

# Processing-in-Memory Primitives for Data Management (PIMPMe)

Kai-Uwe Sattler (DBIS) Daniel Ziener, Andreas Becher (RAES)



## Primary Research Question:

**How can we exploit computational memory effectively and efficiently in distributed database systems and which are the best abstractions?**

Most recent publication of **phase 1**:

Accelerating Aggregation Using a Real Processing-in-Memory System

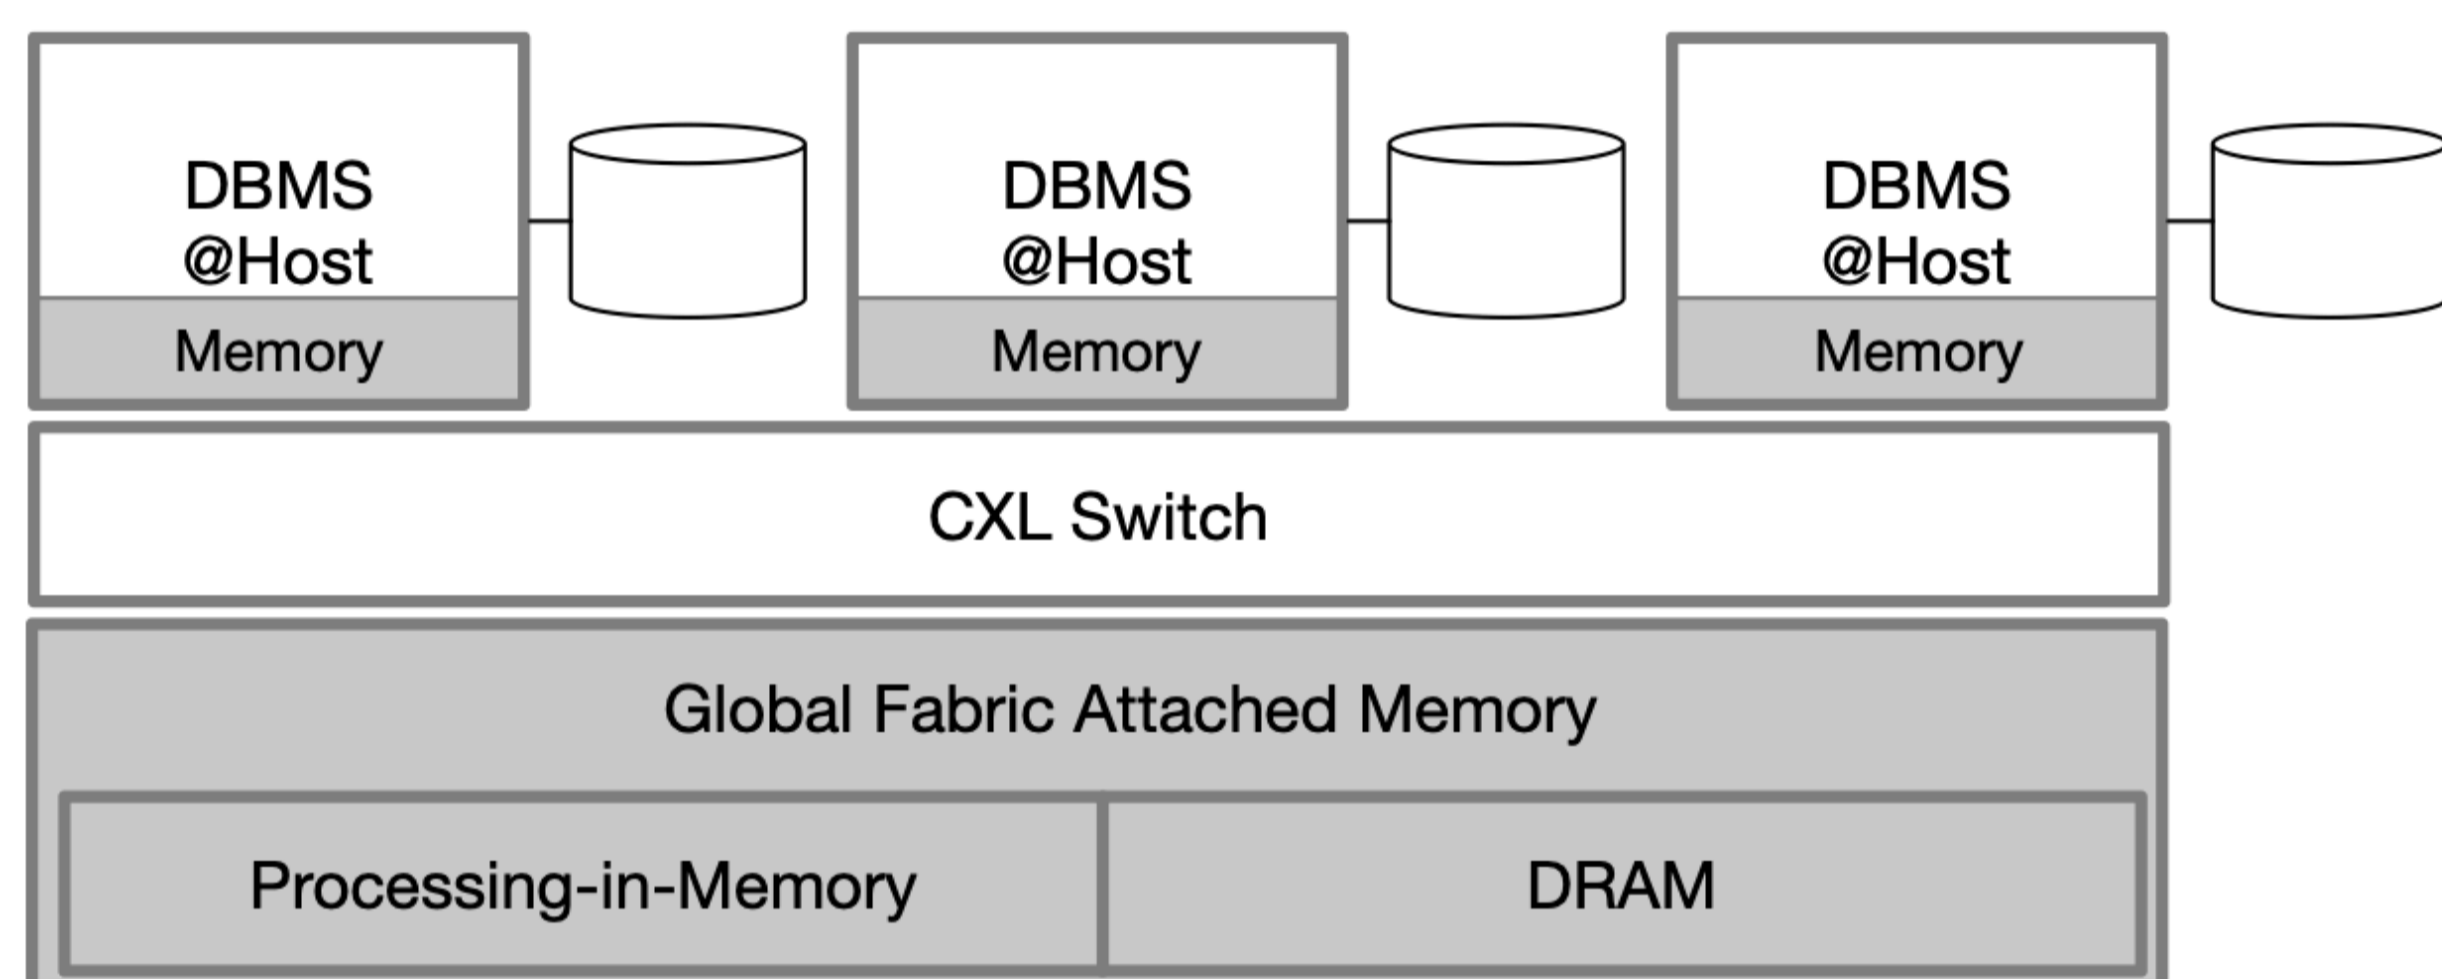
(Muhammad Attahir Jibril, Hani Al-Sayeh, and Kai-Uwe Sattler; DOI: 10.1109/ICDE60146.2024.00300)

## Phase 2

**Overall goal of the second phase:**

Extend on results of first phase by focusing on the following areas

- PIM abstractions beyond UPMEM
- Distributed computing using shared computational memory



## Objective A

**Exploring and revising candidates for offloading distributed database operations:**

Computational Memory Platform

- PIM paradigm with UPMEM and FPGA
- Operators for offloading to computational memory
- Identification of critical distributed DBMS operations/operators

DBIS




## Objective C

**Developing concepts to exploit these operations in distributed database operations:**

Developing distributed query operations exploiting shared computational memory

- Query operators for shared computational memory
- Operators for distributed queries
- Memory-centric transaction processing
- Code generation and query compilation

DBIS



## Objective B

**Defining abstractions for distributed database operations in computational memory:**

Computational memory primitives for distributed data processing

- Unified memory allocation and access
- Unified compute model
- FPGA-based evaluation platform for CXL architecture

