

# Fast and Efficient Update Handling for Graph H<sup>2</sup>TAP

Muhammad Attahir Jibril ▪ Hani Al-Sayeh ▪ Alexander Baumstark ▪ Kai-Uwe Sattler

**(EDBT 2023)**

Sixteenth LDBC TUC Meeting

24.06.2023

# Introduction

- **Graph workloads:**
  - Transactions e.g. Inserts, deletes, lookups
  - Analytics e.g. BFS, SSSP, PR
- **Graph HTAP:** concurrent mix of transactional and analytical workloads

# Design Space

- Single Instance CPU
  - Pros: easier to manage
  - cons: no workload-specific optimizations, no accelerated analytics
- Single Instance GPU
  - more complex and problematic than single instance CPU
- Separate Instances
  - pros: workload-specific optimizations, leverage GPU
  - cons: synchronizing the instances

# Heterogeneous HTAP (H<sup>2</sup>TAP)

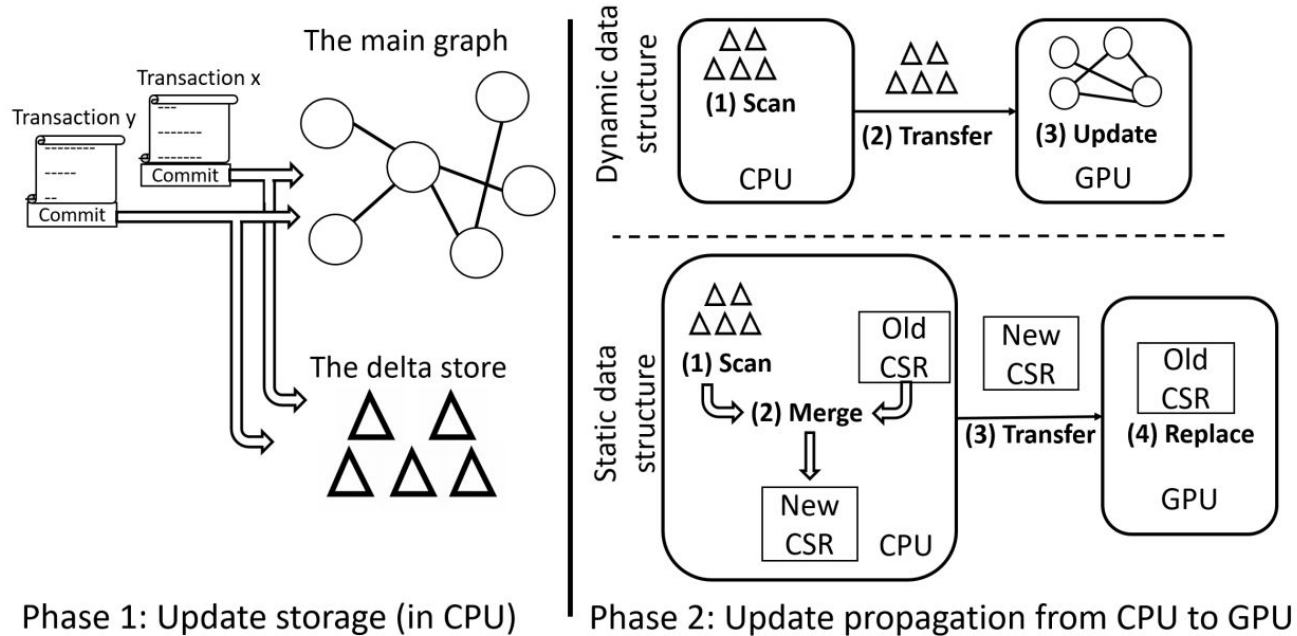
**Leveraging each processor  
for the most suitable workload**

# Challenges

1. Propagating updates
2. Meeting freshness requirements
3. Guaranteeing consistency

## Update Handling

# Update Handling: Overview



# Evaluation


## Setup

- **Machines:**
  - GPU server
  - PMem server
- **Data:**
  - LDBC SNB dataset
  - Graph 500 dataset
- **Workload:**
  - LDBC SNB-like insert and delete queries
  - LDBC Graphalytics


# Evaluation

## Transaction Time


$\blacktriangle$  DELTA<sup>I</sup><sub>LoDeg</sub>     $\square$  DELTA<sup>I</sup><sub>HiDeg</sub>     $\circ$  DELTA<sup>FE</sup><sub>LoDeg</sub>     $\times$  DELTA<sup>FE</sup><sub>HiDeg</sub>




(a) INSERT NODE




(b) DELETE NODE



(c) INSERT RELATIONSHIP



(d) DELETE RELATIONSHIP




(e) MIXED




# Evaluation

## Update Propagation Time


 DELTA<sup>I</sup><sub>LoDeg</sub>    DELTA<sup>I</sup><sub>HiDeg</sub>    DELTA<sup>FE</sup><sub>LoDeg</sub>    DELTA<sup>FE</sup><sub>HiDeg</sub>




(a) INSERT NODE




(b) DELETE NODE



(c) INSERT RELATIONSHIP




(d) DELETE RELATIONSHIP



(e) MIXED


# Evaluation

## Update Propagation Time Components



(a) TOTAL TIME

(b) SCAN AND MERGE TIMES



(c) DELTA MERGE

# Evaluation


## H<sup>2</sup>TAP and Single Instance HTAP (CPU)

		BFS	PR	SSSP
Sortledton	Analytics on CPU	<b>1.48</b>	<b>21.34</b>	<b>57.30</b>
	Update Propagation	5.38	5.38	5.38
DELTA <sup>FE</sup>	Analytics on GPU	0.07	0.30	0.13
	Sum	<b>5.45</b>	<b>5.68</b>	<b>5.51</b>


# Evaluation

## Direct Relational Approach

 R     DELTA<sup>FE</sup>



(a) Transactional Update Time



(b) Delta Store Scan

# Conclusion

- Fast and efficient approach for handling updates in a graph  
H<sup>2</sup>TAP setting
  - minimal overhead
  - freshness
  - consistency
- Outperforms existing approaches

# Thank You