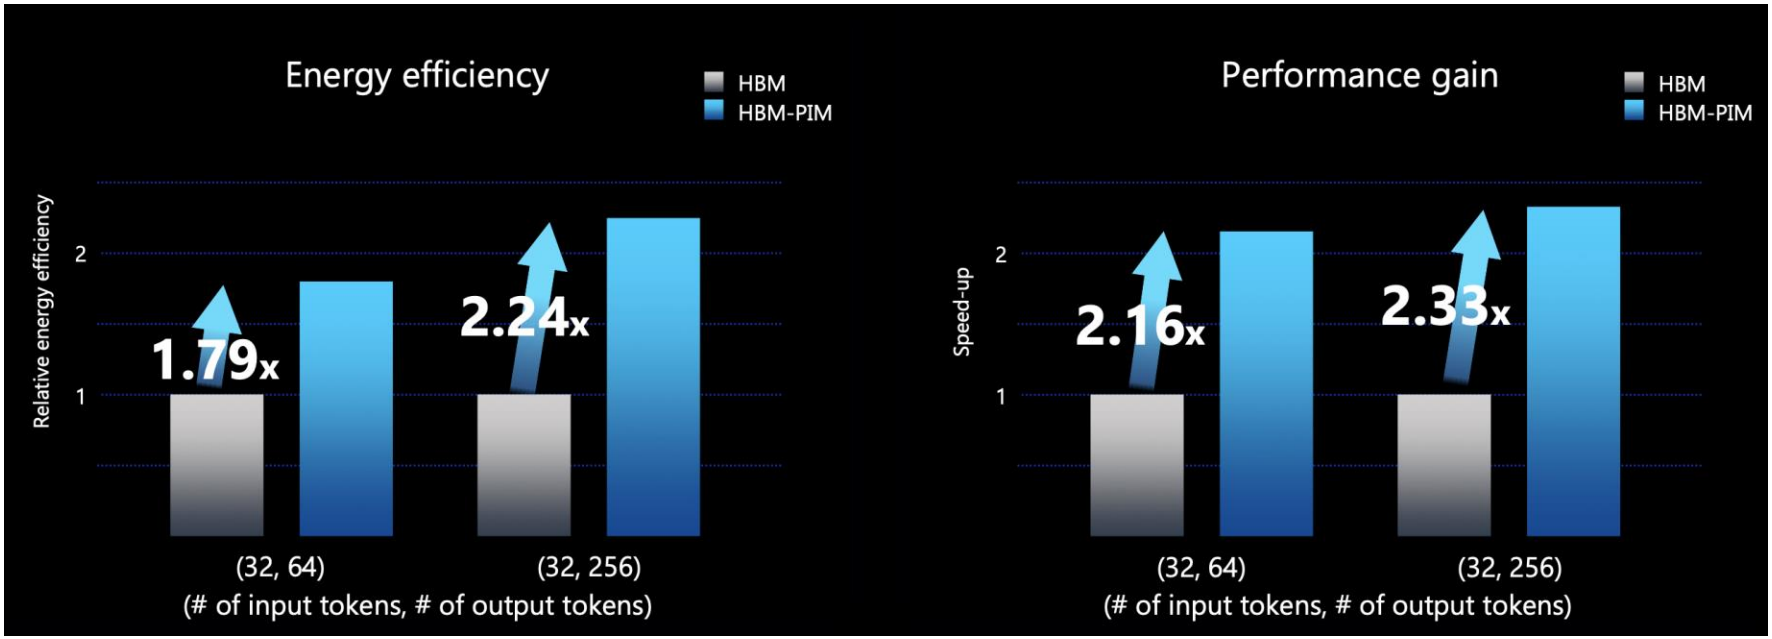

- GPT can be accelerated by more than 2x over baseline

[J. H. Kim+ HC, 2023]

Solution II: Samsung's LPDDR-PIM (2023)

- PIM for on-device generative AI
  - Datacenter costs and power consumption are increasing due to the growing demand for cloud AI
- LPDDR-PIM improves battery life by preventing memory over-provisioning just for bandwidth

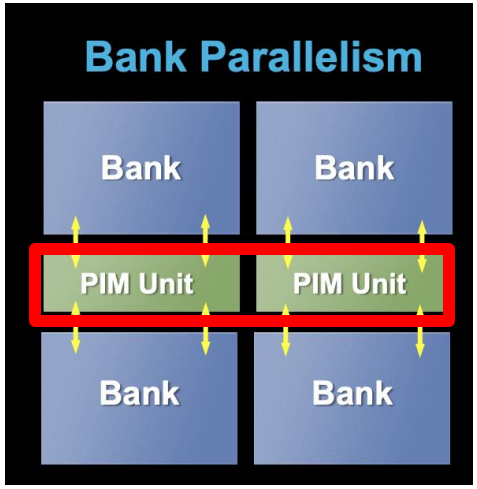

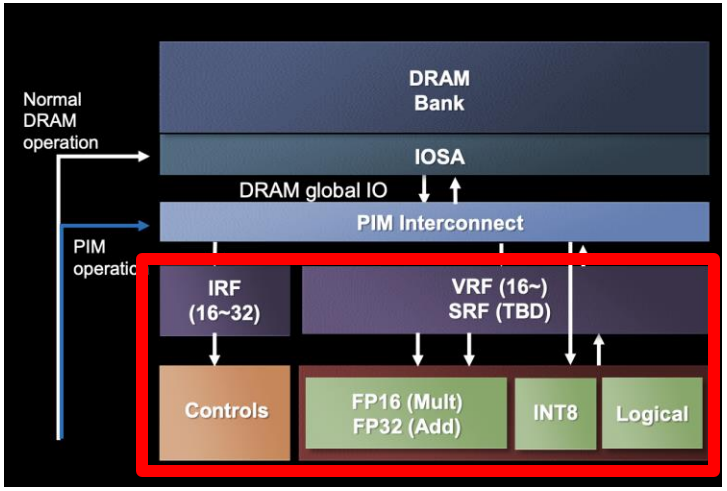

- 4.47x performance gains and 70.6% energy reduction in GPT-2

[J. H. Kim+ HC, 2023]

Solution III: Samsung's CXL-PNM (2023)

- A CXL-based processing-near-memory solution
  - Improves capacity, bandwidth, and power
  - Large-scale large-language models are often capacity-bound

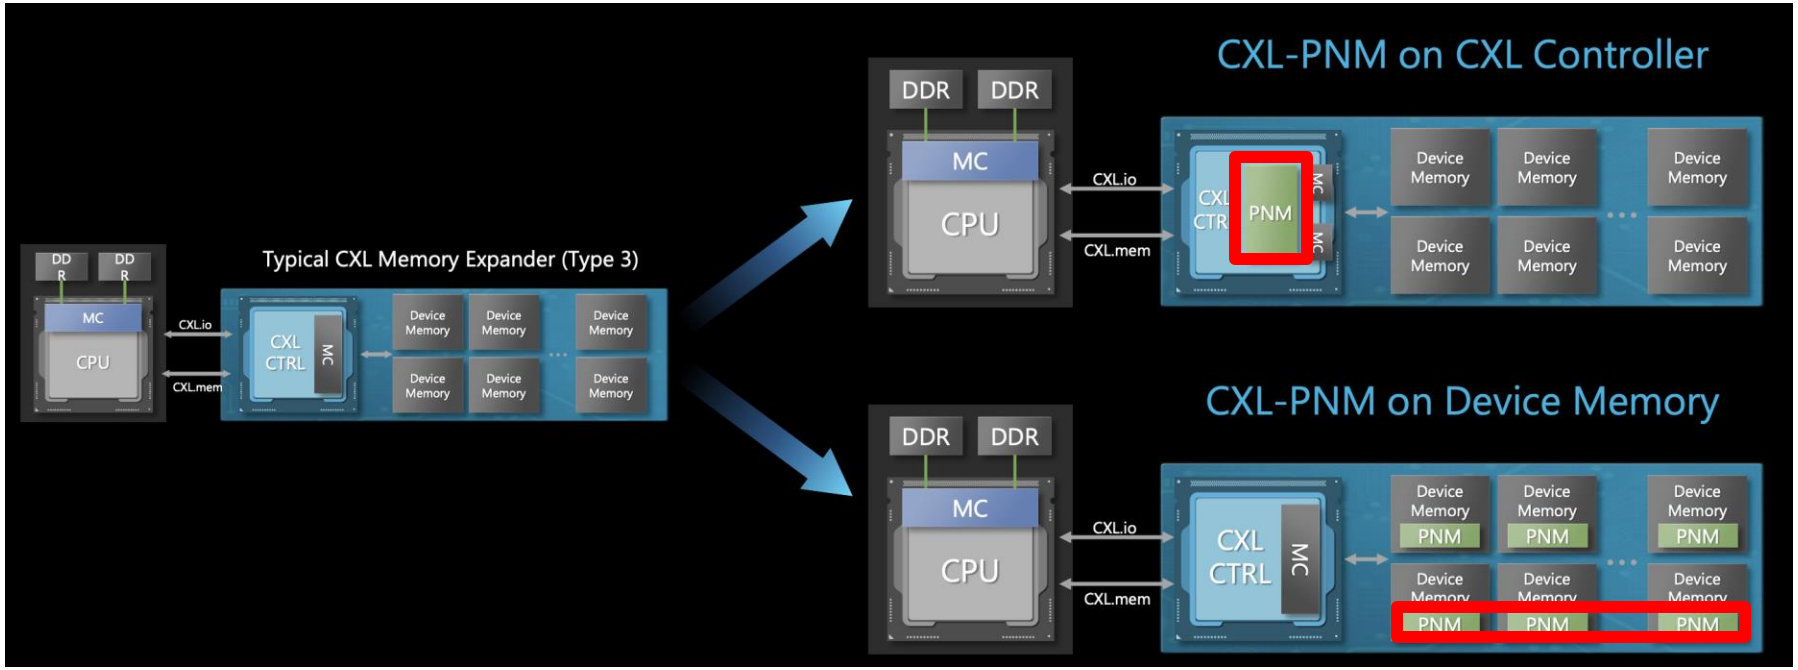

- Multiple CXL-PNM can offer 4.4x higher energy efficiency and 53% higher throughput than multiple GPUs

[J. H. Kim+ HC, 2023]

SK hynix AiM: Chip Implementation (2022)

- 4 Gb AiM die with 16 processing units (PUs)

AiM Die Photograph

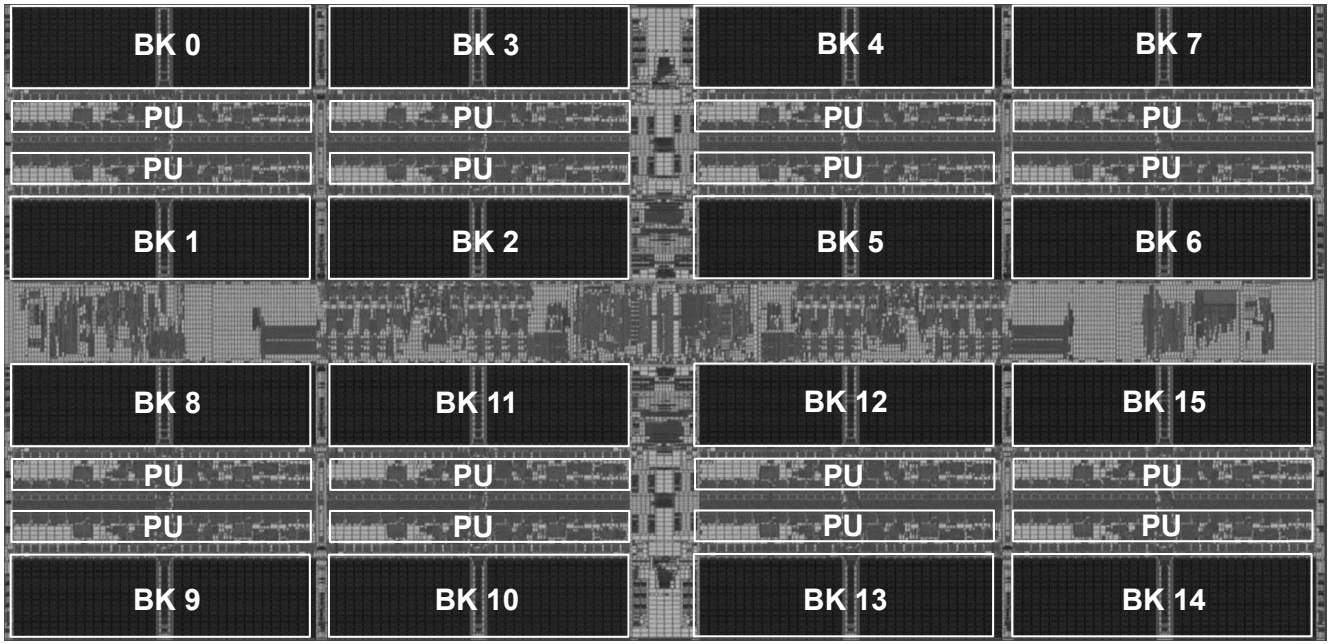

1 Process Unit (PU) Area

Total	0.19mm^2
MAC	0.11mm^2
Activation Function (AF)	0.02mm^2
Reservoir Cap.	0.05mm^2
Etc.	0.01mm^2

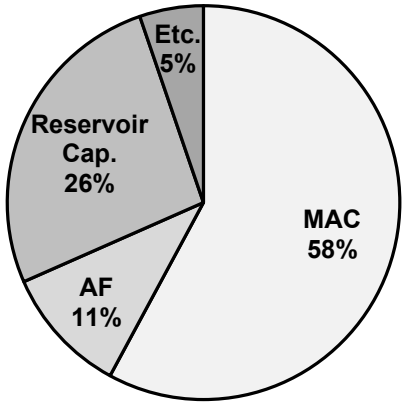

Lee et al., A 1ynm 1.25V 8Gb, 16Gb/s/pin GDDR6-based Accelerator-in-Memory supporting 1TFLOPS MAC Operation and Various Activation Functions for Deep-Learning Applications, ISSCC 2022

SK hynix AiM: System Organization (2022)

- GDDR6-based AiM architecture

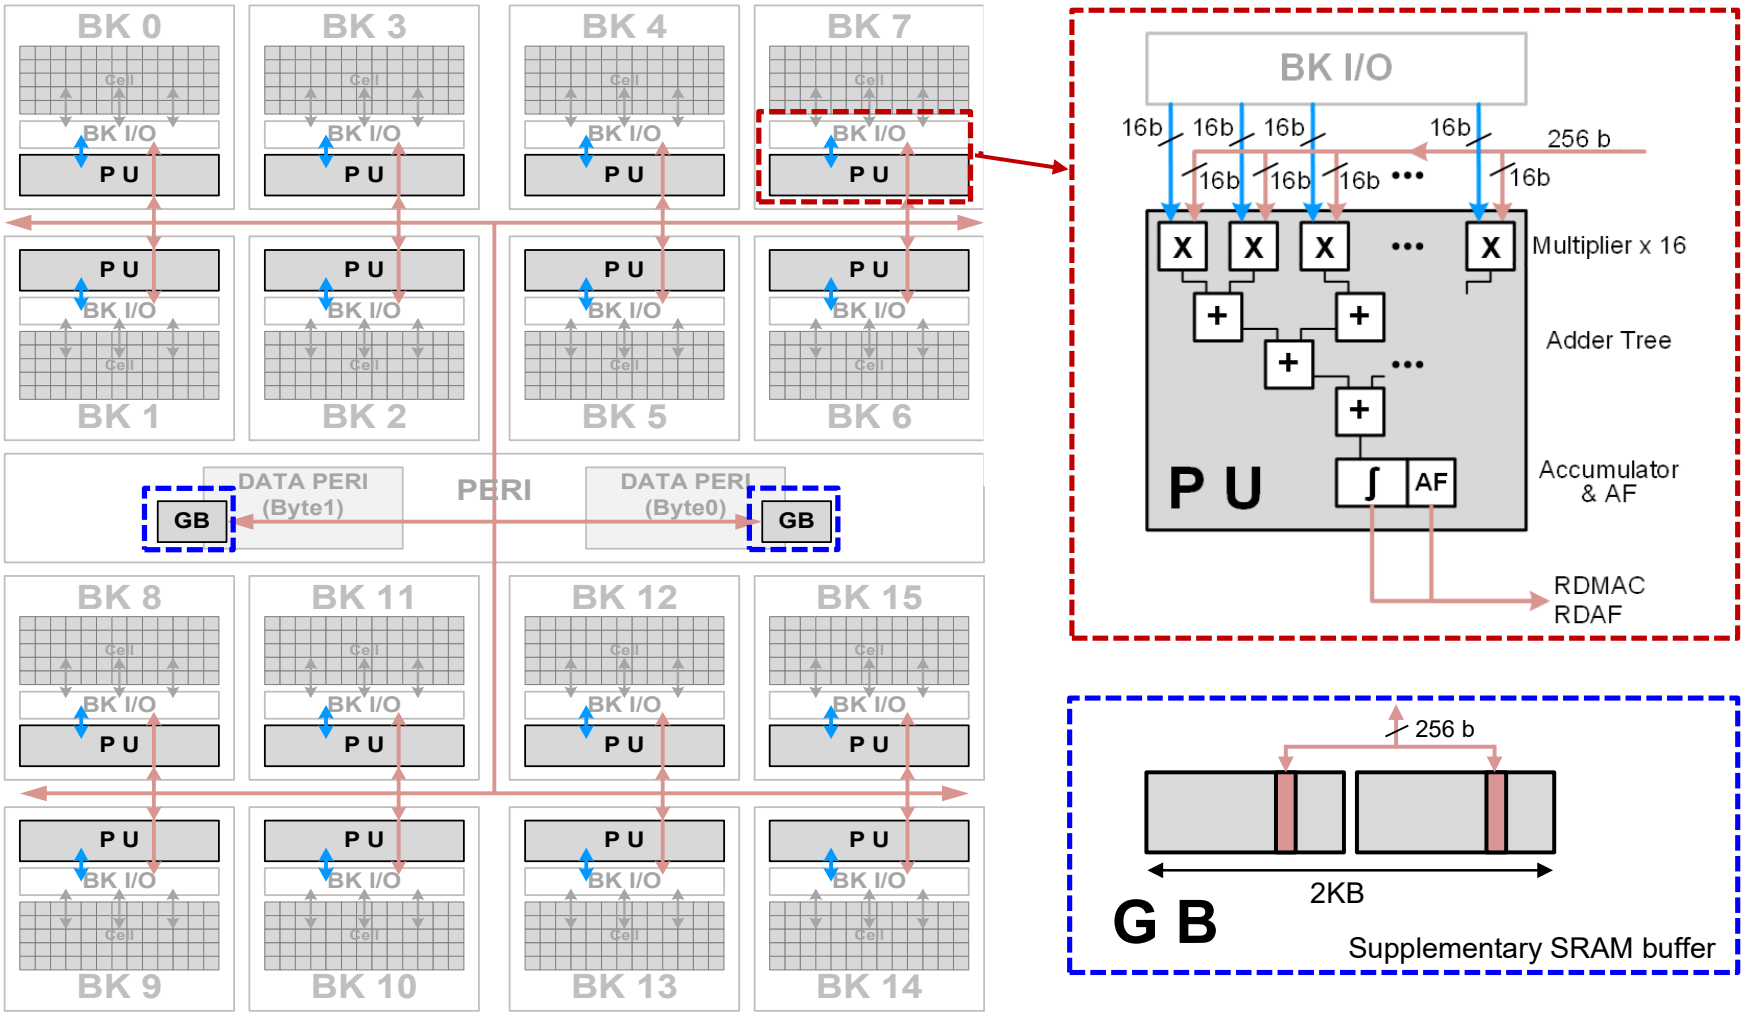

Lee et al., A 1ynm 1.25V 8Gb, 16Gb/s/pin GDDR6-based Accelerator-in-Memory supporting 1TFLOPS MAC Operation and Various Activation Functions for Deep-Learning Applications, ISSCC 2022

Possible PNM Designs

- General-purpose programmable cores
  - Wimpy cores (possibility of running any workload)
  - E.g. from academia: Tesseract PIM for Graph Processing
  - E.g. from industry: UPMEM PIM

- Fixed-function units
  - Hardware/software co-designed PIM for efficiency
  - E.g. from academia: Mensa for NN Edge Inference
  - E.g. from industry: Samsung HBM-PIM, SK hynix AiM

- **Reconfigurable architectures**
  - PNM cores coupled with FPGAs, CGRA
  - E.g. from academia: NERO for Weather Prediction
  - E.g. from industry: Samsung AxDIMM

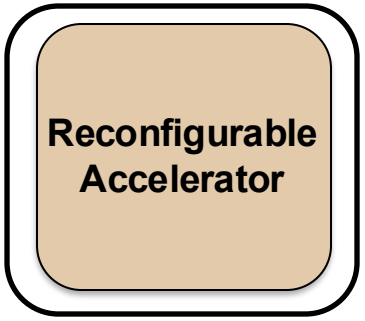

FPGA-based Processing Near Memory

- Gagandeep Singh, Dionysios Diamantopoulos, Christoph Hagleitner, Juan Gómez-Luna, Sander Stuijk, Onur Mutlu, and Henk Corporaal,
  "NERO: A Near High-Bandwidth Memory Stencil Accelerator for Weather Prediction Modeling"
  *Proceedings of the 30th International Conference on Field-Programmable Logic and Applications (**FPL**)*, Gothenburg, Sweden, September 2020.

NERO: A Near High-Bandwidth Memory Stencil Accelerator for Weather Prediction Modeling

Gagandeep Singh[a,b,c] Dionysios Diamantopoulos[c] Christoph Hagleitner[c] Juan Gómez-Luna[b]
Sander Stuijk[a] Onur Mutlu[b] Henk Corporaal[a]
[a]Eindhoven University of Technology [b]ETH Zürich [c]IBM Research Europe, Zurich

Heterogeneous System: CPU+FPGA

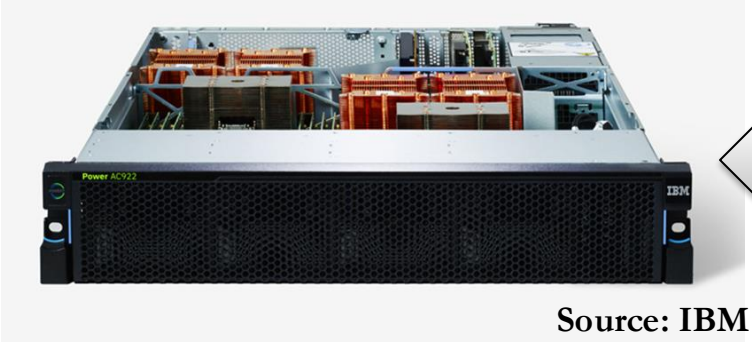

Source: IBM

POWER9 AC922

CAPI2

Source: AlphaData

DDR4-based AD9V3 board

We evaluate two POWER9+FPGA systems:

1. HBM-based board AD9H7

Xilinx Virtex Ultrascale+™ XCVU37P-2

AD9V3

2. DDR4-based board

Xilinx Virtex Ultrascale+™ XCVU3P-2

NERO Design Flow

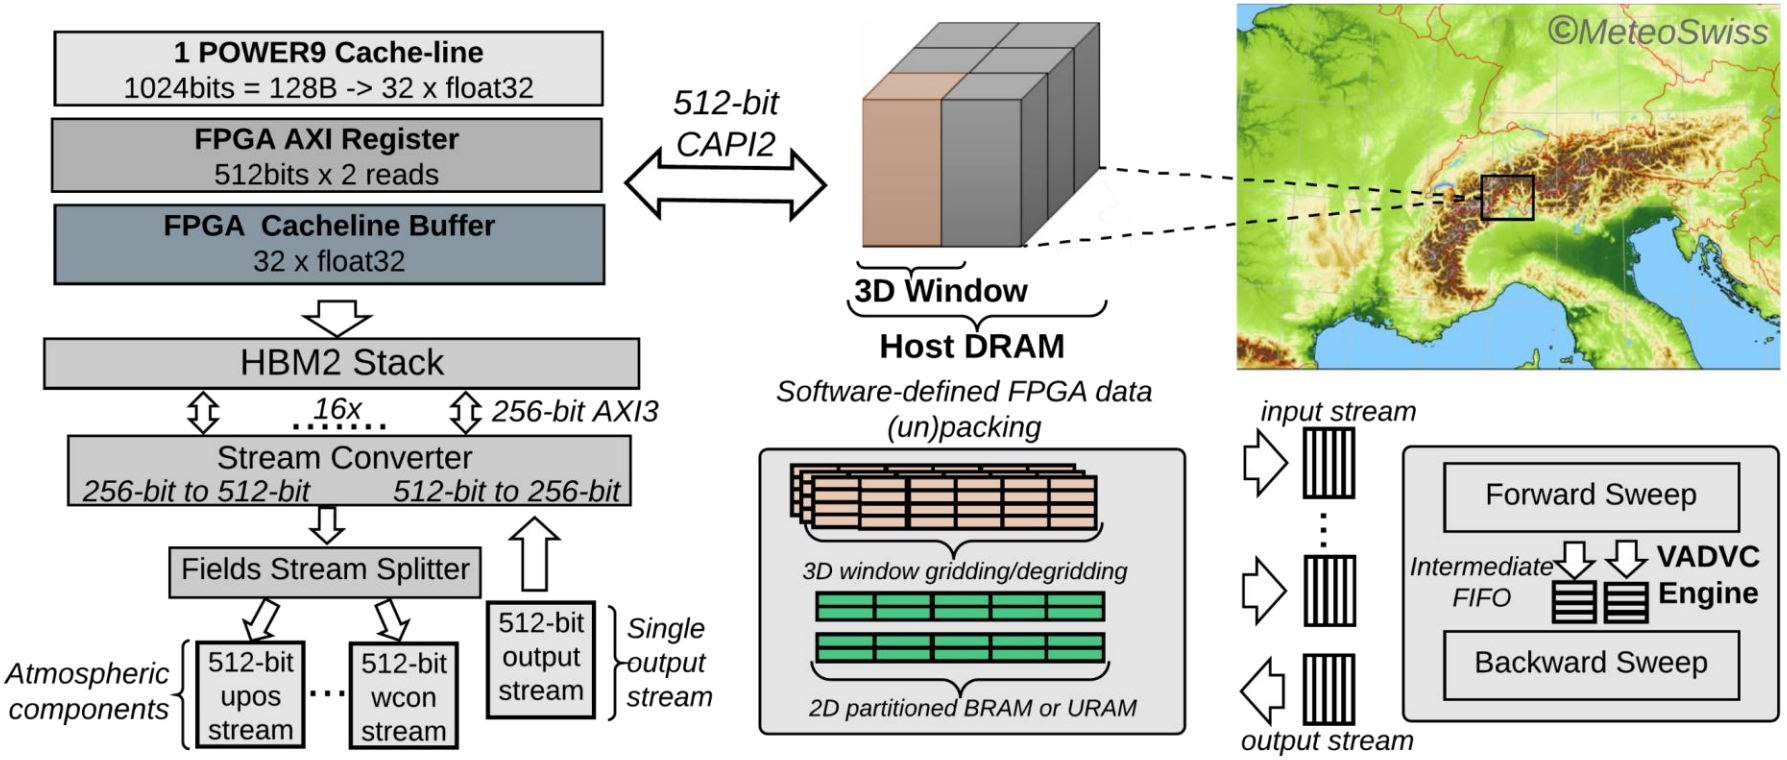

NERO Performance Analysis

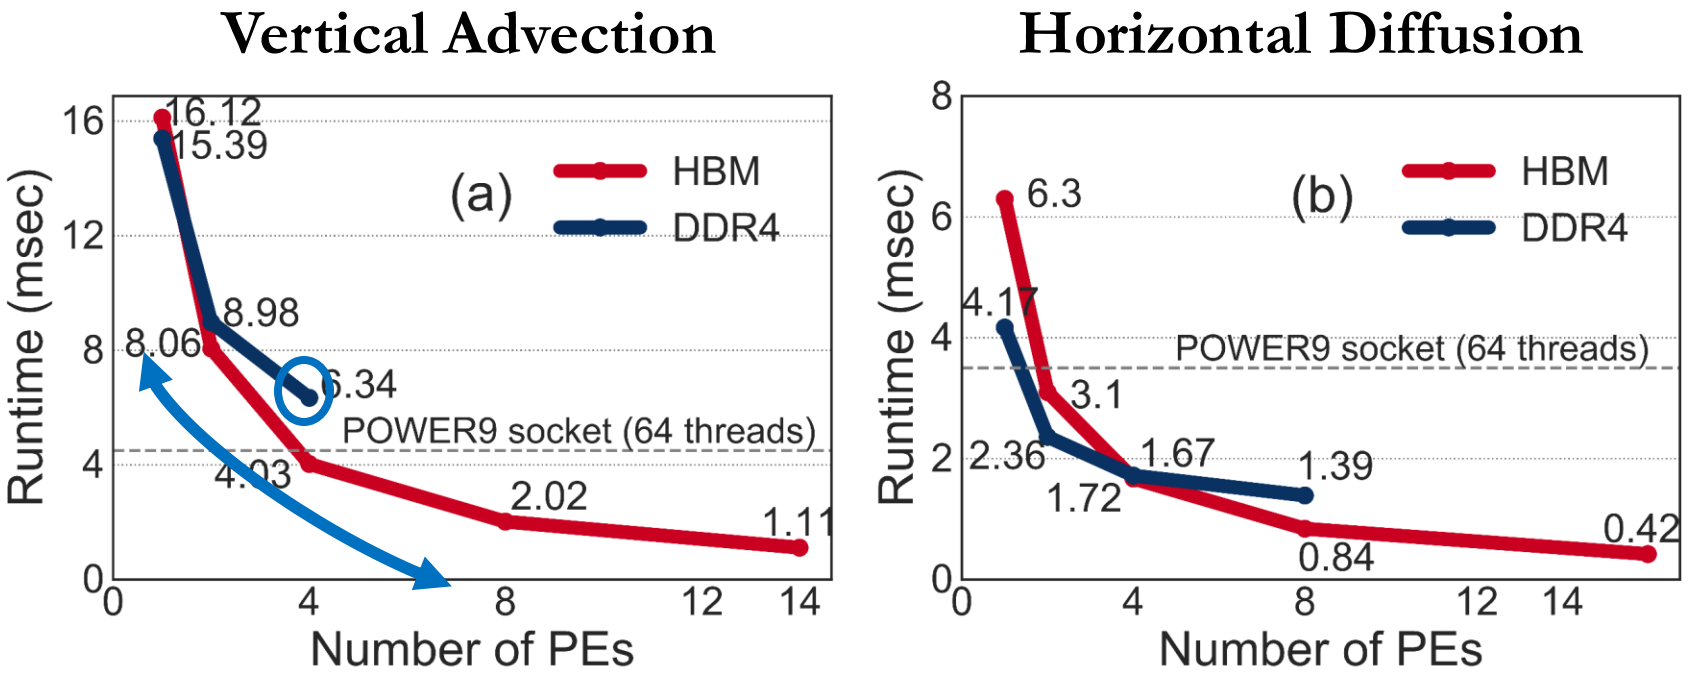

NERO is 4.2x and 8.3x faster than a complete POWER9 socket

Possible PNM Designs

- General-purpose programmable cores
  - Wimpy cores (possibility of running any workload)
  - E.g. from academia: Tesseract PIM for Graph Processing
  - E.g. from industry: UPMEM PIM

- Fixed-function units
  - Hardware/software co-designed PIM for efficiency
  - E.g. from academia: Mensa for NN Edge Inference
  - E.g. from industry: Samsung HBM-PIM, SK hynix AiM

- **Reconfigurable architectures**
  - PNM cores coupled with FPGAs, CGRA
  - E.g. from academia: NERO for Weather Prediction
  - E.g. from industry: Samsung AxDIMM

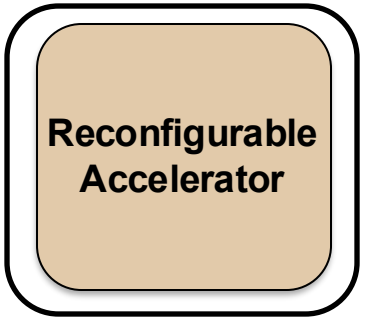

Samsung AxDIMM (2021)

Samsung Newsroom

CORPORATE | PRODUCTS | INSIGHTS | PRESS RESOURCES

Samsung Brings In-Memory Processing Power to Wider Range of Applications

Korea on August 24, 2021 Audio Share

Integration of HBM-PIM with the Xilinx Alveo AI accelerator system will boost overall system performance by 2.5X while reducing energy consumption by more than 60%

PIM architecture will be broadly deployed beyond HBM, to include mainstream DRAM modules and mobile memory

Samsung Electronics, the world leader in advanced memory technology, today showcased its latest advancements with processing-in-memory (PIM) technology at Hot Chips 33—a leading semiconductor conference where the most notable microprocessor and IC innovations are unveiled each year. Samsung's revelations include the first successful integration of its PIM-enabled High Bandwidth Memory (HBM-PIM) into a commercialized accelerator system, and broadened PIM applications to embrace DRAM modules and mobile memory, in accelerating the move toward the convergence of memory and logic.

DRAM Modules Powered by PIM

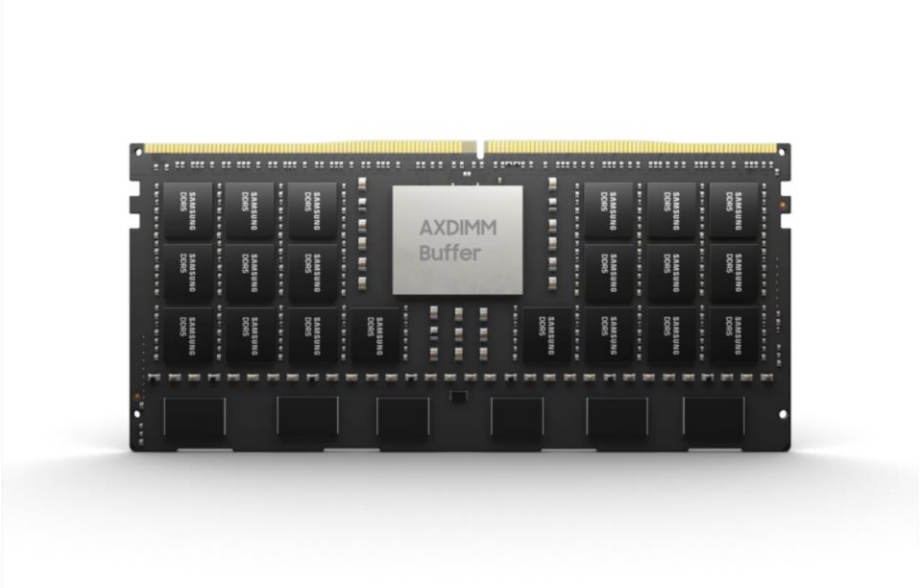

The Acceleration DIMM (AXDIMM) brings processing to the DRAM module itself, minimizing large data movement between the CPU and DRAM to boost the energy efficiency of AI accelerator systems. With an AI engine built inside the buffer chip, the AXDIMM can perform parallel processing of multiple memory ranks (sets of DRAM chips) instead of accessing just one rank at a time, greatly enhancing system performance and efficiency. Since the module can retain its traditional DIMM form factor, the AXDIMM facilitates drop-in replacement without requiring system modifications. Currently being tested on customer servers, the AXDIMM can offer approximately twice the performance in AI-based recommendation applications and a 40% decrease in system-wide energy usage.

Samsung AxDIMM (2021)

- DIMM-based PIM
  - DLRM recommendation system

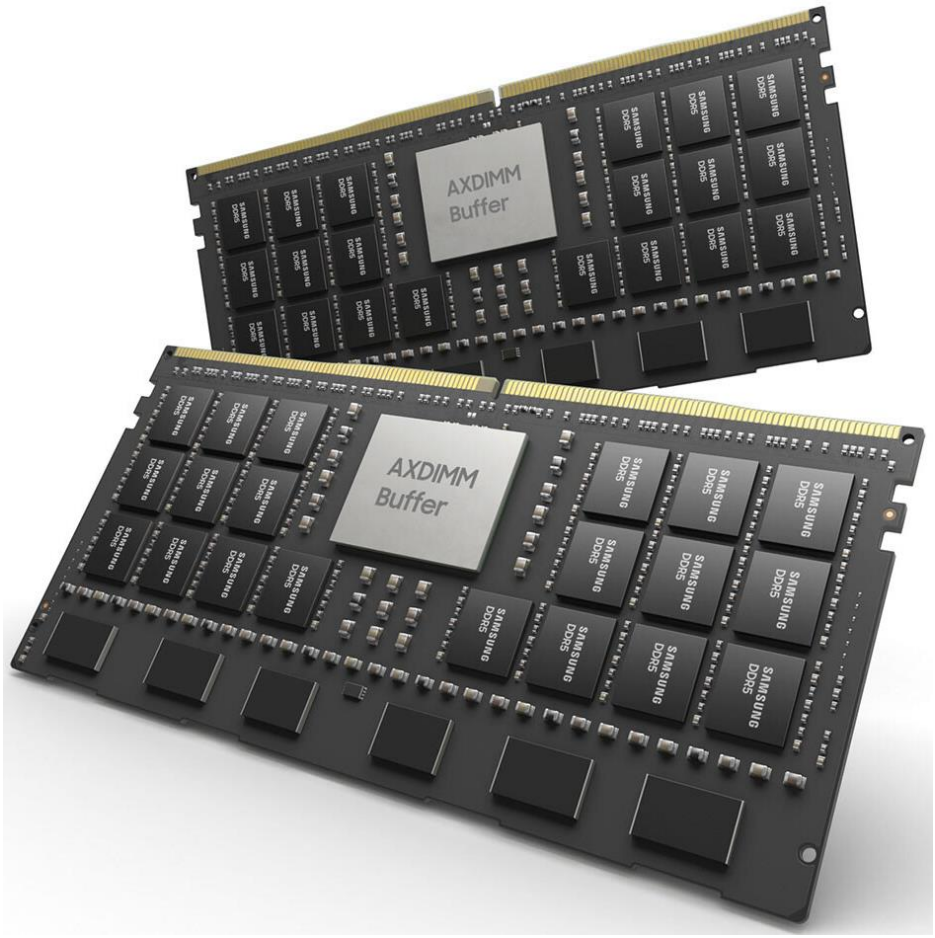

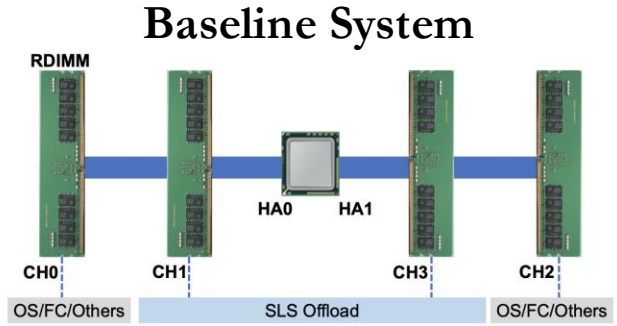

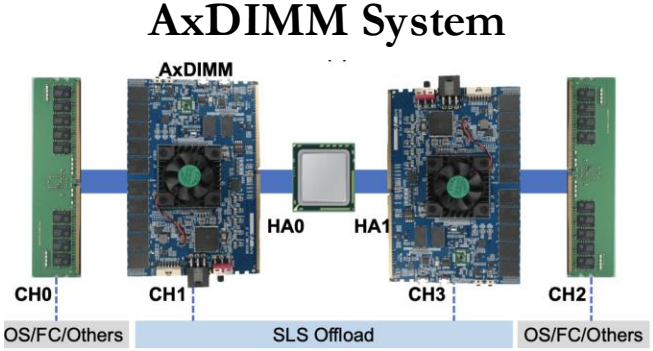

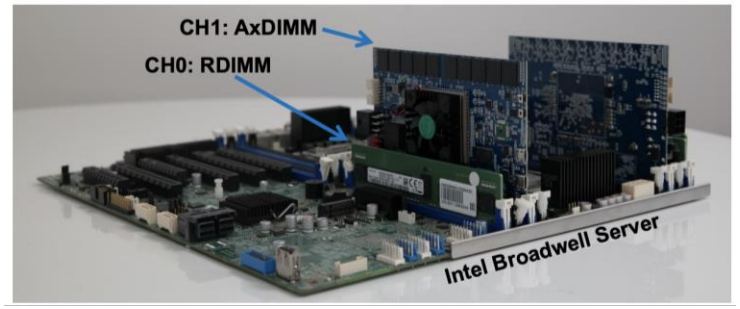

[Ke+, IEEE Micro'2021]

AxDIMM Design: Hardware Architecture

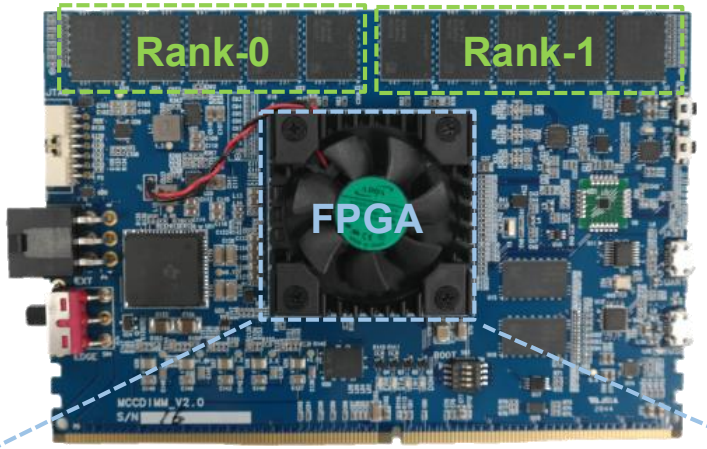

FPGA board with standard DIMM interface:
It serves as a real-system
near-memory processing implementation

AxDIMM Design: Hardware Architecture

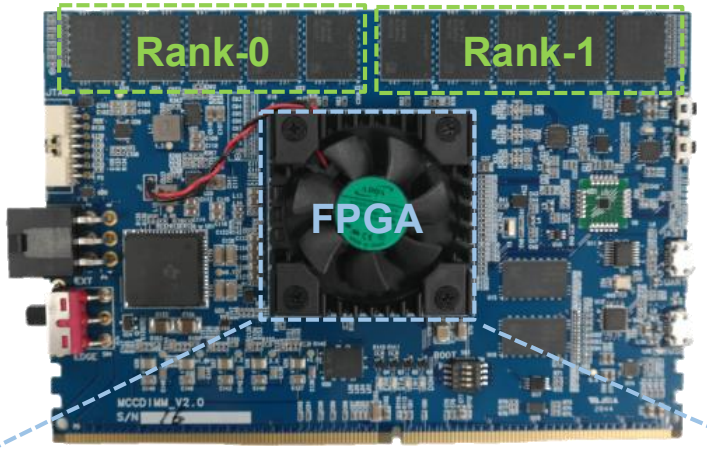

Standard DIMM Interface

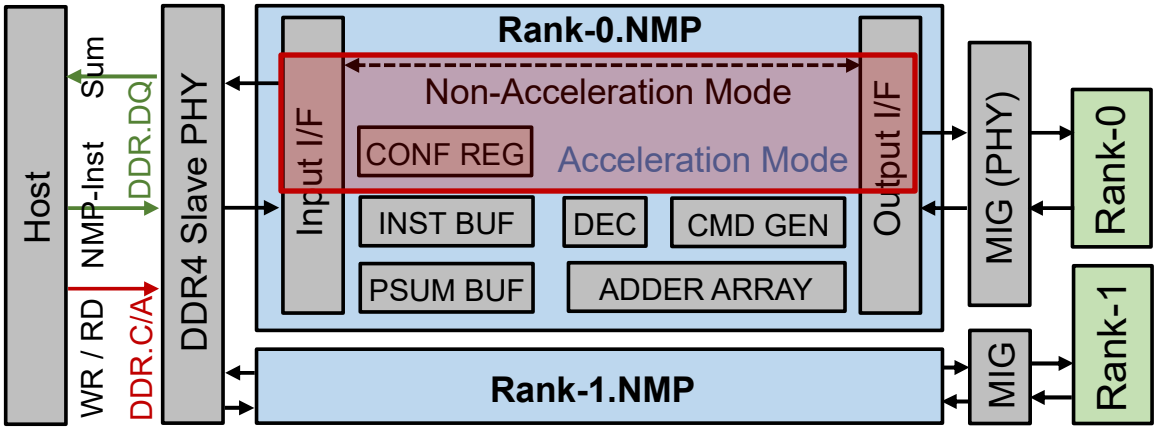

Two execution modes:
(1) non-acceleration mode
(2) acceleration mode (blocking)

[Ke+, IEEE Micro'2021]

Sparse Length Sum with AxDIMM (IEEE Micro 2021)

Near-Memory Processing in Action: Accelerating Personalized Recommendation with AxDIMM

Liu Ke*†, Xuan Zhang†, Jinin So‡, Jong-Geon Lee‡, Shin-Haeng Kang‡, Sukhan Lee‡, Songyi Han‡, YeonGon Cho‡, JIN Hyun Kim‡, Yongsuk Kwon‡, KyungSoo Kim‡, Jin Jung‡, Ilkwon Yun‡, Sung Joo Park‡, Hyunsun Park‡, Joonho Song‡, Jeonghyeon Cho‡, Kyomin Sohn‡, Nam Sung Kim‡, Hsien-Hsin S. Lee*

*Facebook, †Washington University in St. Louis, ‡Samsung

https://doi.org/10.1109/MM.2021.3097700

Sparse Length Sum with AxDIMM (AICAS 2022)

An Architecture of Sparse Length Sum Accelerator in AxDIMM

Shin-haeng Kang
DRAM Design Team 1
Samsung Electronics
Hwasung, South Korea
s-h.kang@samsung.com

Byeongho Kim
DRAM Design Team 1
Samsung Electronics
Hwasung, South Korea
bh1122.kim@samsung.com

Sukhan Lee
DRAM Design Team 1
Samsung Electronics
Hwasung, South Korea
sh1026.lee@samsung.com

Kyomin Sohn
DRAM Design Team 1
Samsung Electronics
Hwasung, South Korea
kyomin.sohn@samsung.com

https://doi.org/10.1109/AICAS54282.2022.9869896

Database Operations with AxDIMM (DaMoN 2022)

Improving In-Memory Database Operations with Acceleration DIMM (AxDIMM)

Donghun Lee
Minseon Ahn
Jungmin Kim
dong.hun.lee@sap.com
minseon.ahn@sap.com
jungmin.kim@sap.com
SAP Labs Korea

Oliver Rebholz
oliver.rebholz@sap.com
SAP SE

Jinin So
Jong-Geon Lee
Jeonghyeon Cho
Vishnu Charan Thummala
jinin.so@samsung.com
jg1021.lee@samsung.com
caleb1@samsung.com
vishnu.c.t@samsung.com
Samsung Electronics

Ravi Shankar JV
Sachin Suresh Upadhya
Mohammed Ibrahim Khan
Jin Hyun Kim
venkata.ravi@samsung.com
sachin1.s@samsung.com
ibrahim.khan@samsung.com
kjh5555@samsung.com
Samsung Electronics

https://doi.org/10.1145/3533737.3535093

Longer Lecture on AxDIMM

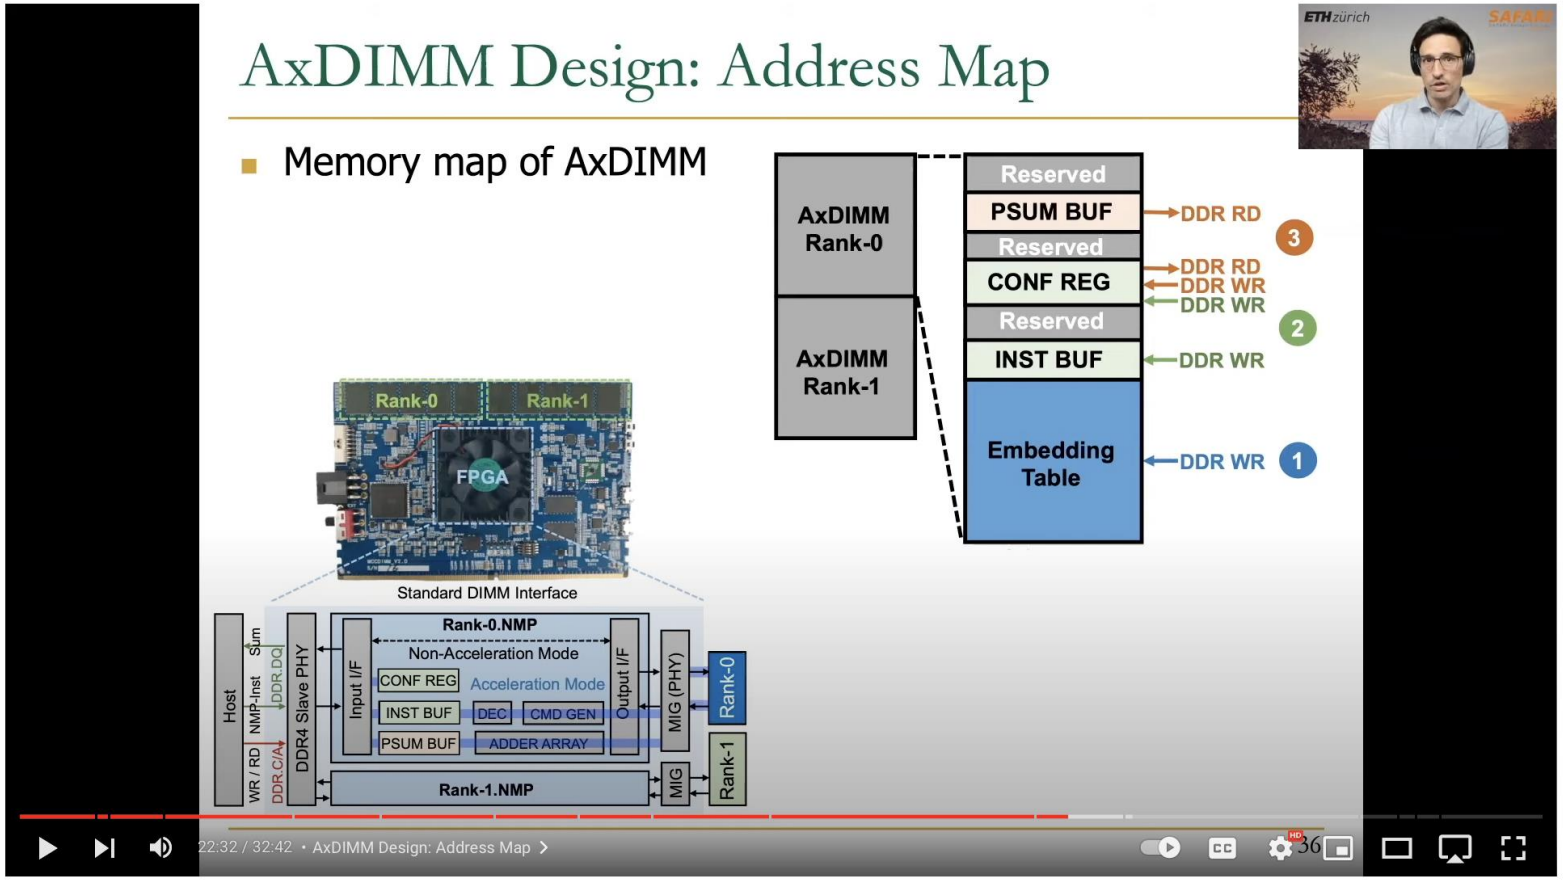

PIM Course: Lecture 7: Real-world PIM: Samsung AxDIMM - Fall 2022

Onur Mutlu Lectures
32.4K subscribers

Subscribed | 21 | Share | Clip | Save

846 views 4 months ago Livestream - P&S Data-Centric Architectures: Fundamentally Improving Performance and Energy (Fall 2022)
Projects & Seminars, ETH Zürich, Fall 2022
Data-Centric Architectures: Fundamentally Improving Performance and Energy
(https://safari.ethz.ch/projects_and_s...) Show more

https://youtu.be/SXdzQZAKG-Y

Another Longer Lecture on AxDIMM

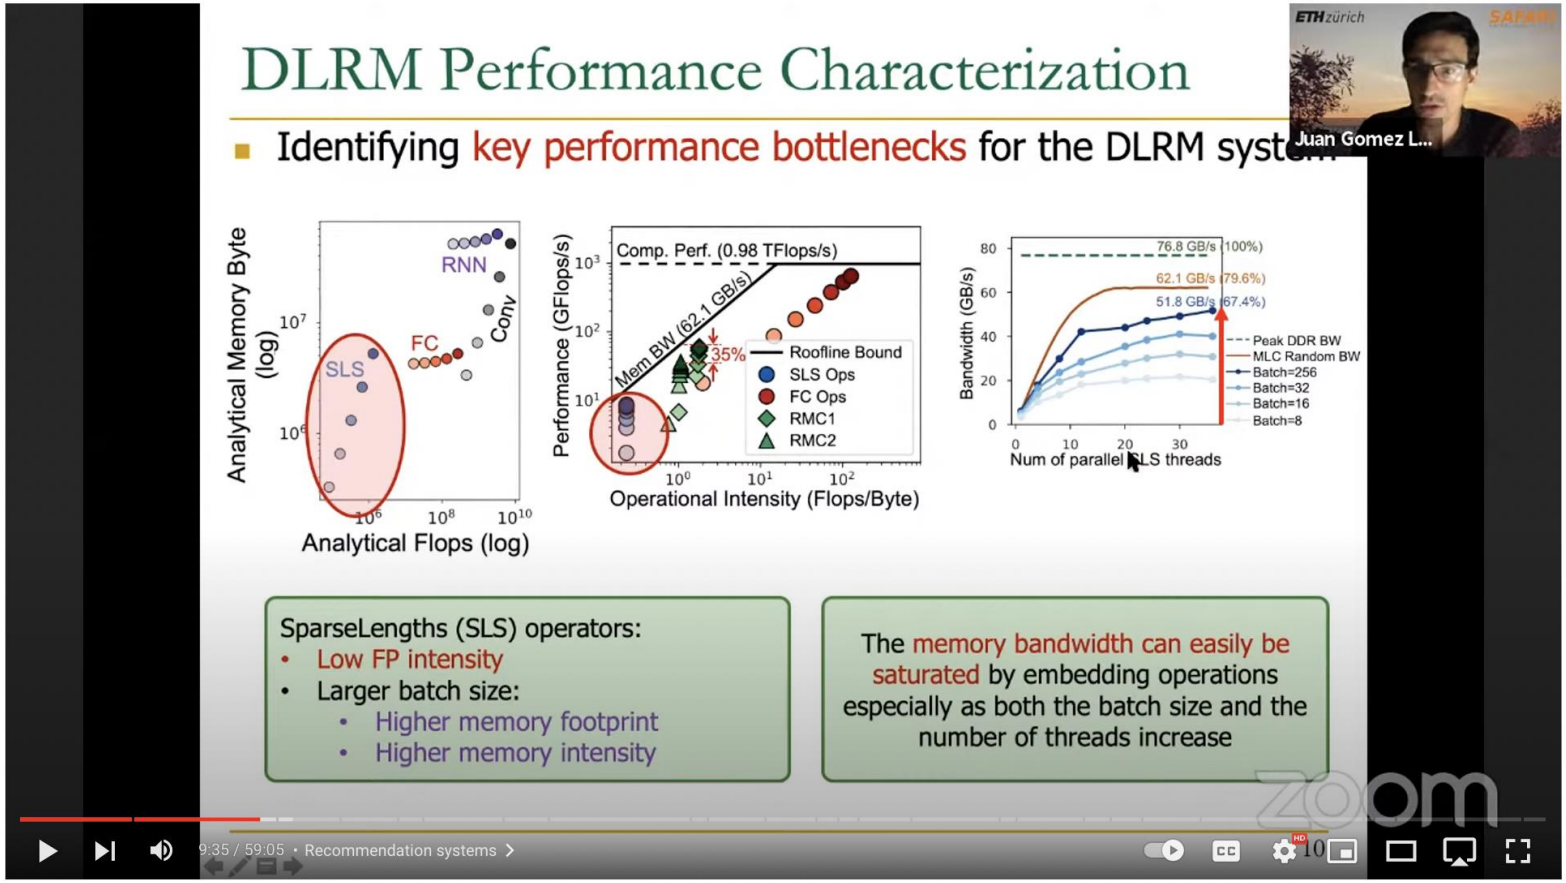

Processing in Memory Course: Meeting 5: Real-world PIM architectures IV - Fall'21

Onur Mutlu Lectures
32.4K subscribers

Subscribed | 18 | Share | Clip | Save

676 views Streamed 1 year ago Livestream - P&S Exploring the Processing-in-Memory Paradigm for Future Computing Systems (Fall 2021)
Project & Seminar, ETH Zürich, Fall 2021
Exploring the Processing-in-Memory Paradigm for Future Computing Systems (https://safari.ethz.ch/projects_and_s...)
Show more

https://youtu.be/2FMQg786GKs

Processing-in-Memory Landscape Today

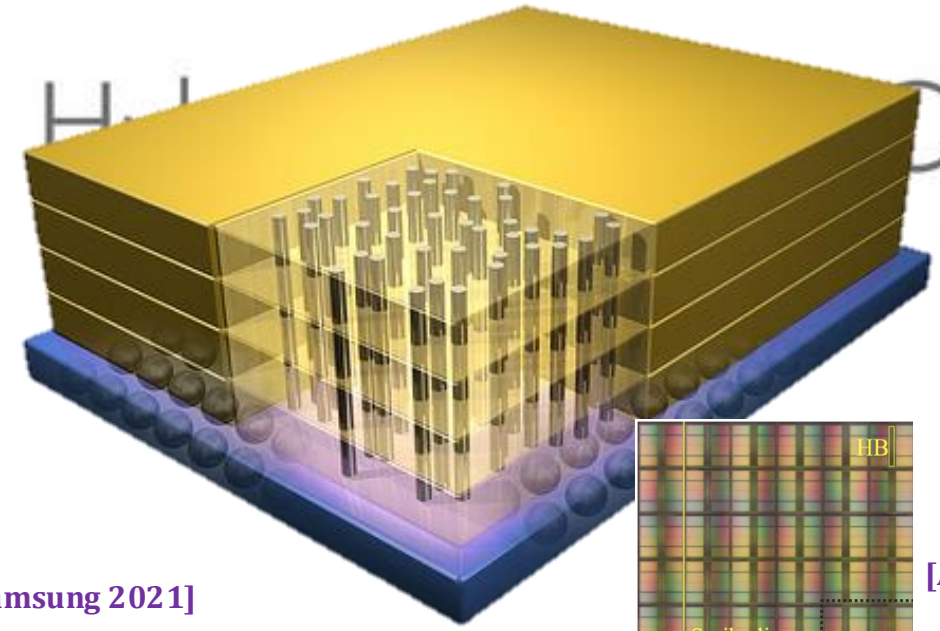

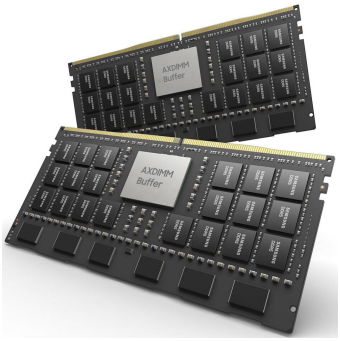

[Samsung 2021]

[Alibaba 2022]

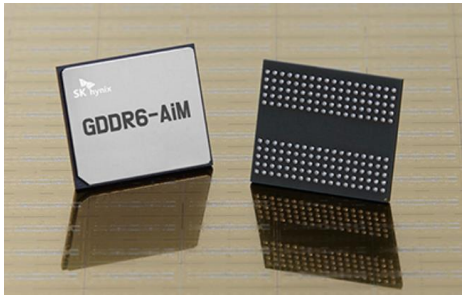

[SK Hynix 2022]

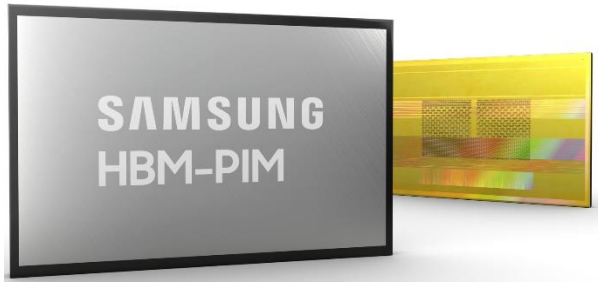

[Samsung 2021]

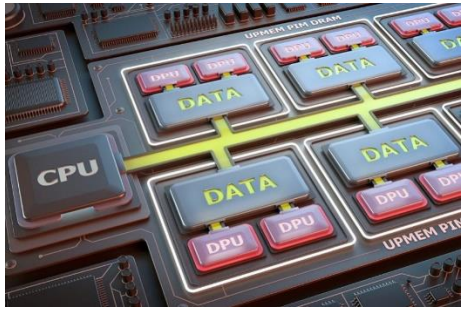

[UPMEM 2019]

This does not include many experimental chips and startups

Possible PNM Designs

- General-purpose programmable cores
  - Wimpy cores (possibility of running any workload)
  - E.g. from academia: Tesseract PIM for Graph Processing
  - E.g. from industry: UPMEM PIM

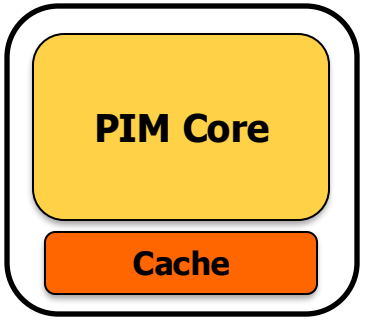

- Fixed-function units
  - Hardware/software co-designed PIM for efficiency
  - E.g. from academia: Mensa for NN Edge Inference
  - E.g. from industry: Samsung HBM-PIM, SK hynix AiM

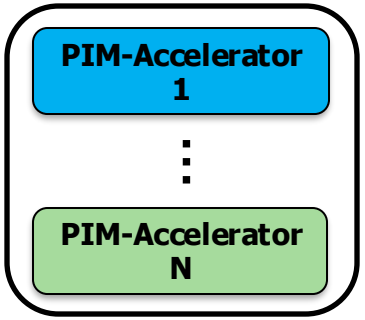

- Reconfigurable architectures
  - PNM cores coupled with FPGAs, CGRA
  - E.g. from academia: NERO for Weather Prediction
  - E.g. from industry: Samsung AxDIMM

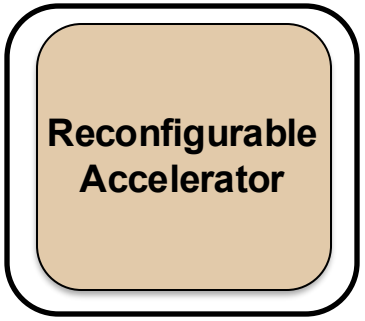

Research Tools PNM: DAMOV-SIM

- Geraldo F. Oliveira, Juan Gomez-Luna, Lois Orosa, Saugata Ghose, Nandita Vijaykumar, Ivan fernandez, Mohammad Sadrosadati, and Onur Mutlu, **"DAMOV: A New Methodology and Benchmark Suite for Evaluating Data Movement Bottlenecks"**
  IEEE Access, 8 September 2021.
  Preprint in ***arXiv***, 8 May 2021.
  [arXiv preprint]
  [IEEE Access version]
  [DAMOV Suite and Simulator Source Code]
  [SAFARI Live Seminar Video (2 hrs 40 mins)]
  [Short Talk Video (21 minutes)]

DAMOV: A New Methodology and Benchmark Suite for Evaluating Data Movement Bottlenecks

GERALDO F. OLIVEIRA, ETH Zürich, Switzerland
JUAN GÓMEZ-LUNA, ETH Zürich, Switzerland
LOIS OROSA, ETH Zürich, Switzerland
SAUGATA GHOSE, University of Illinois at Urbana–Champaign, USA
NANDITA VIJAYKUMAR, University of Toronto, Canada
IVAN FERNANDEZ, University of Malaga, Spain & ETH Zürich, Switzerland
MOHAMMAD SADROSADATI, ETH Zürich, Switzerland
ONUR MUTLU, ETH Zürich, Switzerland

Step 3: Memory Bottleneck Classification (2/2)

- **Goal:** identify the specific sources of data movement bottlenecks

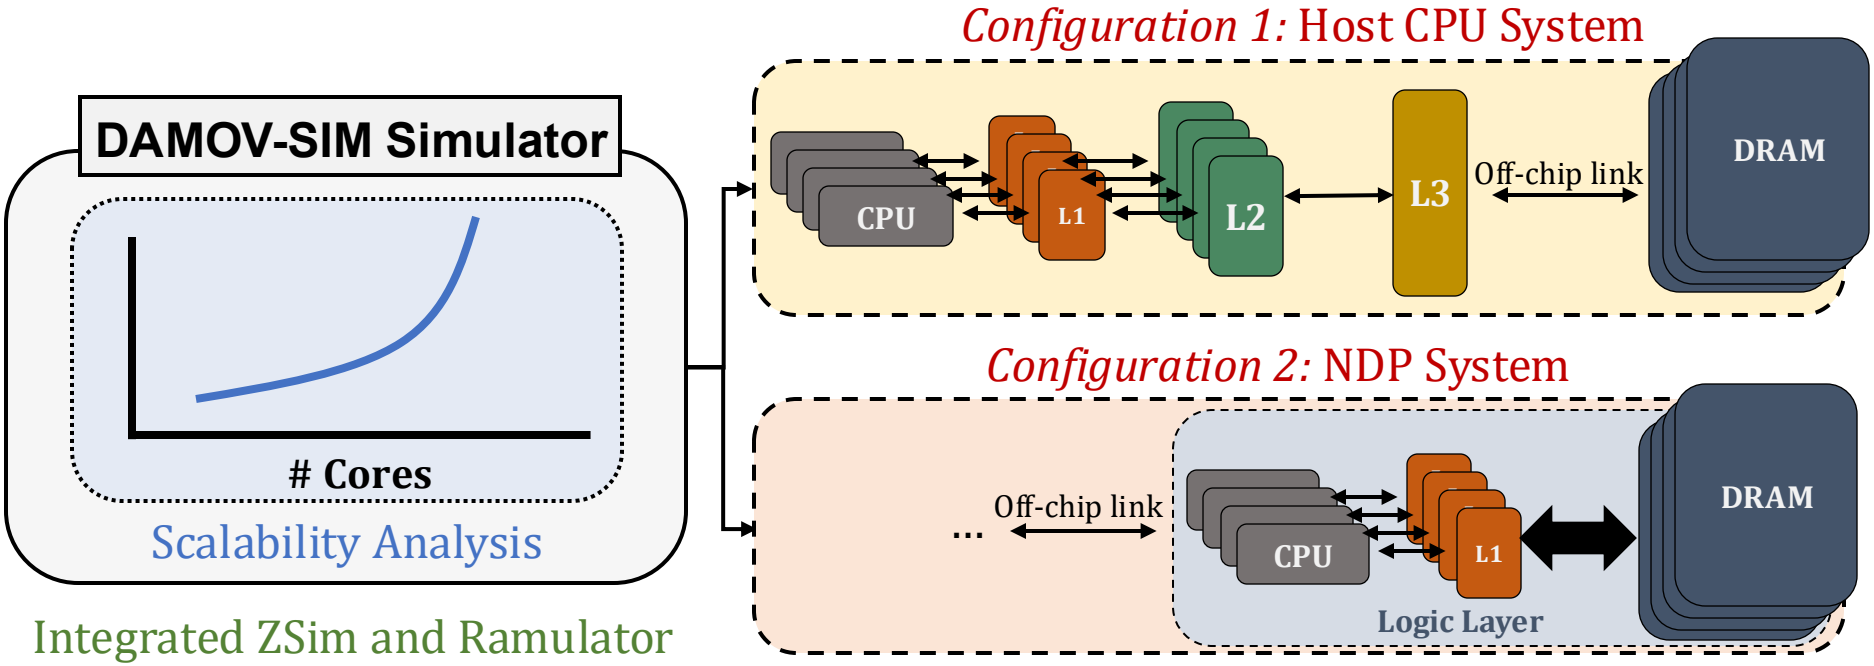

- Scalability Analysis:
  - 1, 4, 16, 64, and 256 out-of-order/in-order host and NDP CPU cores
  - 3D-stacked memory as main memory

DAMOV-SIM: https://github.com/CMU-SAFARI/DAMOV

DAMOV is Open Source

- We open-source our benchmark suite and our toolchain

DAMOV-SIM

DAMOV Benchmarks

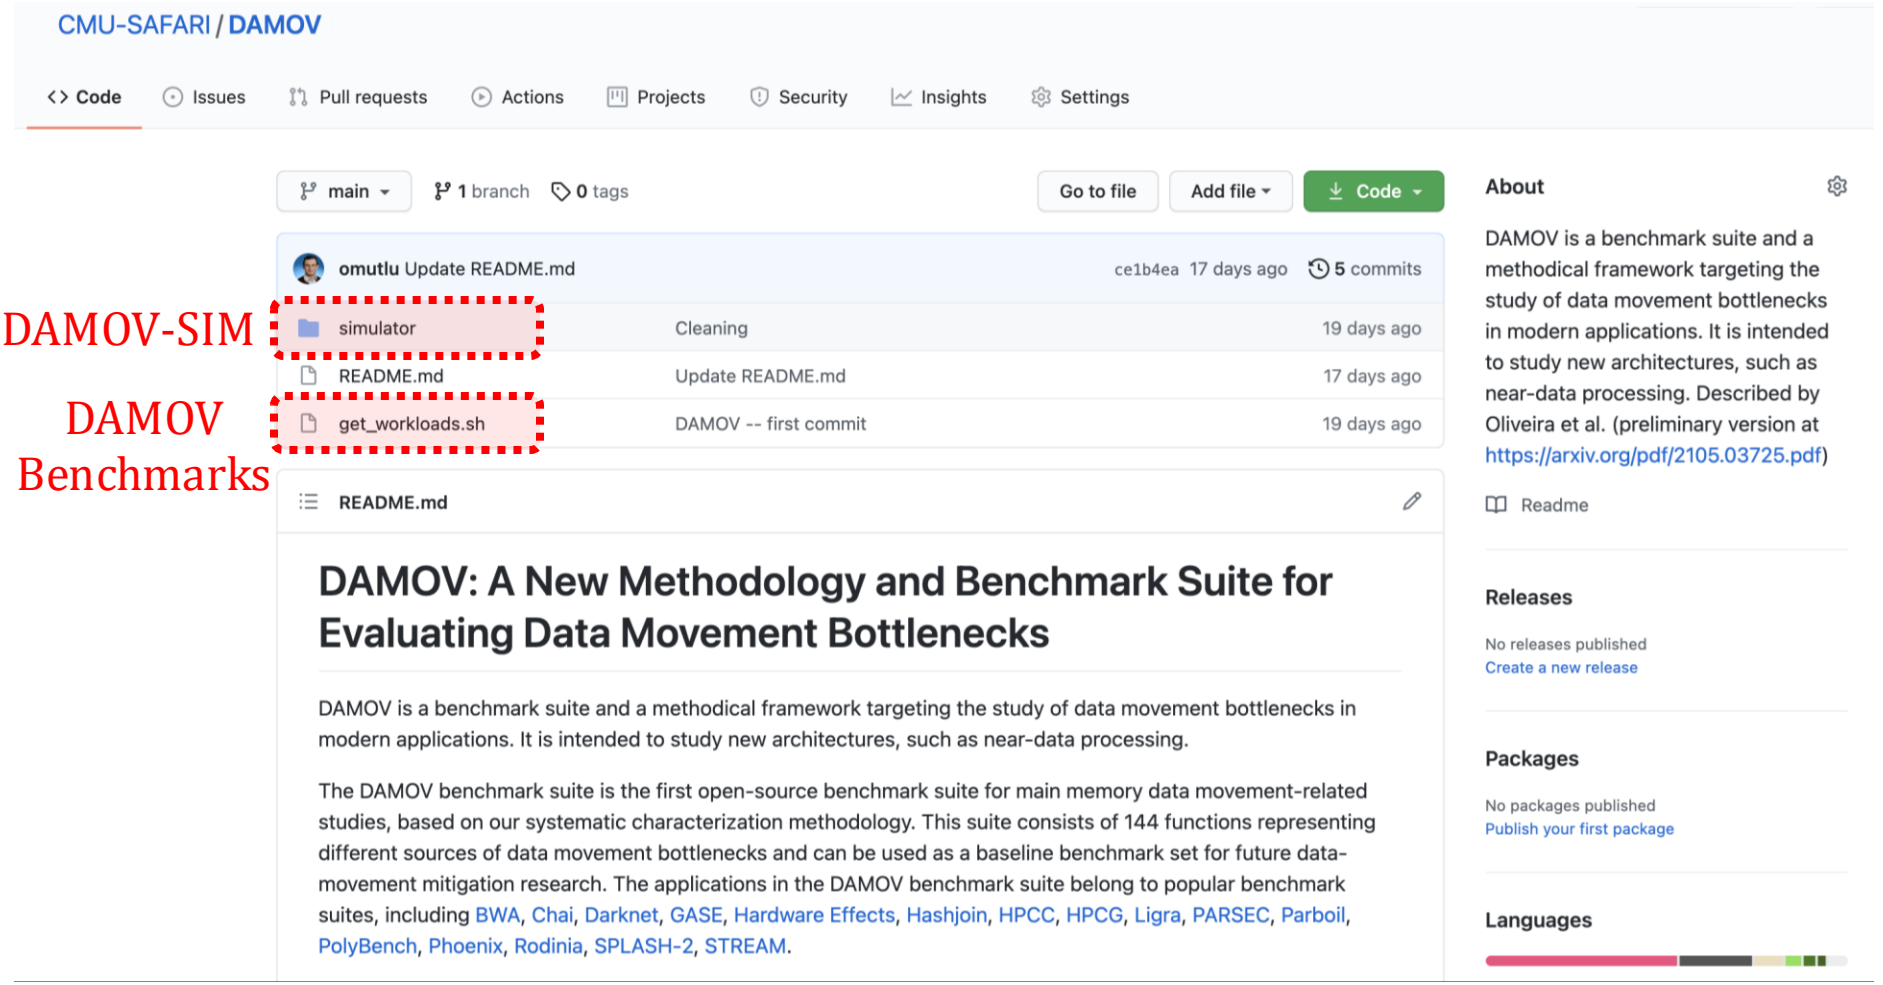

DAMOV is Open Source

- We open-source our benchmark suite and our toolchain

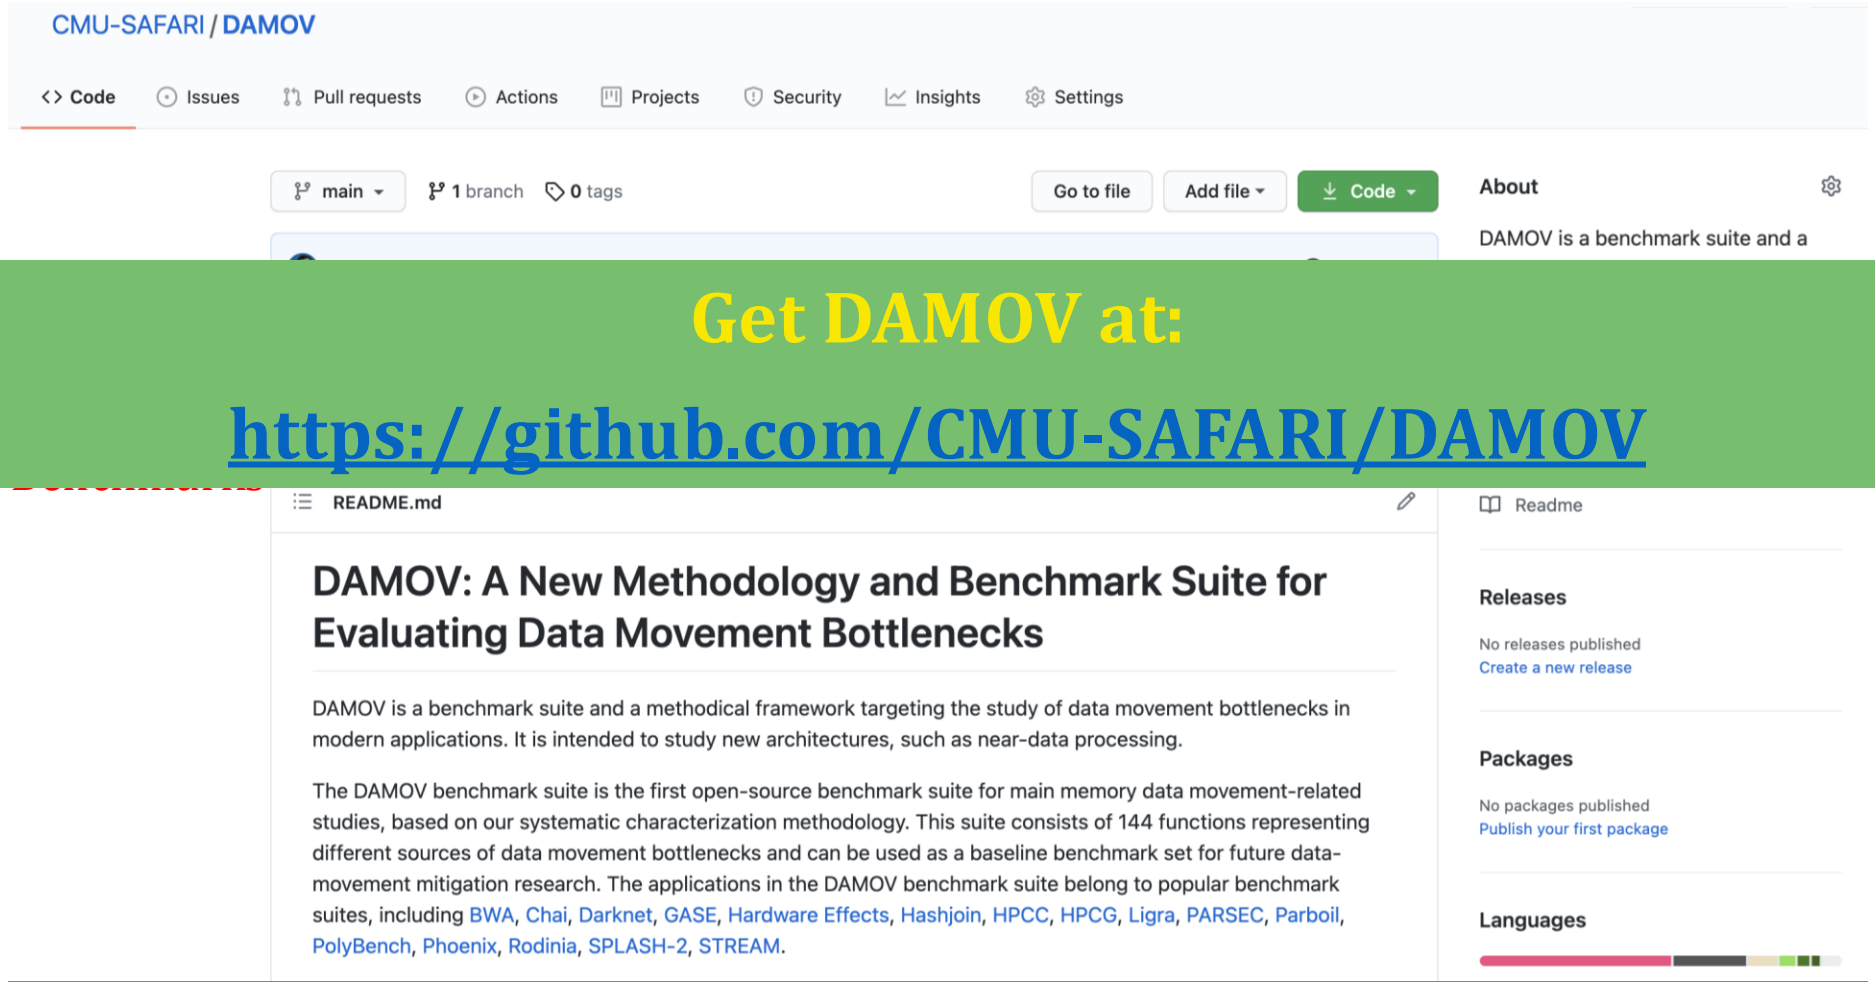

More on DAMOV Analysis Methodology & Workloads

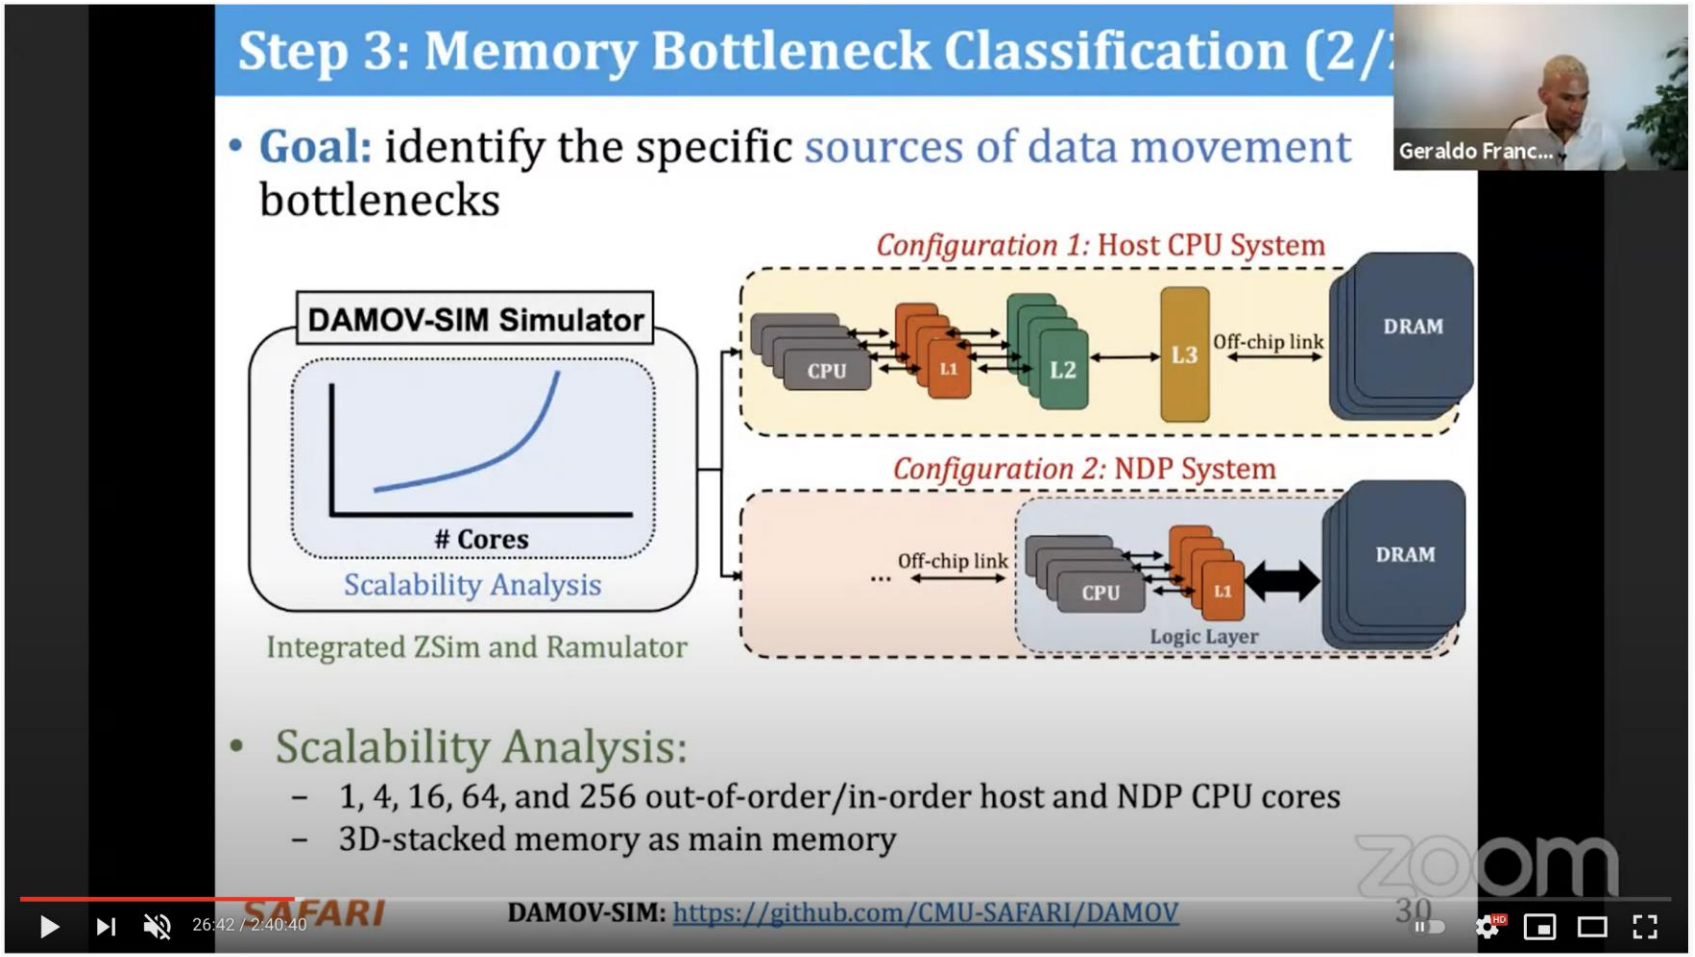

SAFARI Live Seminar: DAMOV: A New Methodology & Benchmark Suite for Data Movement Bottlenecks

352 views • Streamed live on Jul 22, 2021

18 0 SHARE SAVE ...

Onur Mutlu Lectures
17.7K subscribers

ANALYTICS EDIT VIDEO

More on DAMOV Methods & Benchmarks

- Geraldo F. Oliveira, Juan Gomez-Luna, Lois Orosa, Saugata Ghose, Nandita Vijaykumar, Ivan fernandez, Mohammad Sadrosadati, and Onur Mutlu,
  "DAMOV: A New Methodology and Benchmark Suite for Evaluating Data Movement Bottlenecks"
  IEEE Access, 8 September 2021.
  Preprint in ***arXiv***, 8 May 2021.
  [arXiv preprint]
  [IEEE Access version]
  [DAMOV Suite and Simulator Source Code]
  [SAFARI Live Seminar Video (2 hrs 40 mins)]
  [Short Talk Video (21 minutes)]

DAMOV: A New Methodology and Benchmark Suite for Evaluating Data Movement Bottlenecks

GERALDO F. OLIVEIRA, ETH Zürich, Switzerland
JUAN GÓMEZ-LUNA, ETH Zürich, Switzerland
LOIS OROSA, ETH Zürich, Switzerland
SAUGATA GHOSE, University of Illinois at Urbana–Champaign, USA
NANDITA VIJAYKUMAR, University of Toronto, Canada
IVAN FERNANDEZ, University of Malaga, Spain & ETH Zürich, Switzerland
MOHAMMAD SADROSADATI, ETH Zürich, Switzerland
ONUR MUTLU, ETH Zürich, Switzerland

Research Tools PNM: Samsung HBM-PIM

https://github.com/SAITPublic/PIMSimulator

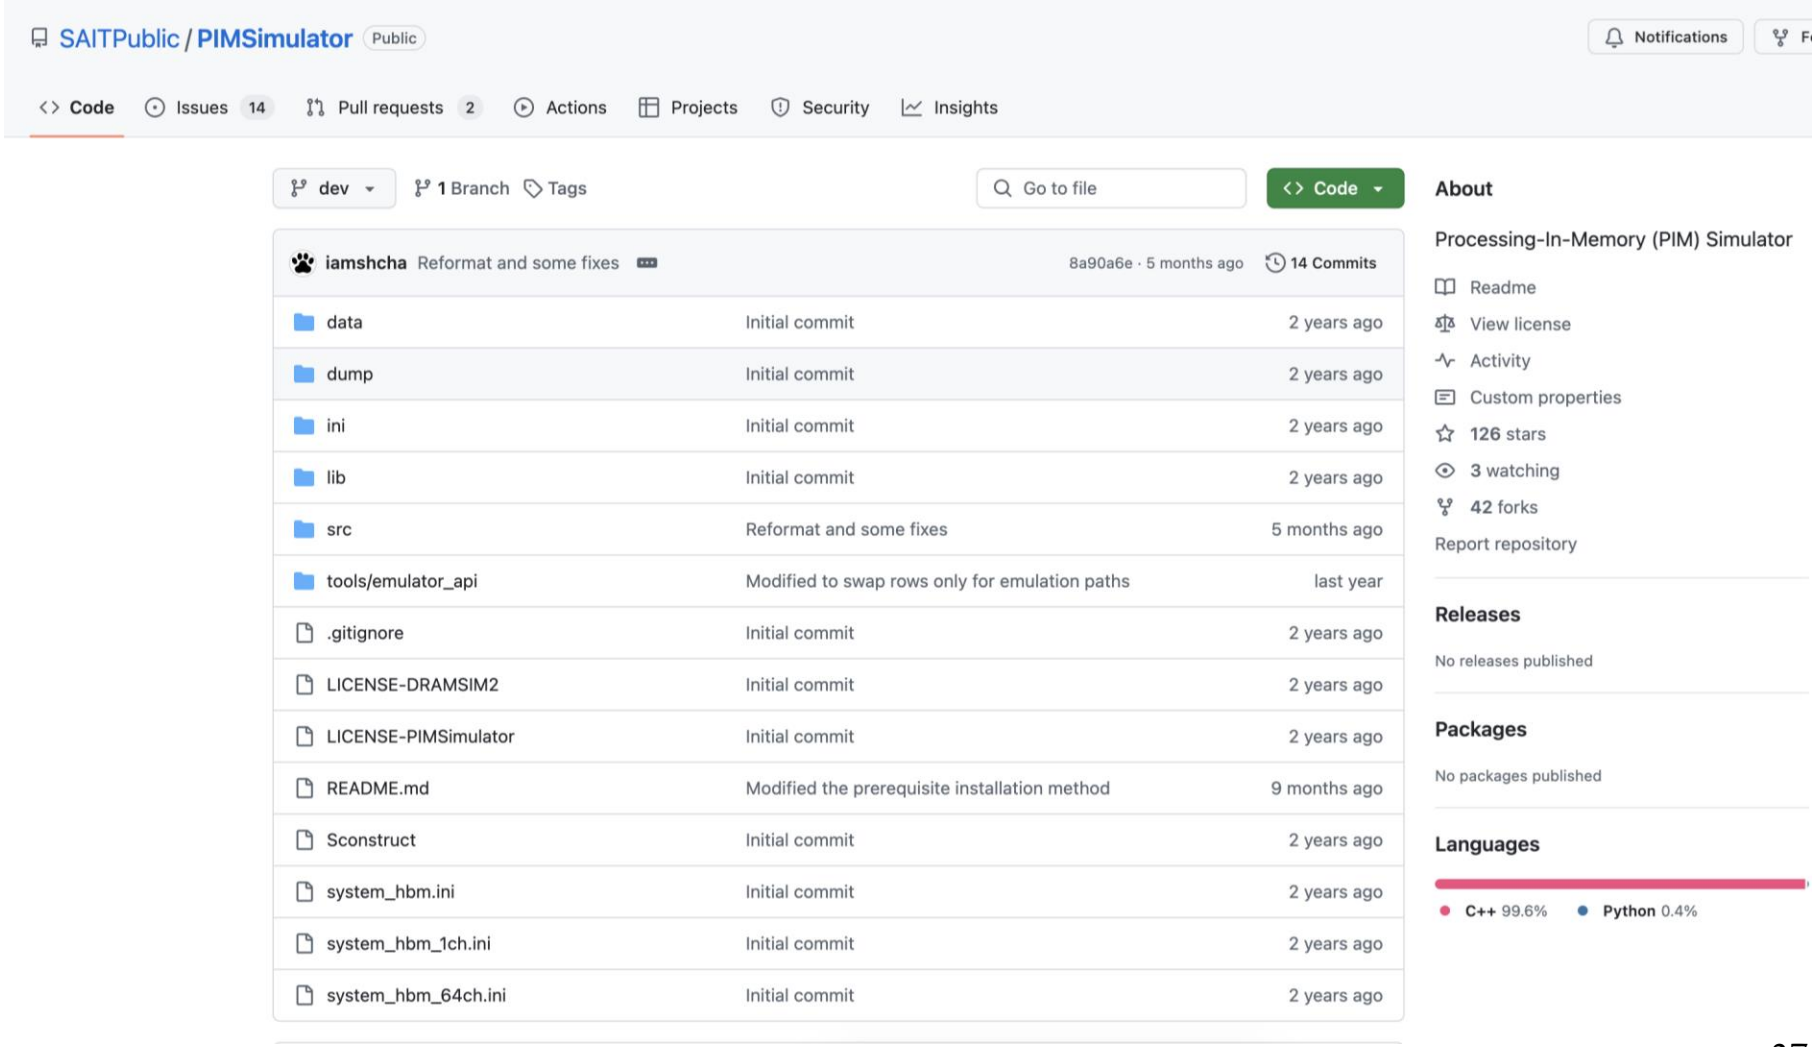

Research Tools PNM: UPMEM PIM (I)

https://github.com/VIA-Research/uPIMulator

README MIT license

📖 Introduction

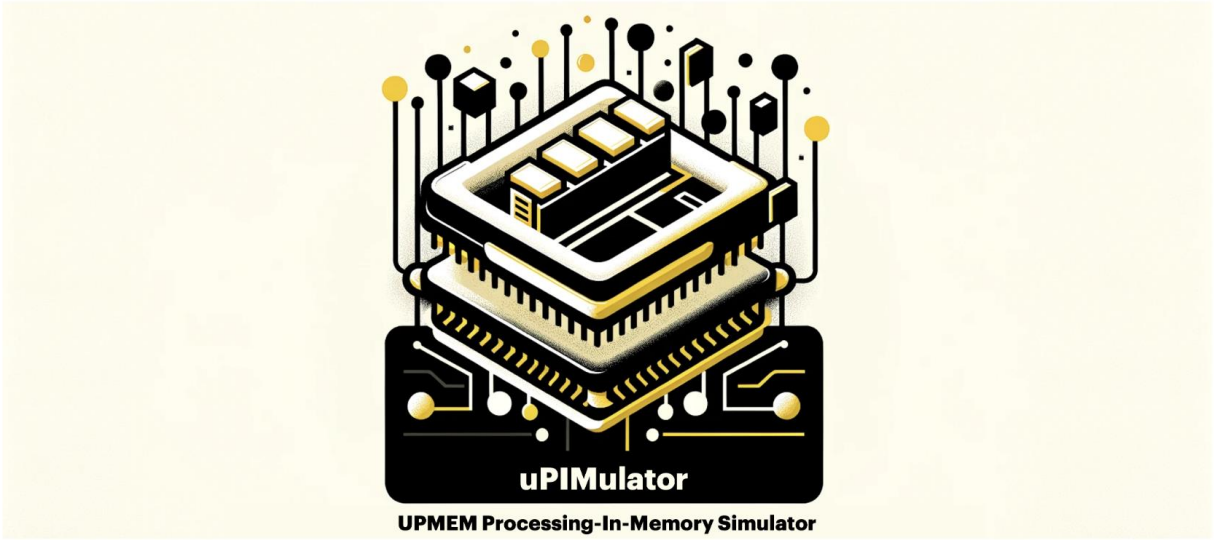

Welcome to the uPIMulator Framework Documentation!

This documentation serves as your comprehensive guide to the uPIMulator framework, catering to both novice and experienced researchers. Here, you'll find the resources necessary to leverage uPIMulator effectively for your research projects.

We provide in-depth coverage of uPIMulator's features, from foundational concepts to advanced functionalities. Explore this documentation to unlock the full potential of uPIMulator and elevate your research endeavors.

Research Tools PNM: UPMEM PIM (II)

https://ieeexplore.ieee.org/document/10476411/

2024 IEEE International Symposium on High-Performance Computer Architecture (HPCA)

Pathfinding Future PIM Architectures by Demystifying a Commercial PIM Technology

Bongjoon Hyun Taehun Kim Dongjae Lee Minsoo Rhu
KAIST
{bongjoon.hyun, taehun.kim, dongjae.lee, mrhu}@kaist.ac.kr

2nd Workshop on Memory-Centric Computing: Processing-Near-Memory - Part I

Dr. Geraldo F. Oliveira

https://geraldofojunior.github.io

ICS 2025

08 June 2025

SAFARI

ETH zürich