

Bachelor's / Master's / Semester Project

Improving UPMEM, a Real-world Computation-in-Memory Architecture

Context & background:

Data movement among computing and memory units of our current systems has already become a significant performance and energy bottleneck for today's applications. An example of such applications and workloads are those for machine learning, graph processing, and computational biology.

To alleviate this problem, computation in memory (CIM)¹ is a reignited computation paradigm that advocates a shift from traditional processor-centric systems to a more data-centric one where we place some compute units closer to or inside the memory, where our data resides.

UPMEM is the first real-world CIM-enabled architecture that is currently publicly available. The UPMEM CIM-enabled architecture enhances the existing DRAM chips with some DRAM Processing Units (DPUs), enabling an unprecedented memory bandwidth for workloads that need it.

In this direction, we aim to program and optimize various workloads on the UPMEM architecture and/or propose and implement architecture improvements for the future generation of UPMEM like CIM-enabled architectures.

Expected effort:

This project can take shape as a semester, BSc, or MSc project.

In this project, the student will

- investigate different applications and workloads suited for UPMEM and build an analytical model to show the maximum potential benefits
- program the workload on real-world UPMEM architecture and compare the results with the same workload on traditional systems (e.g., CPUs and GPUs)
- propose novel ideas (in hardware or software) to improve the workload for the new platform

For preliminary readings to get familiar with UPMEM and CIM studies please see:

- UPMEM SDK documentation (<https://sdk.upmem.com/2020.4.0/>)
- Gómez-Luna, et al., "[Benchmarking a New Paradigm: Experimental Analysis and Characterization of a Real Processing-in-Memory System.](#)" IEEE Access, 2022.
- Ghose, et al., "[Processing-in-memory: A workload-driven perspective.](#)" IBM Journal of Research and Development, 2019.
- Hamdioui, et al., "[Memristor based computation-in-memory architecture for data-intensive applications.](#)" DATE, 2015.

¹ Interchangeably known as Processing-in-Memory (PIM)

Requirements:

- Outstanding programming skills (C/C++)
- Good knowledge of design and analysis of algorithms and data structures
- An interest in developing and evaluating new ideas, discovering why things do or do not work, making systems efficient and usable
- Strong work ethic

Plus Points:

- Background in CIM
- Background in computer architecture

Contact information:

If you are interested, please email: Said Hamdioui via S.Hamdioui@tudelft.nl, Stephan Wong via J.S.S.M.Wong@tudelft.nl, and Michael Shahroodi via tahamichaelshahroodi@gmail.com.