

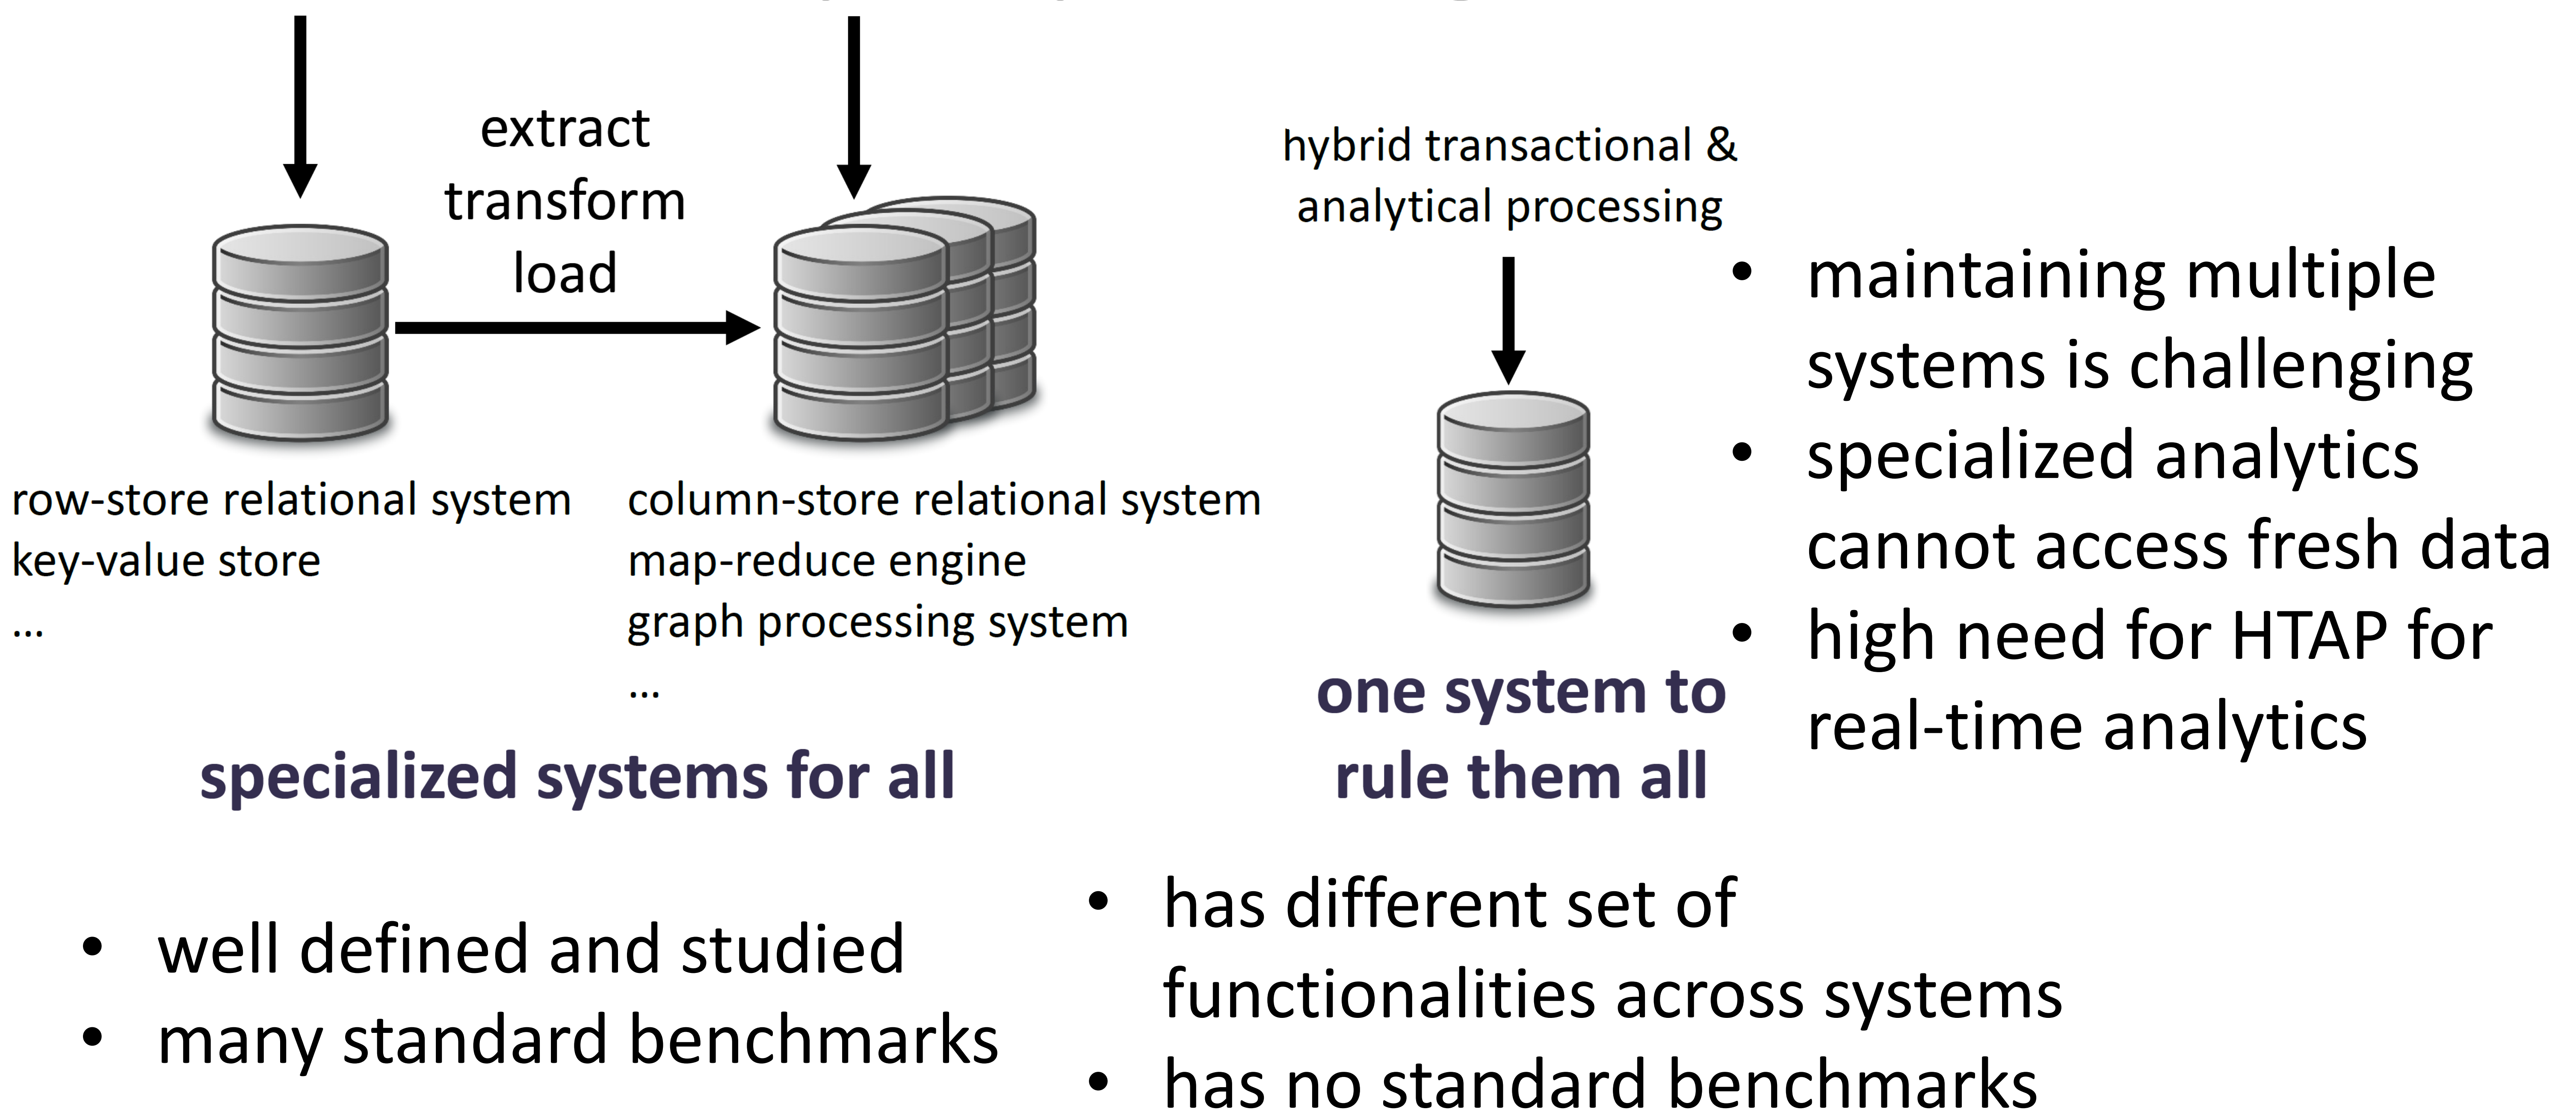
An Analysis of Open-Source HTAP Platforms – Survey

Jeppe Koch Schwensen

Supervisor: Pinar Tözün

Hybrid Transactional & Analytical Processing

transaction vs. analytical processing



How to standardize HTAP?

Open-Source HTAP Systems – Functionality Comparison

Aspect	Industry					Academia		Apache
	Streaming Ledger on Apache Flink	Aerospike	MongoDB	VoltDB	TiDB	Peloton	Ermia	Kudu
Engine	Streaming Processor	NoSQL Database	NoSQL Database	Relational Database	Relational Database	Relational Database	Storage Engine	Storage Engine
Targets	Transactions	Transactions	Transactions	Transactions	HTAP	HTAP	Transactions	HTAP
Cluster Deployment	Distributed †	Distributed	Distributed	Distributed	Distributed	Standalone	Standalone	Distributed
Storage Architecture	External Storage	Disk-oriented	Hybrid †	In-memory	Disk-oriented	In-memory	In-memory	Disk-oriented
Storage Model	External Storage	Row-oriented Layout	Row-oriented, Column-oriented, LSM Tree	Row-oriented Layout	Row-oriented Layout	Hybrid-oriented Layout	Append-only Row-oriented	Column-oriented Layout
Data Model	External Storage	Key-Value Store	Document-Store	Relational	Key-Value Store	Relational	Relational	Relational
Logging	Depends on external storage	Write-ahead Logging	Write-ahead Logging	Write-ahead Logging	Write-ahead Logging	Write-behind Logging	Single round-trip Write-ahead Logging	Write-ahead Logging
Concurrency Control	Timestamps, watermarks and out-of-order scheduling	Partitioned, Single-Threaded Data Structures	Optimistic Concurrency Control	Deterministic Concurrency Control	2-Phase Commit MVCC	Multi-version Timestamp Ordering	Snapshot-based Concurrency Control with Serial Safety Net	Multi-version Concurrency Control
Semantics	Multi-row, multi-table serializable ACID	Single-row Atomicity and Durability	Multi-row, multi-document serializable ACID	Multi-row, multi-partition serializable ACID	Multi-row, multi-table serializable ACID	Single-row serializable ACID	Single-row serializable ACID	Single-row serializable ACID
Indexing	None	Hashed B+-tree	B+-tree	B+-tree	B+-tree	Lock-free Bw-tree	MassTree	MassTree
Available Benchmarks	YCSB	YCSB Custom Workload	TPC-C YCSB	TPC-C YCSB	TPC-C TPC-H TPC-DS	TPC-C YCSB CH-benCHmark	TPC-C TPC-E YCSB	TPC-H YCSB
Implementation Languages	Java	C	C++ JavaScript	Java C++	Go Rust	C++ LLVM	C++	C++ Java
User-facing Languages	Java Scala	Java C	Java C++	Java C++	MySQL Spark SQL	PostgreSQL	C++	Apache Spark Apache Impala
Documentation Quality	Acceptable	Comprehensive	Comprehensive	Comprehensive	Massive	Comprehensive	Low	Acceptable
Notes				Supports true HTAP		Self-driving Components †		

† The feature is hidden behind a pay-wall or is not a part of the default database solution.

As a result of this survey, Kudu, TiDB, and VoltDB are picked to compare via experiments using HTAPBench.