Near-Data Processing in DBMS on Native Computational Storage under HTAP Workloads









Tobias Vinçon, Christian Knödler, Arthur Bernhardt, Ilia Petrov
Data Management Lab
Reutlingen University, Germany

Leonardo S.-Vasquez, Sajjad Tamimi, Lukas Weber, Florian Stock, Andreas Koch
Embedded Systems and Applications Group
TU Darmstadt, Germany

Motivation

Control Flow →
Data Flow →

TECHNISCHE UNIVERSITÄT DARMSTADT | Hochschule Reutlingen Reutlingen University

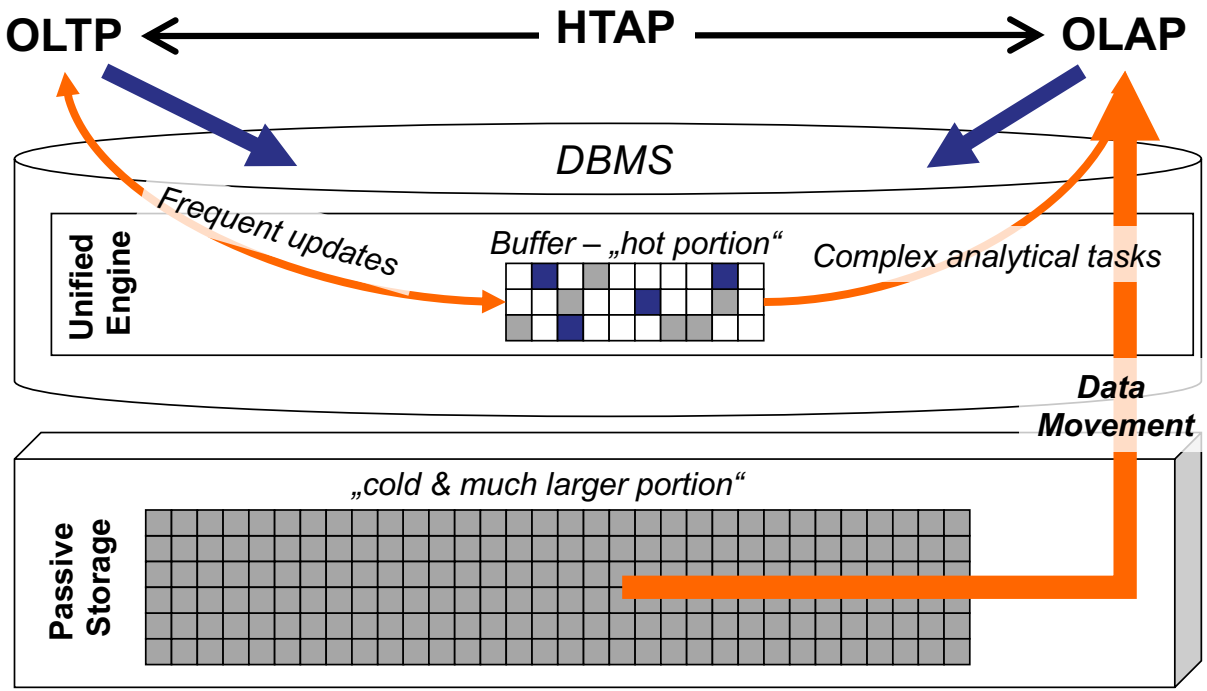

Performance Expectation

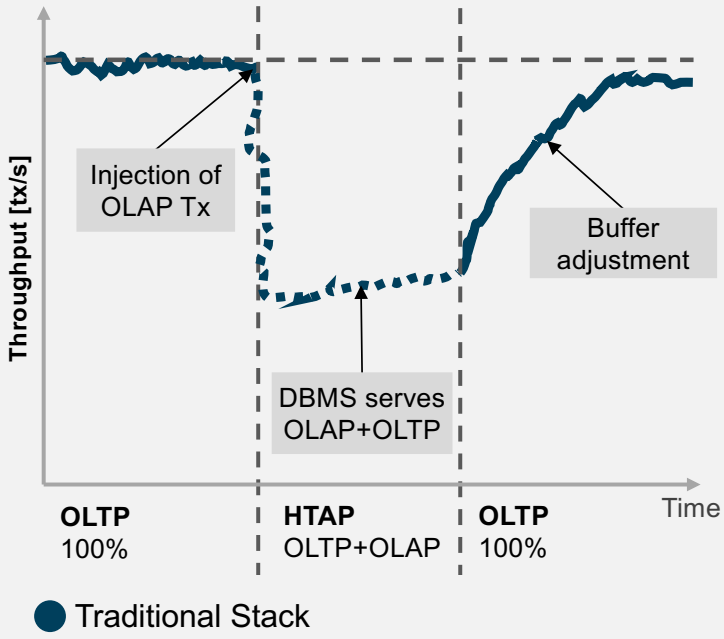

Traditional data-to-code system architectures entail non-robust performance, scalability issues & poor resource efficiency.

Which alternatives are possible?

Motivation

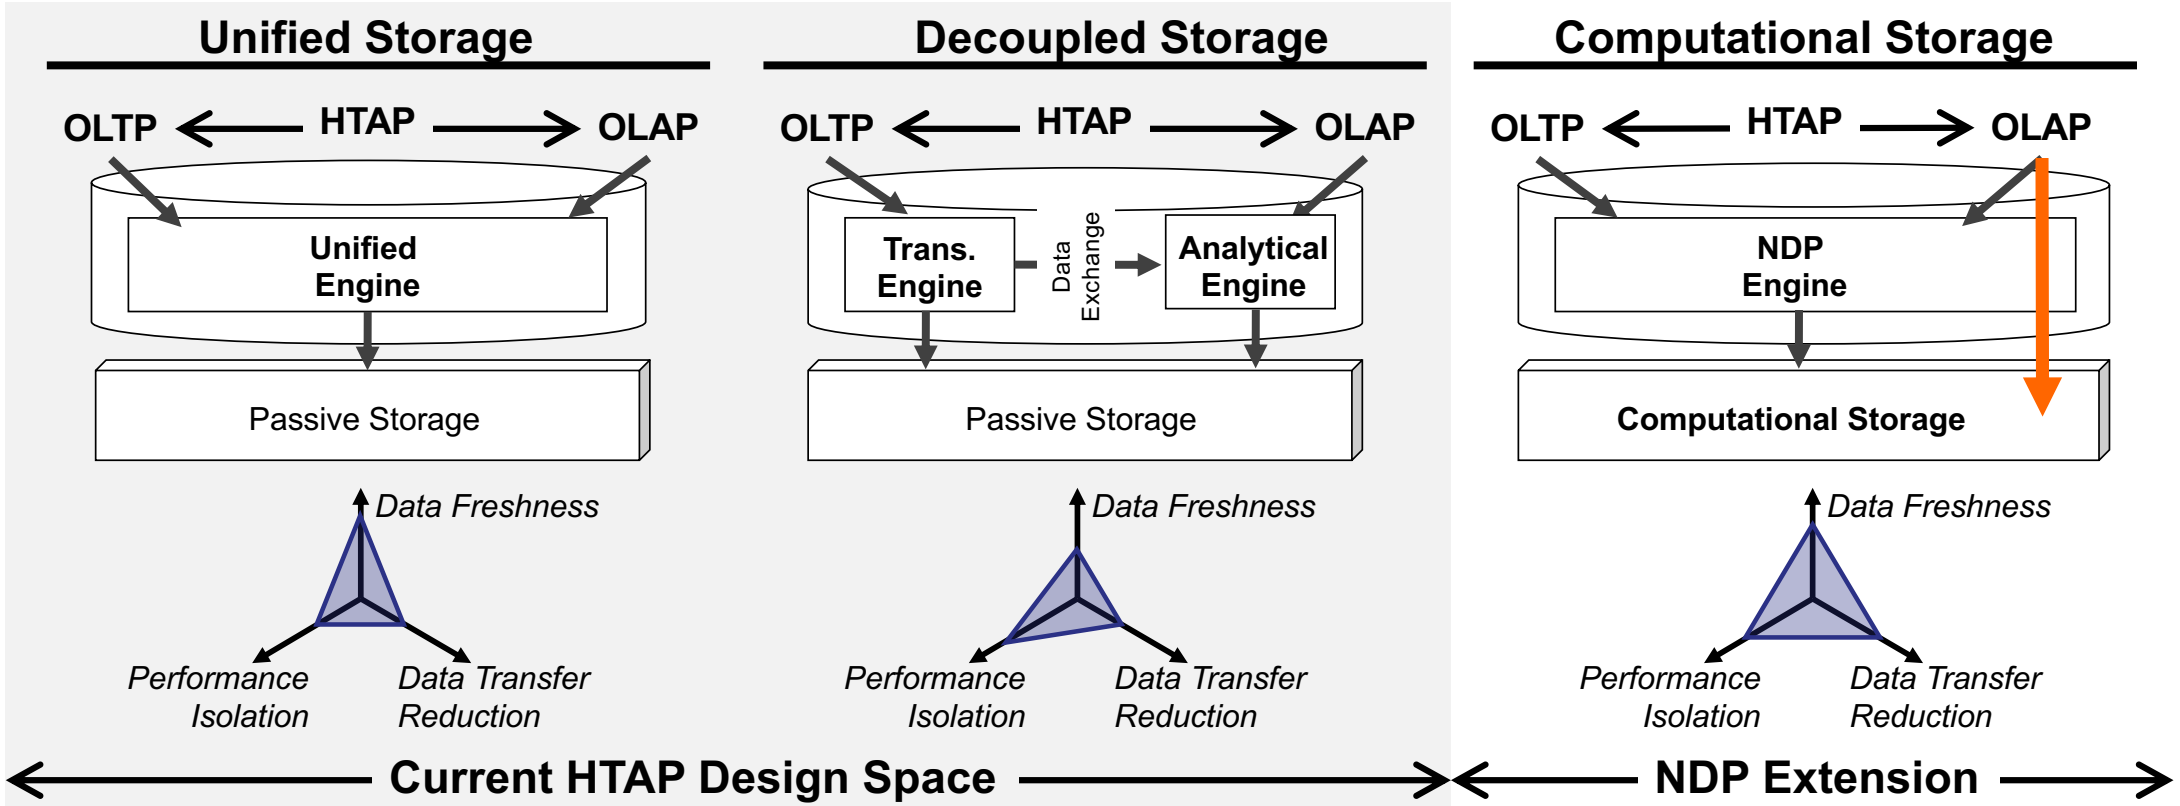

Current approaches in the HTAP design space suffer in either performance isolation or data freshness

Our Proposal: Extend the HTAP design space with NDP to optimize on all dimensions?

Motivation

Control Flow
Data Flow

TECHNISCHE UNIVERSITÄT DARMSTADT
Hochschule Reutlingen
Reutlingen University

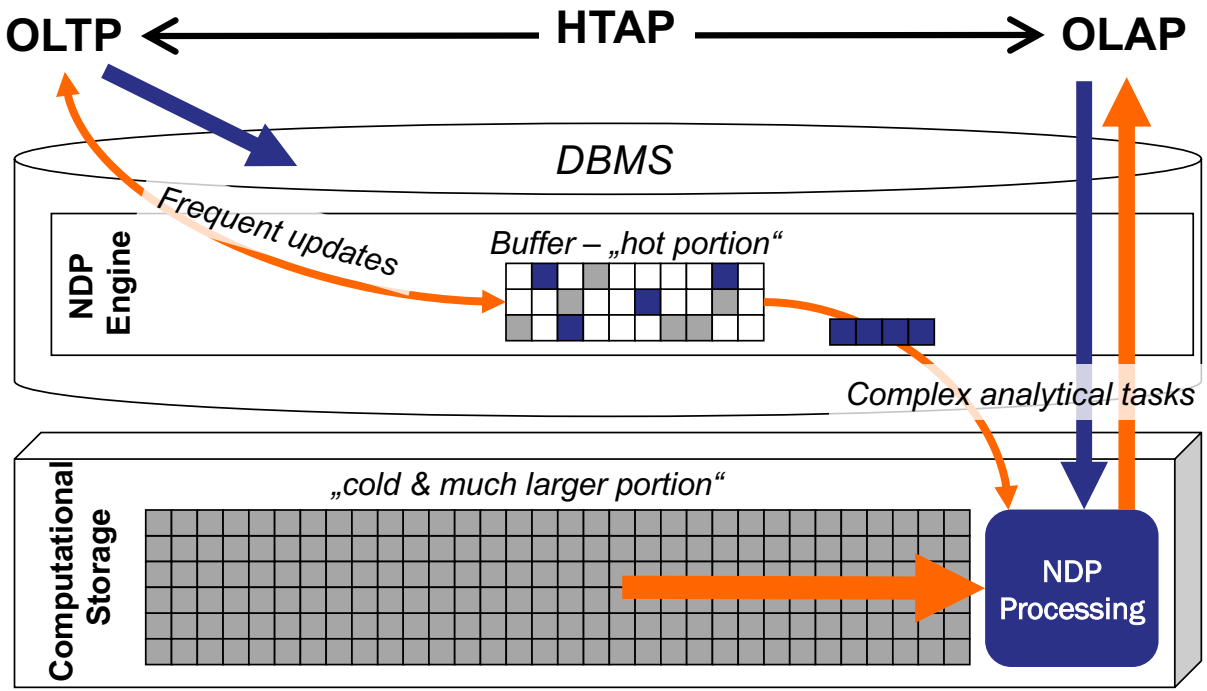

Performance Expectation

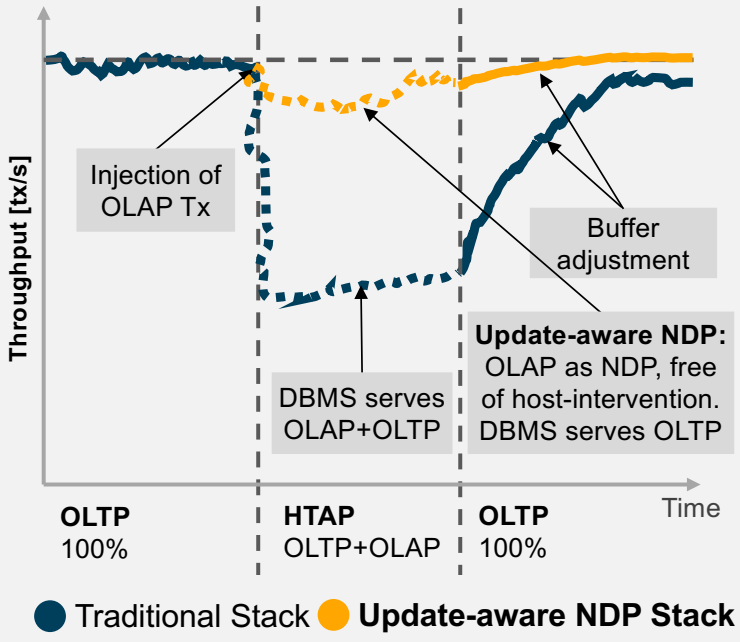

With Update-aware NDP a higher throughput during HTAP is expected as OLAP queries are executed in an transactinally consistent manner on-device while OLTP is served by the DBMS.

How can we achieve intervention-free NDP executions with transactional guarantees?

ES Embedded Systems & Applications
DBLab Data Management Lab

Shared State & Transaction Managment

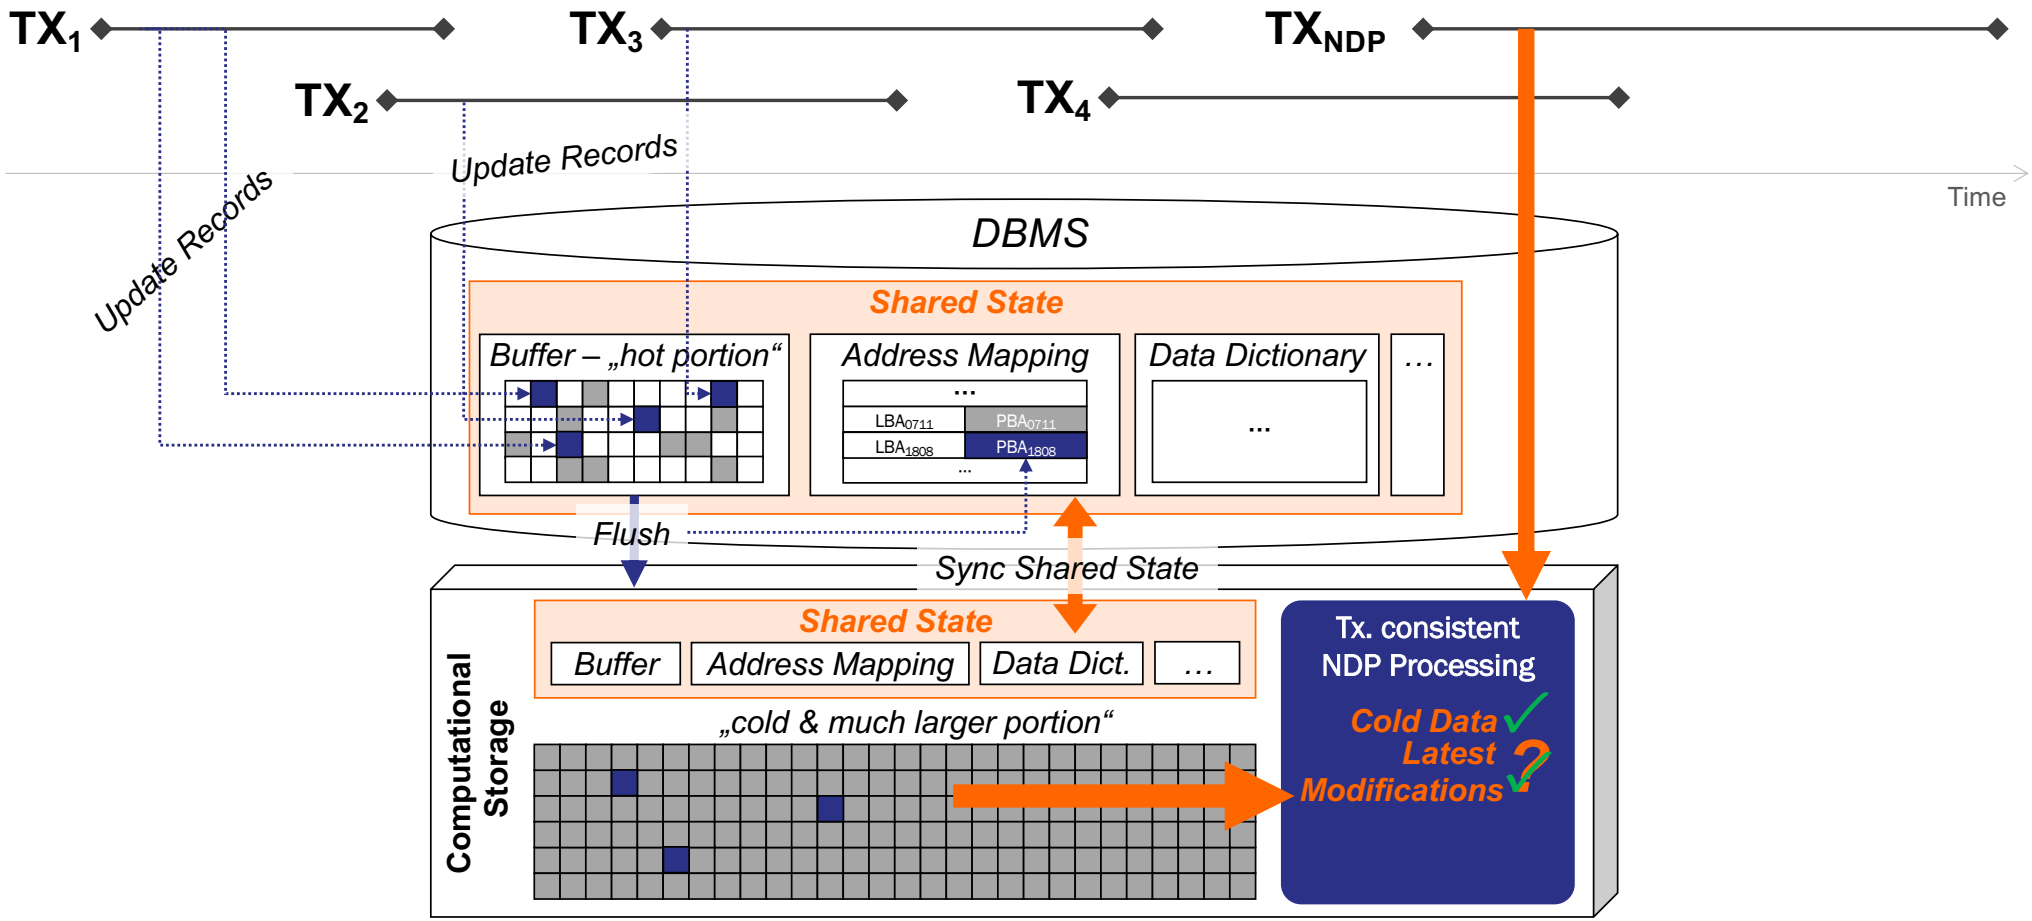

Experimental Setups

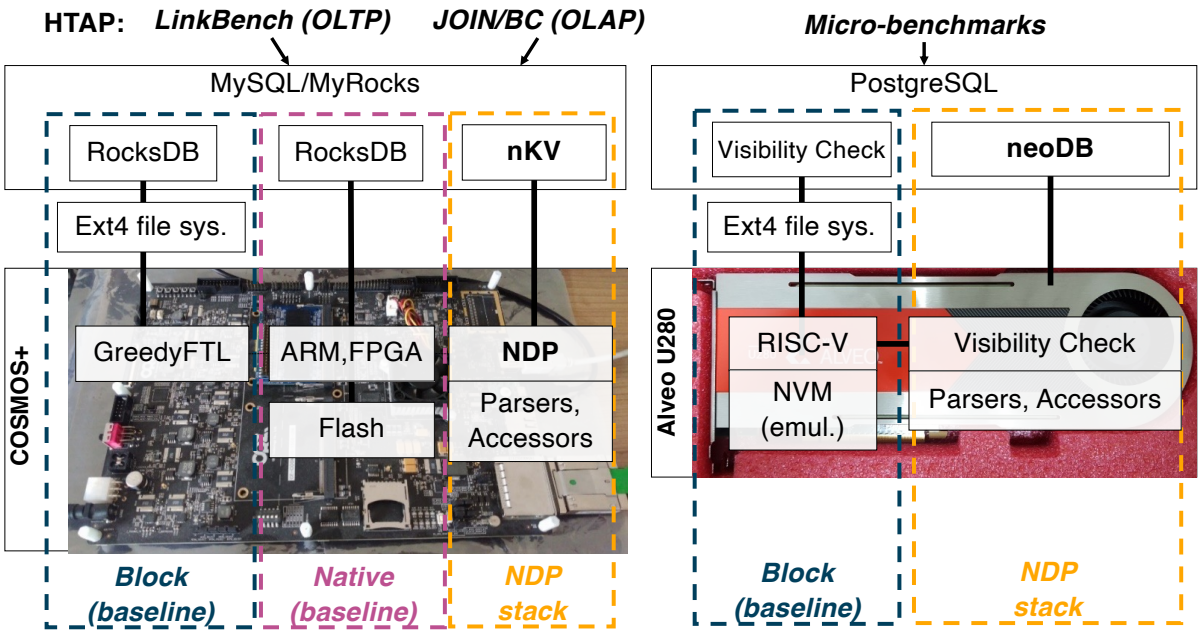

- **Update-aware NDP architecture is integrated in 2 Systems**
  - **nKV**: Snapshot-based MyRocks with nKV as Storage Manager
  - **neoDBMS**: MVCC-based PostgreSQL
- **Experiments are conducted on real hardware**
  - COSMOS+ emulates consumer-grade Computational Storage Device
  - Alveo U280 emulates enterprise-grade Computational Storage Device
- **Extended LinkBench for HTAP workload**
  - OLAP Queries (e.g. Betweenness Centralitly or Join/Group By) are executed in parallel to the classical OLTP workload
  - Dataset comprises 10M Nodes ≙ 20 GB data
  - Several control parameters: $OLTP_{SKEW}$, $OLAP_{SEL}$, $OLAP_{PAUSE}$

Exp #1 – Overall performance:

Update-aware NDP enables transactionally consistent NDP executions of OLAP operations in presence of OLTP updates in HTAP systems, without performance drops.

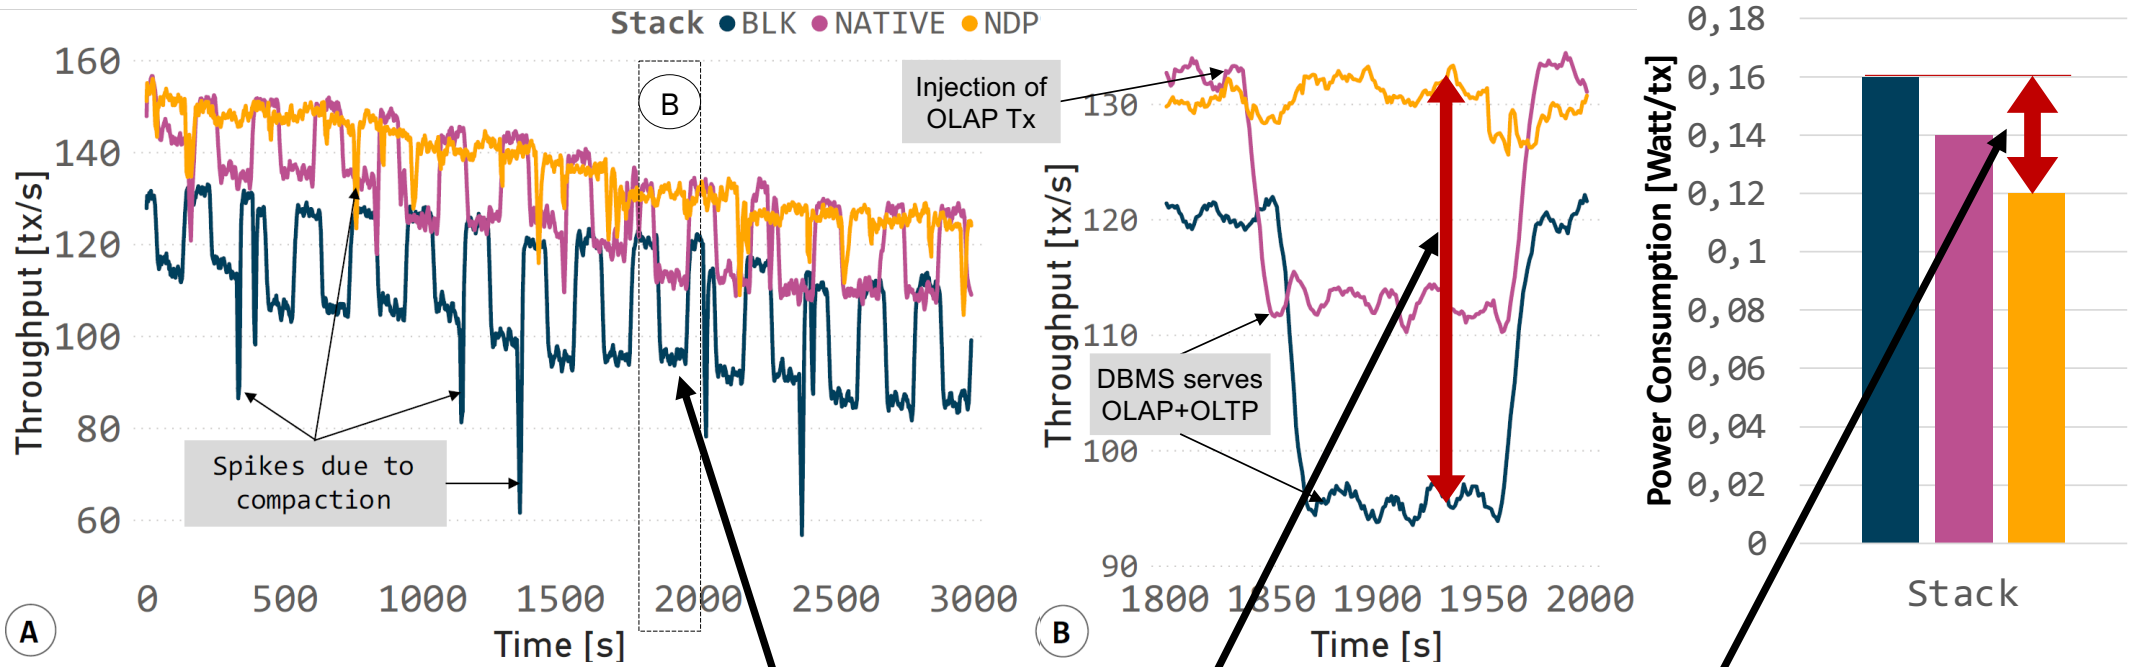

Intervention-free, update-aware NDP improves average throughput by 30% with lower power consumption.

Exp #2 – Robustness:

Update-aware NDP is intervention-free, yielding robust performance.

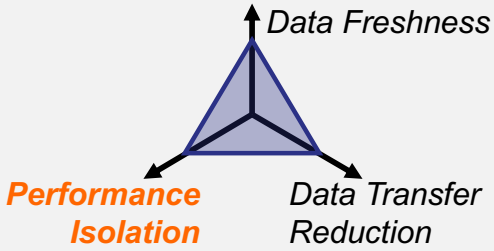

Drops in OLTP are entailed by Buffer Pollution that correlates with the Buffer Size

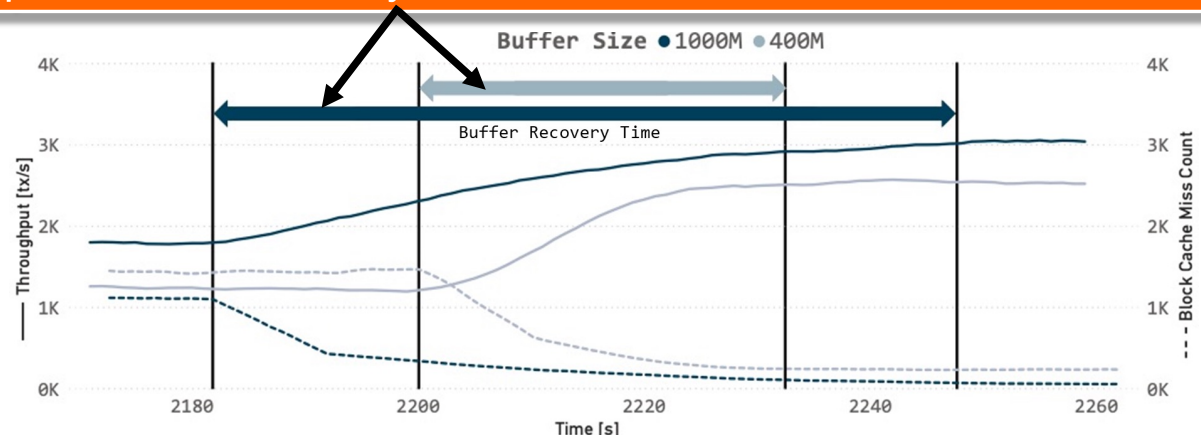

Update-NDP can continuously leverage the CPU more efficiently

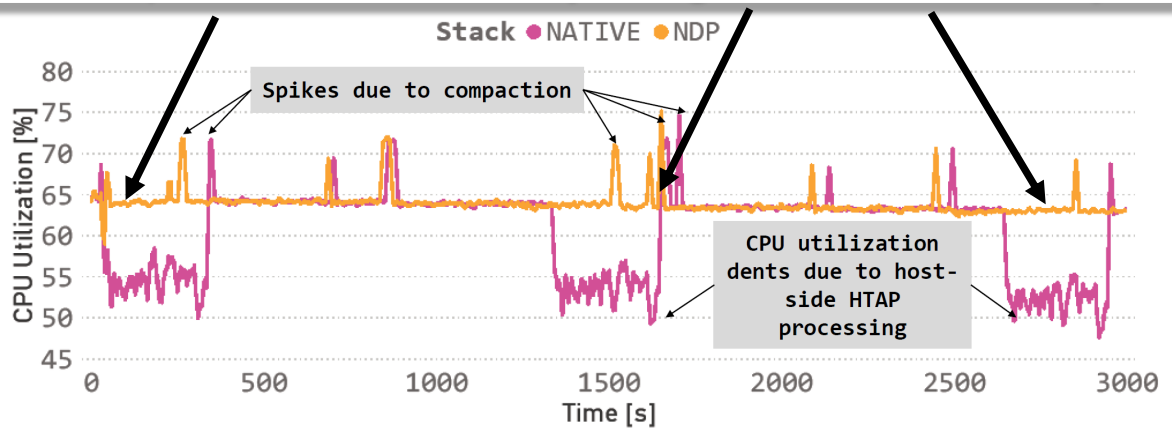

Exp #4 – Data Transfer Reduction:

Update-aware NDP reduces data transfers.

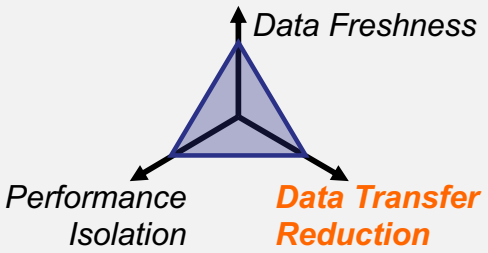

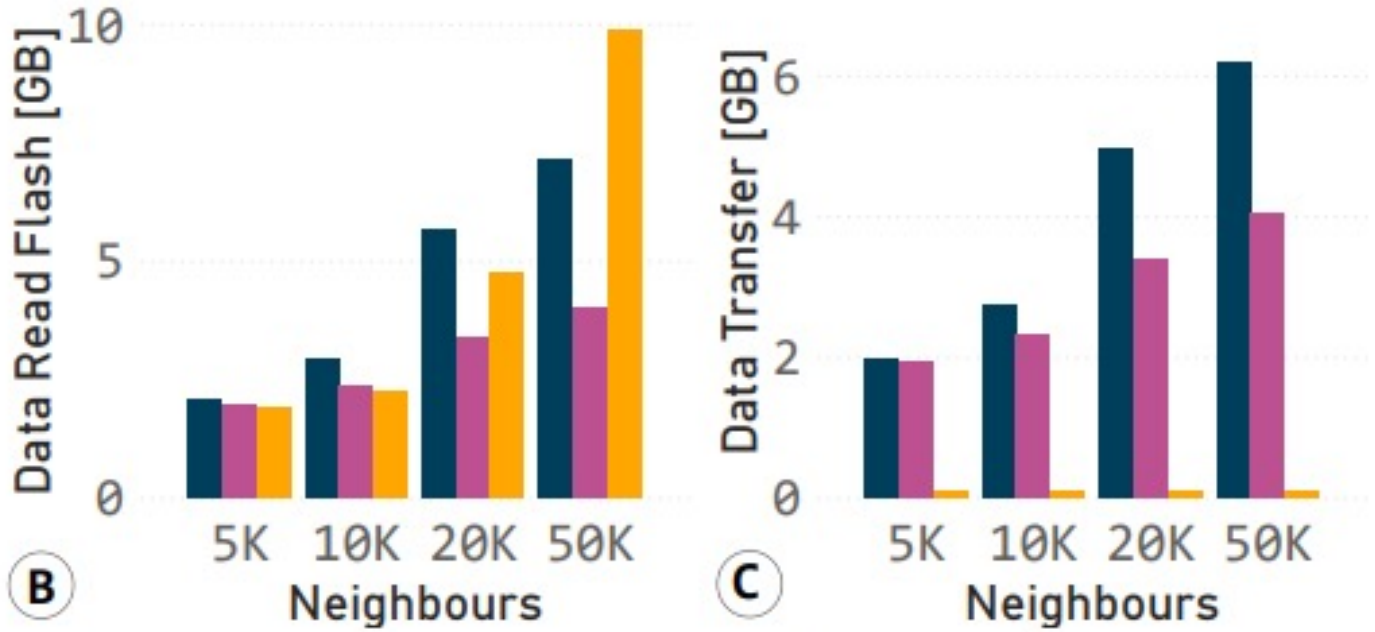

Even though Update-aware NDP requires to read more on-device (B), it reduces data transfers to the host significantly (C) and thus, improves the execution duration (A)

Exp #5 – Fresh data with low overhead:

Update-aware NDP can operate on fresh data with low overhead.

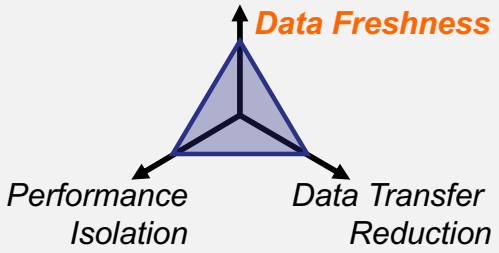

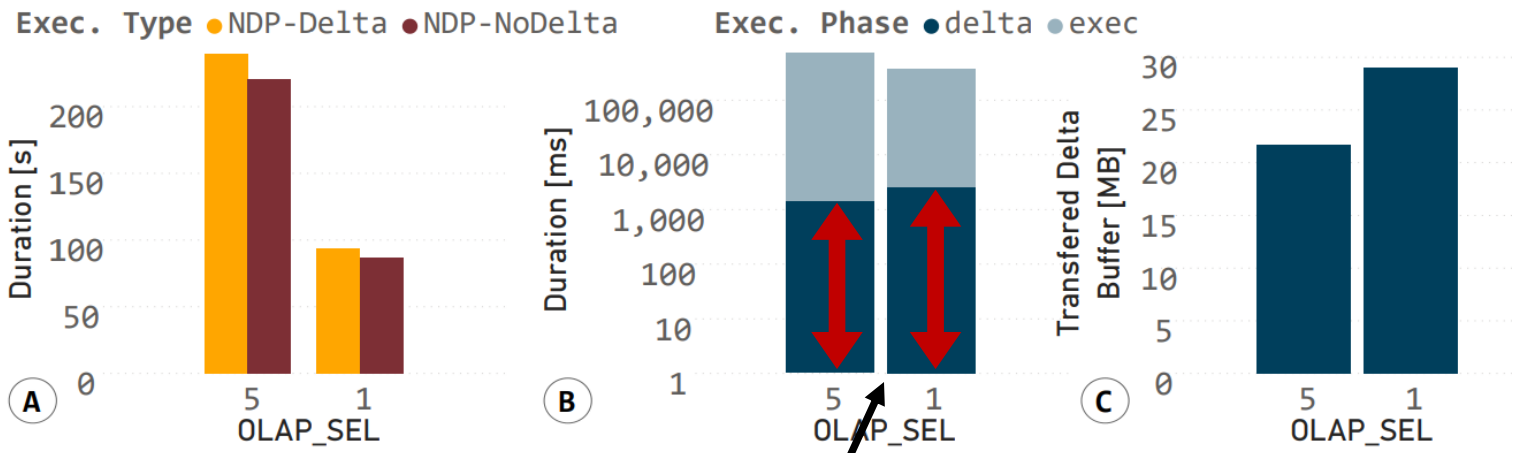

The expenses for transferring the Shared State to achieve transactional consistency amount to negligible 0.7% of the entire execution time

Conclusion

Near-Data Processing in Database Systems on Native Computational Storage under HTAP Workloads

TECHNISCHE UNIVERSITÄT DARMSTADT | Hochschule Reutlingen Reutlingen University

Contributions:

- **Introduction of generic update-aware NDP architecture**
  - Definition of a Shared State as delta between host and device
  - Proposal of intervention-free execution to avoid device-host roundtrips
  - Enablement of transactionally consistent NDP operations
- **Integration in 2 Systems**
  - Integration of concepts in Snapshot-based and MVCC system
  - Evaluation on real hardware
- **Evaluation of Performance Improvements**
  - Reduces data transfers between host and device
  - Improves robustness and performance (+30%)
  - Reduces overall power consumption (-26%)
  - Shared State is transfered with marginal cost (+0.7%)

More in the paper:

doi:10.14778/3547305.3547307

- **Concepts about**
  - NDP Interface
  - Parsers & Accessors incl. Physical Page Pointer
  - NDP Pipelines and Operations
  - Result-Set Handling
- **Experiments on**
  - Different types of OLAP Operations
  - With different memory setups
  - NDP Visibility Checks
  - Power Consumption

https://dblab.reutlingen-university.de
https://www.esa.informatik.tu-darmstadt.de

TECHNISCHE UNIVERSITÄT DARMSTADT

Hochschule Reutlingen
Reutlingen University

Embedded Systems & Applications

DBLab Data Management Lab