

Bachelor's / Master's / Semester Project

Programming and Improving a Real-world Processing-in-Memory Architecture

Data movement between the memory units and the compute units of current computing systems is a major performance and energy bottleneck. For example, data movement between the main memory and the processing cores accounts for 62% of the total system energy in consumer applications. As a result, the **data movement bottleneck** is a huge burden that greatly limits the energy efficiency and performance of modern computing systems.

Many modern and important workloads such as **machine learning, computational biology, and graph processing** suffer greatly from the data movement bottleneck. In order to alleviate this data movement bottleneck, **Processing-in-Memory (PIM)** represents a **paradigm shift** from the traditional processor-centric design, where all computation takes place in the compute units, to a more data-centric design where processing elements are placed closer to or inside where the data resides.

After many years of research proposals from Industry and Academia, **a real-world processing-in-memory architecture is publicly available**. The UPMEM PIM architecture integrates DRAM Processing Units (DPUs) inside DRAM chips. As a result, workloads can take advantage of an **unprecedented memory bandwidth**.

Projects in this line of research span **software and hardware as well as the software/hardware interface**. We are looking for enthusiastic students who want to work hands-on (1) **programming and optimizing workloads** on the UPMEM PIM architecture, and/or (2) proposing and implementing **hardware and architecture improvements** for future PIM architectures.

Requirements

- Outstanding programming skills (C/C++)
- Computer architecture background
- Interest in future computer architectures and computing paradigms
- Interest in discovering why things do or do not work and solving problems
- Interest in making systems efficient and usable
- Strong work ethic

Preliminary readings to get familiar with PIM and the UPMEM architecture:

- Introduction to PIM (summary papers):
 - https://people.inf.ethz.ch/omutlu/pub/ProcessingDataWhereItMakesSense_micpro19-invited.pdf
 - https://people.inf.ethz.ch/omutlu/pub/processing-in-memory_workload-driven-perspective_IBMjrd19.pdf
- UPMEM SDK documentation (<https://sdk.upmem.com/2020.4.0/>)

If you are interested, please email **Professor Onur Mutlu** and **Dr. Juan Gómez-Luna**:

omutlu@gmail.com

el1goluj@gmail.com

<https://safari.ethz.ch>

<https://people.inf.ethz.ch/omutlu/>