The mixed workload CH-BenCHmark

➔Hybrid OLTP&OLAP Database Systems
➔Real-Time Business Intelligence
➔Analytical information at your fingertips

Richard Cole (ParAccel), Florian Funke (TU München), Leo Giakoumakis (Microsoft), Wey Guy (Microsoft), **Alfons Kemper (TU München),** Stefan Krompass (TU München), Harumi Kuno (HP Labs), Raghunath Nambiar (Cisco), Thomas Neumann (TU München), Meikel Poess (Oracle), Kai-Uwe Sattler (TU Ilmenau), Eric Simon (SAP), Florian Waas (Greenplum)

Outcome of the Dagstuhl Seminar Fall 2010

- Robust Query Processing
  - Organized by Goetz Graefe et al.
- Breakout Working Group
  - Workload Management
  - Headed by: Harumi Kuno

State of the Art: Separate Transaction (OLTP) and Query (OLAP) Systems

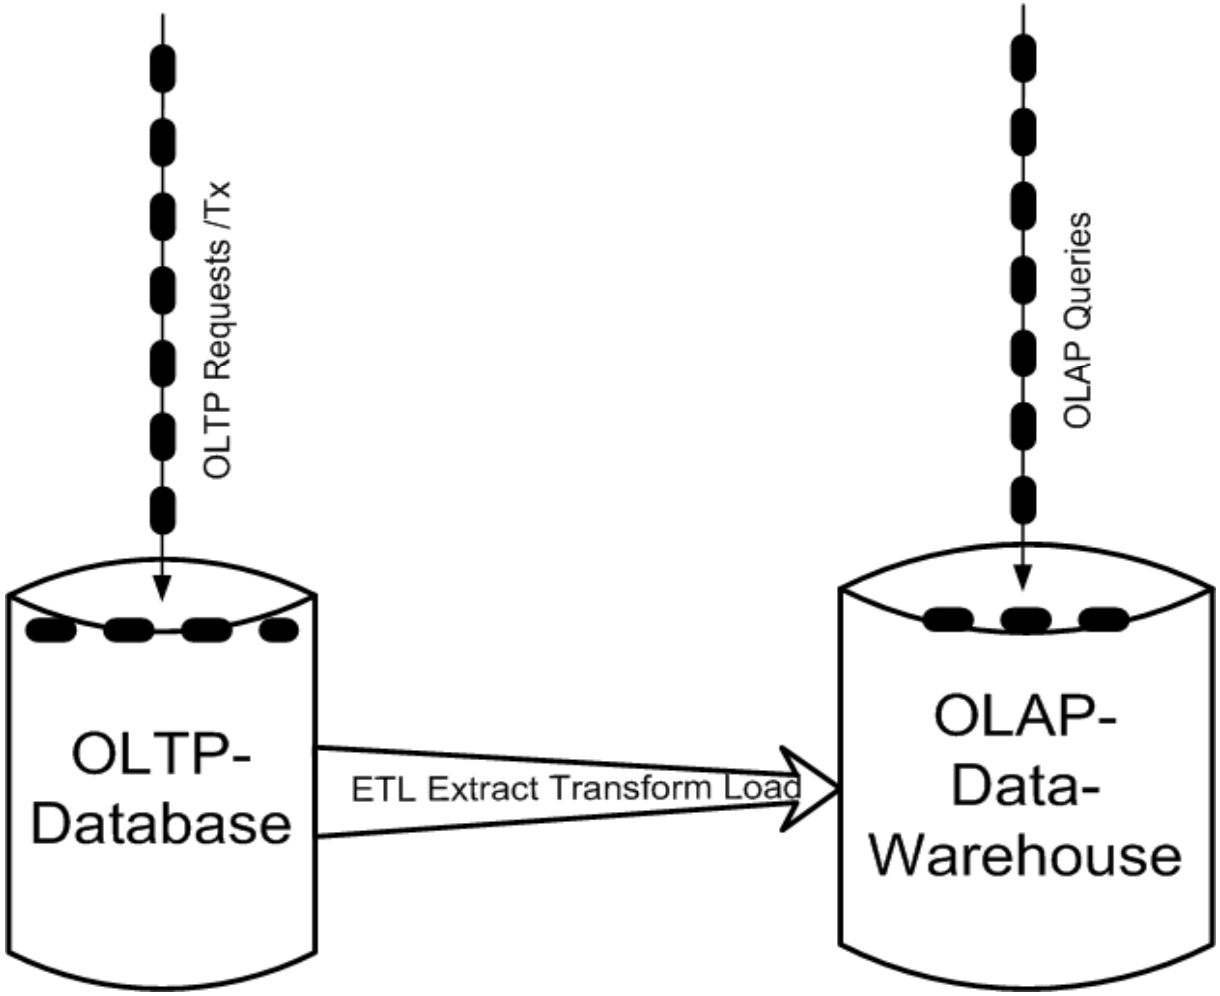

Goal: Real Time Business Intelligence
➔ Querying the Transactional Data

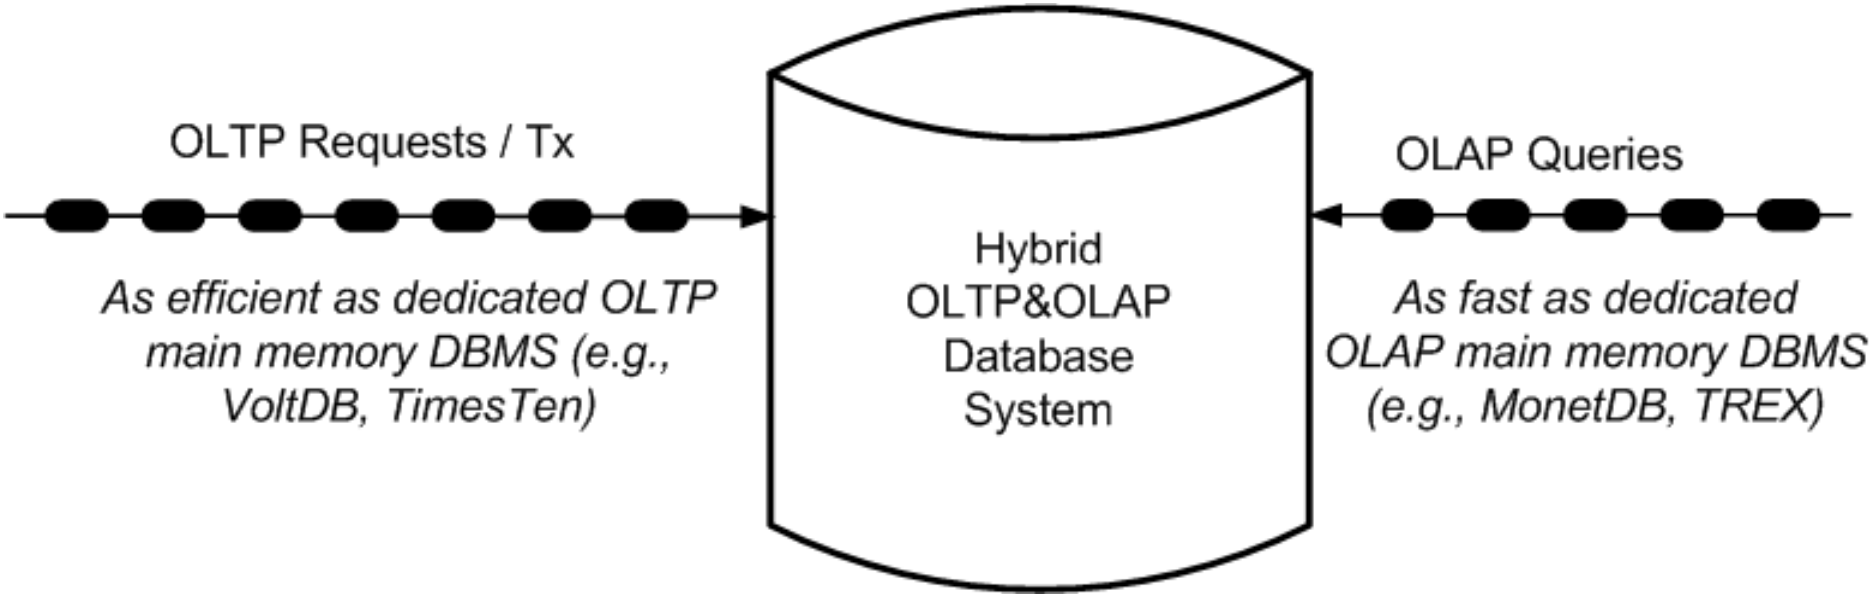

Hasso Plattner (SAP): Keynote at SIGMOD 09

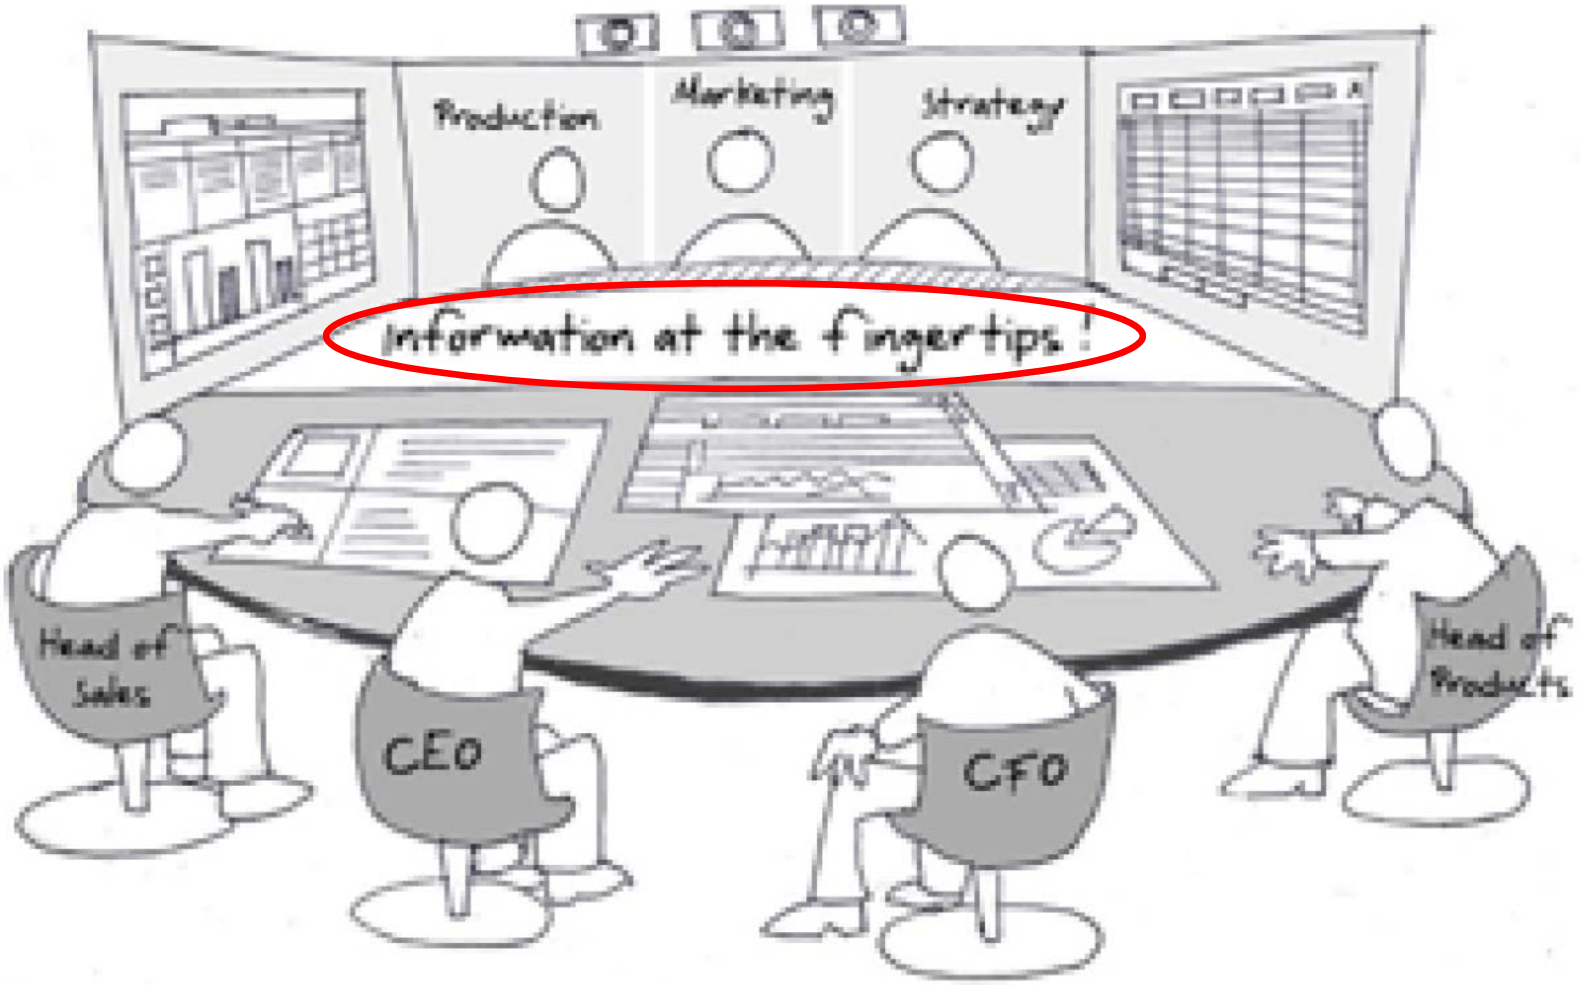

Use cases for low latency analytics [Curt Monash's Blog (April 11, 2011), Teradata]

- BI dashboards
  - 7 X 24 real time operations
- Financial peak periods
  - Month end, quarter end
- Cyber Security
  - Short and long term threats
- Operational reporting
  - Claims processed
  - Inventory instant status
- Machine generated data
  - Rapid response
  - Fast analytics

Frankly, I think low-latency monitoring is going to be one of the hot areas over the next few years. "Real-time" is cool, and big monitors with constantly changing graphics are cooler yet. [C.M.]

The Best of Both Worlds …

…. one size fits all – again??

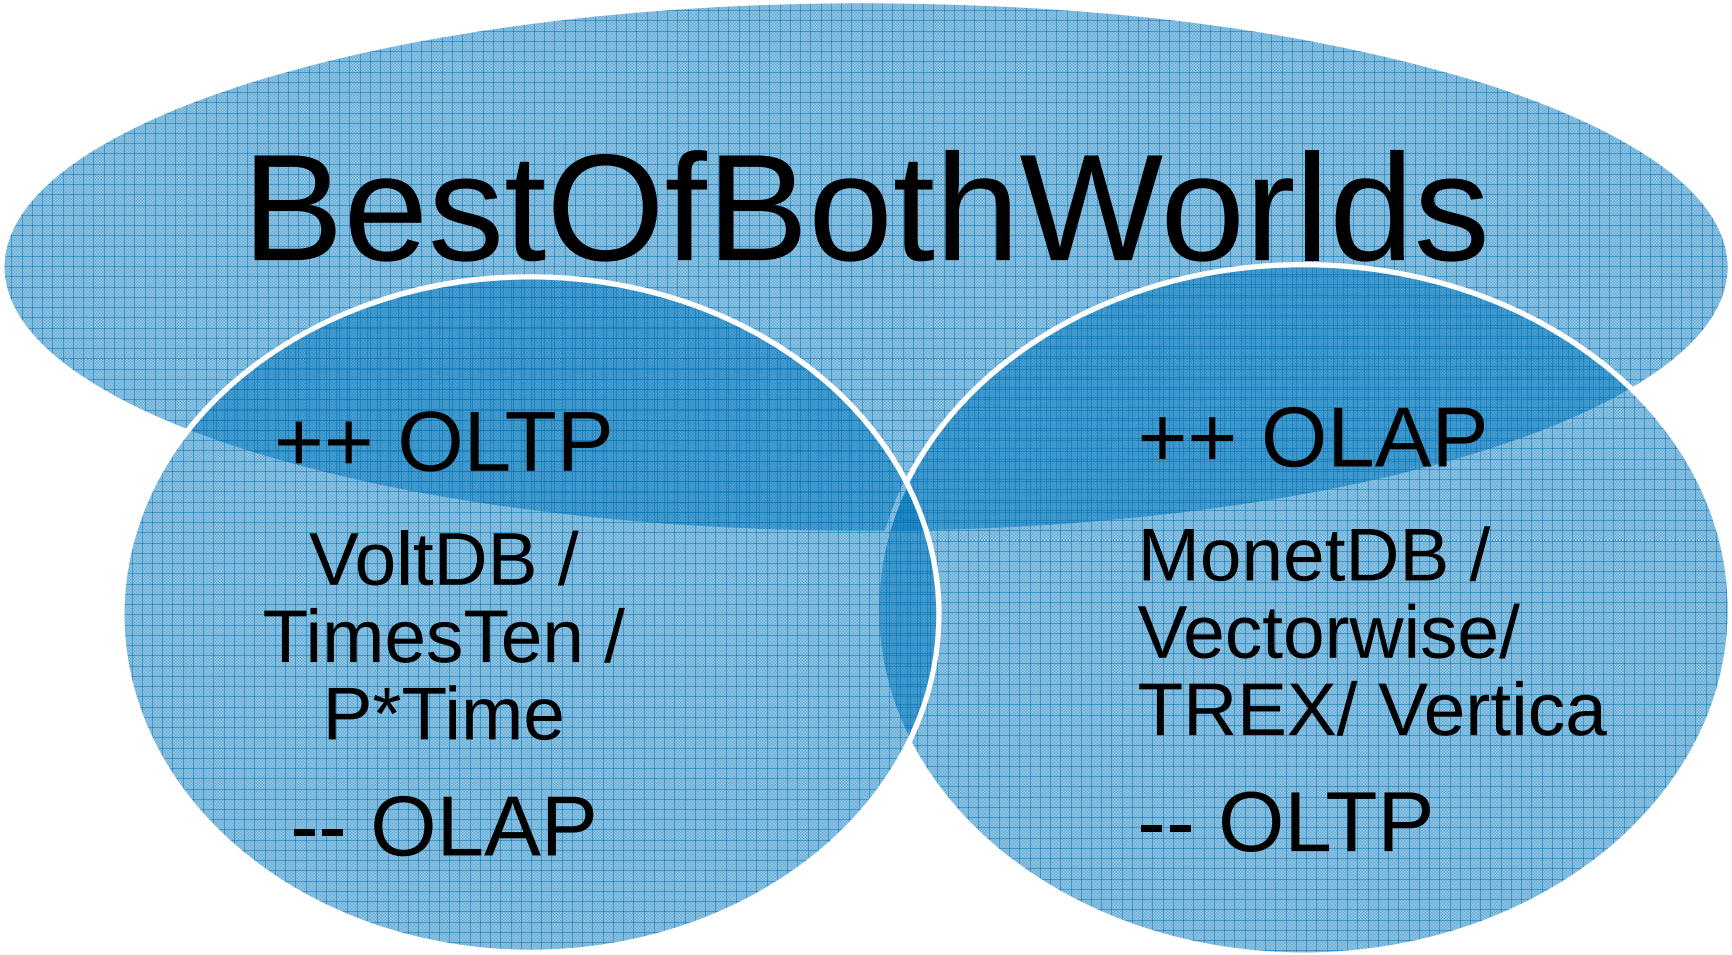

TPC-C and TPC-H Schemas

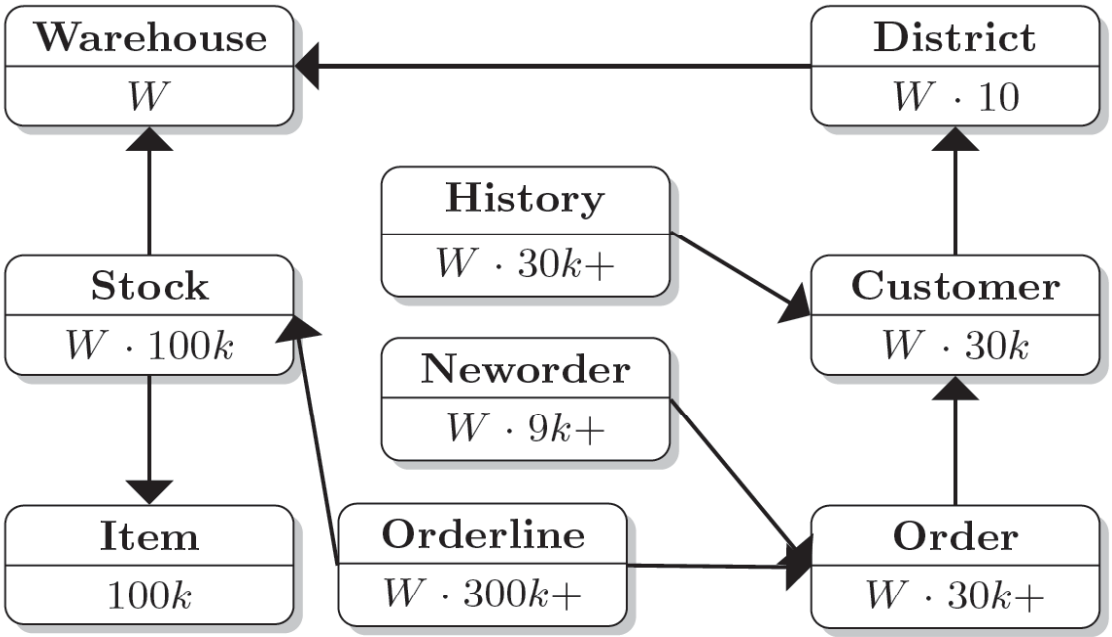

(a) TPC-C

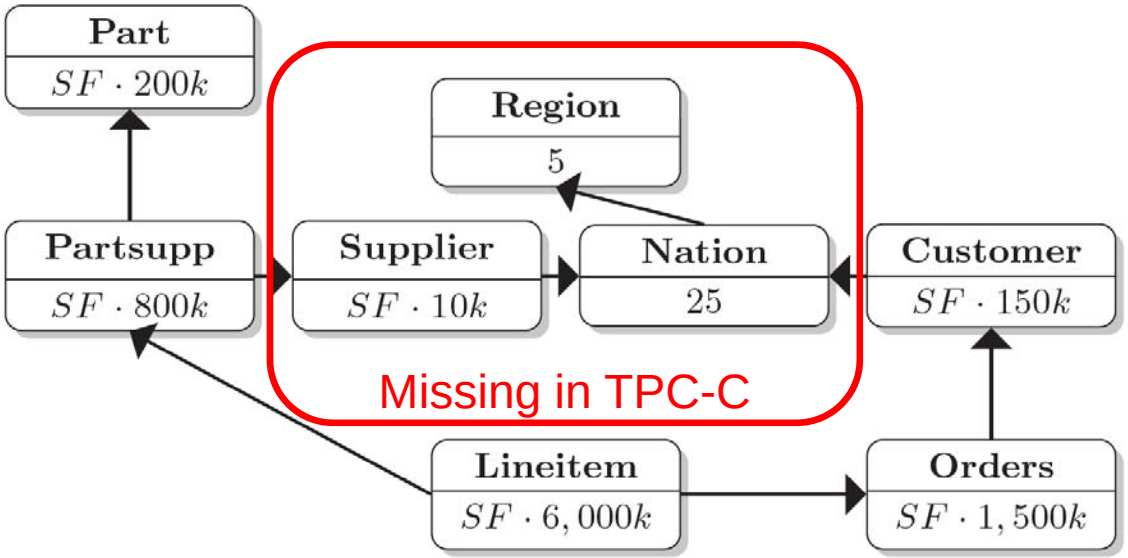

(b) TPC-H

C&H BenCHmark schema

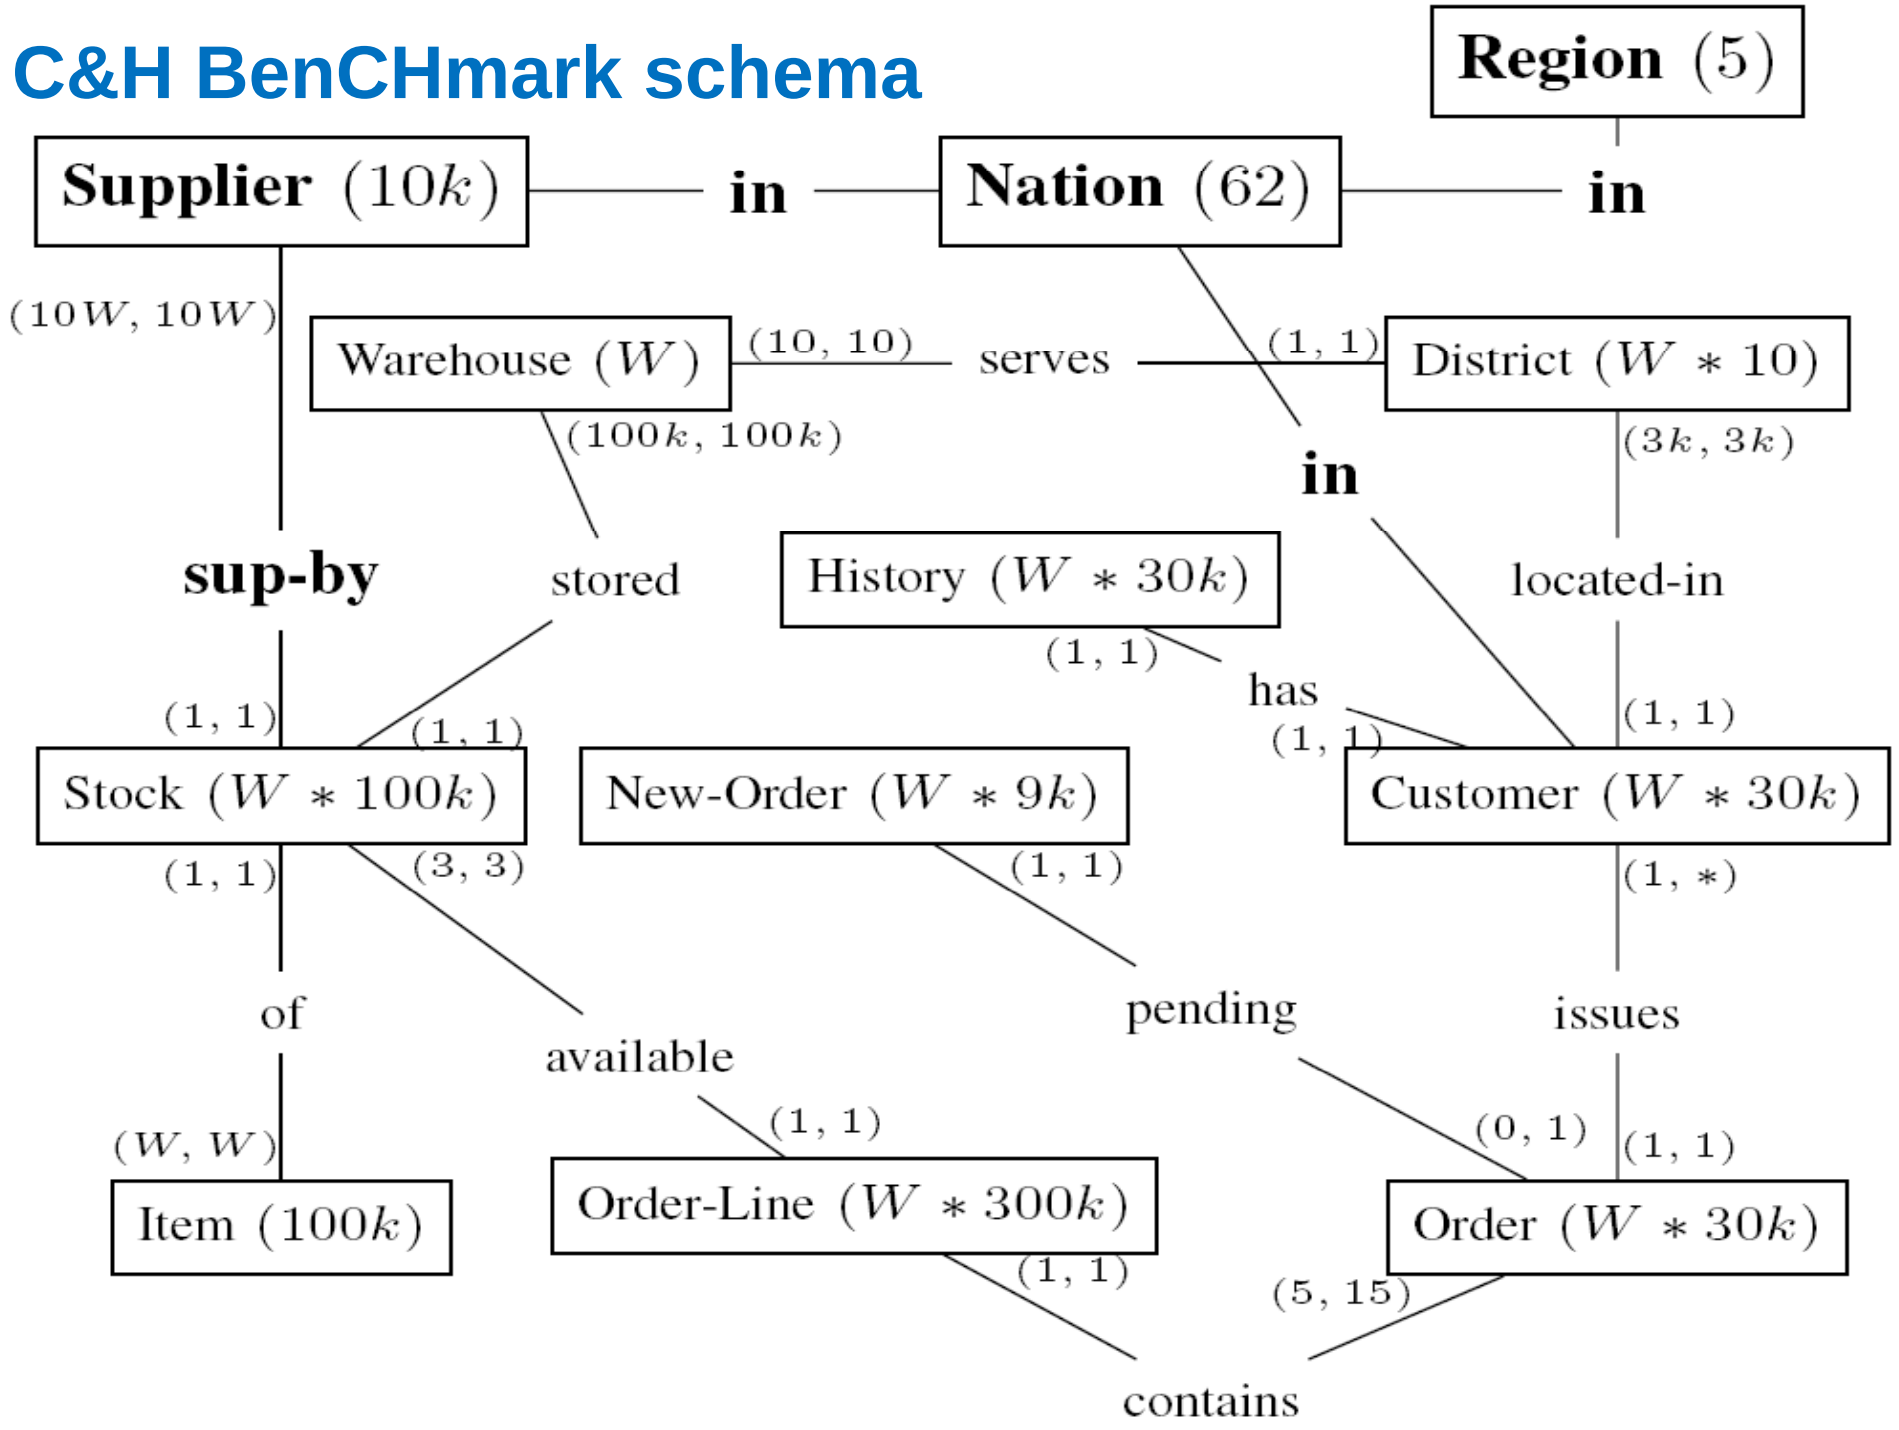

Mixed OLTP&OLAP Workload

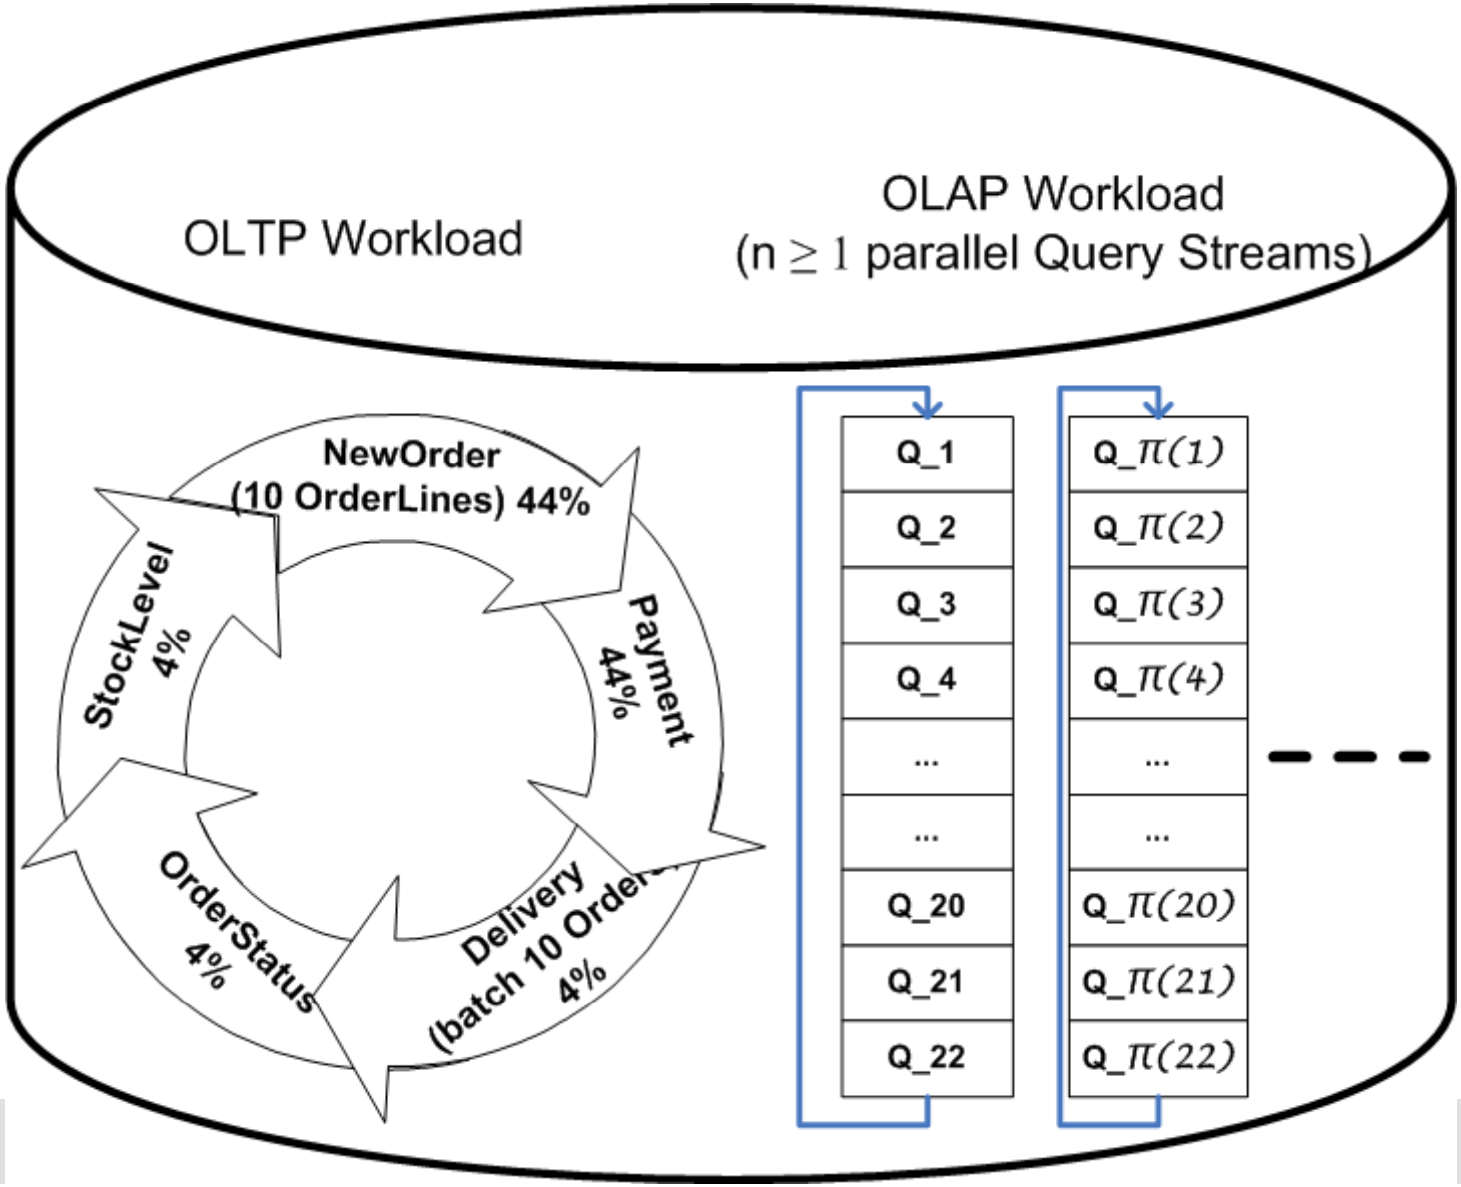

Re-use existing TPC-C-Benchmark Kit

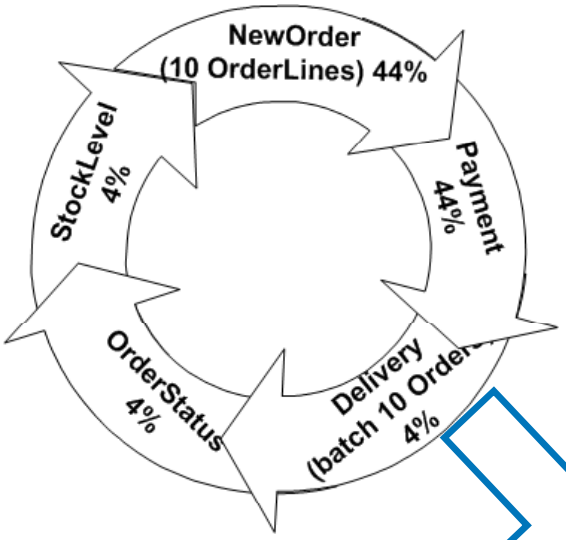

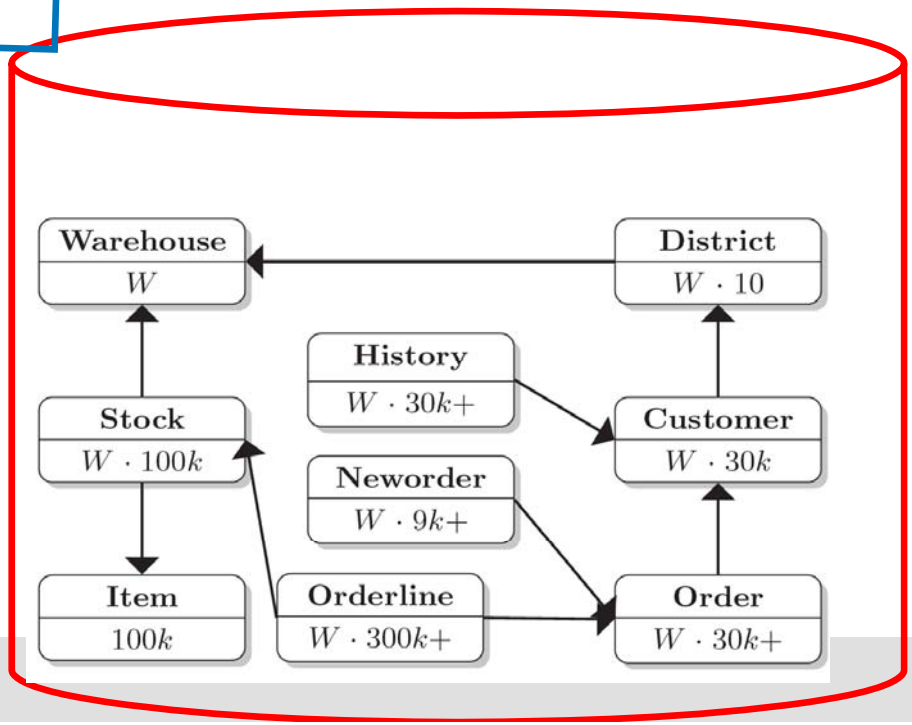

Re-use existing TPC-C-Benchmark Kit

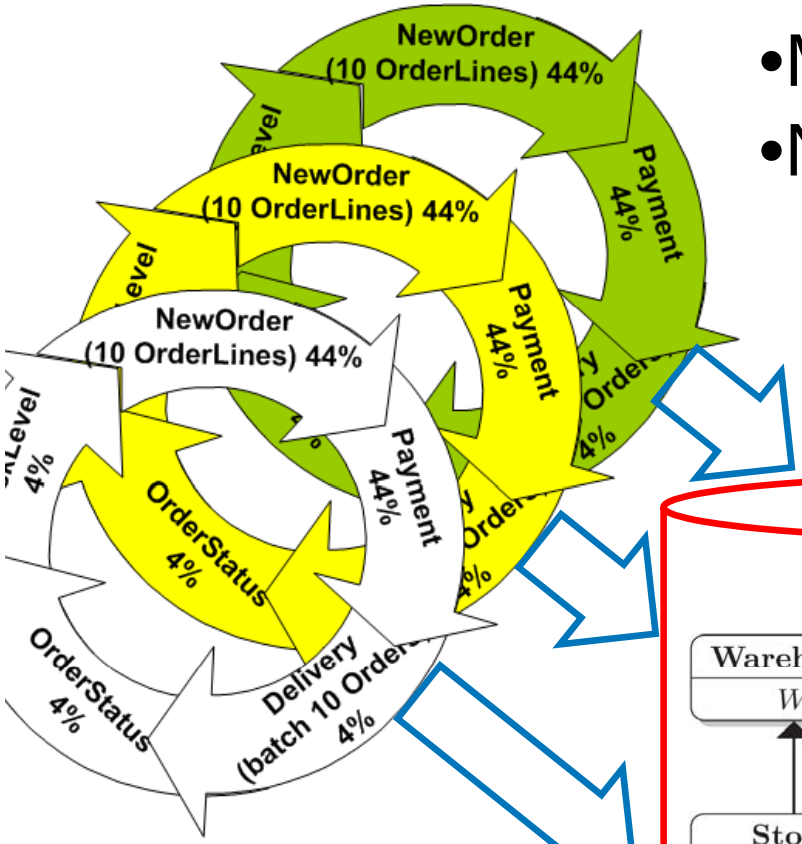

- Multiple OLTP clients
- No wait-time in between requests
  - Deviating from original TPC-C
  - High throughput for smaller DB

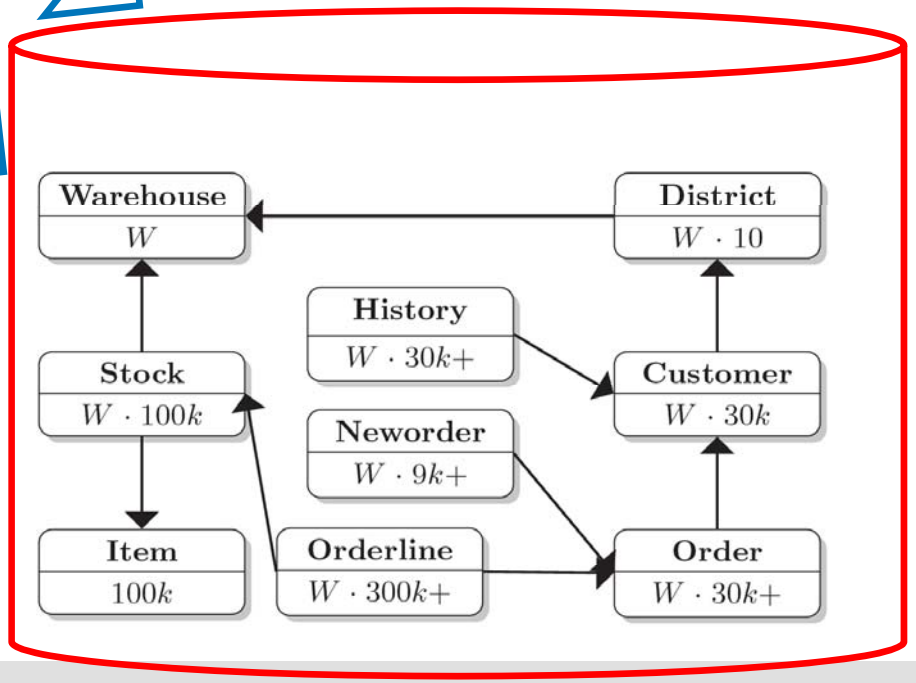

No Keying/Think-Time ➔ Clients generate one request after another as fast as possible

Transaction Type	Minimum % of mix	Minimum Keying Time	90th Percentile Response Time Constraint	Minimum Mean of Think Time Distribution
New-Order	n/ a	18 sec.	5 sec.	12 sec.
Payment	43.0	3 sec.	5 sec.	12 sec.
Order-Status	4.0	2 sec.	5 sec.	10 sec.
Delivery [1]	4.0	2 sec.	5 sec.	5 sec.
Stock-Level	4.0	2 sec.	20 sec.	5 sec.

0 sec.

0 sec.

- ~~10 clients=terminals per Warehouse~~

Re-use existing TPC-C-Benchmark Kit

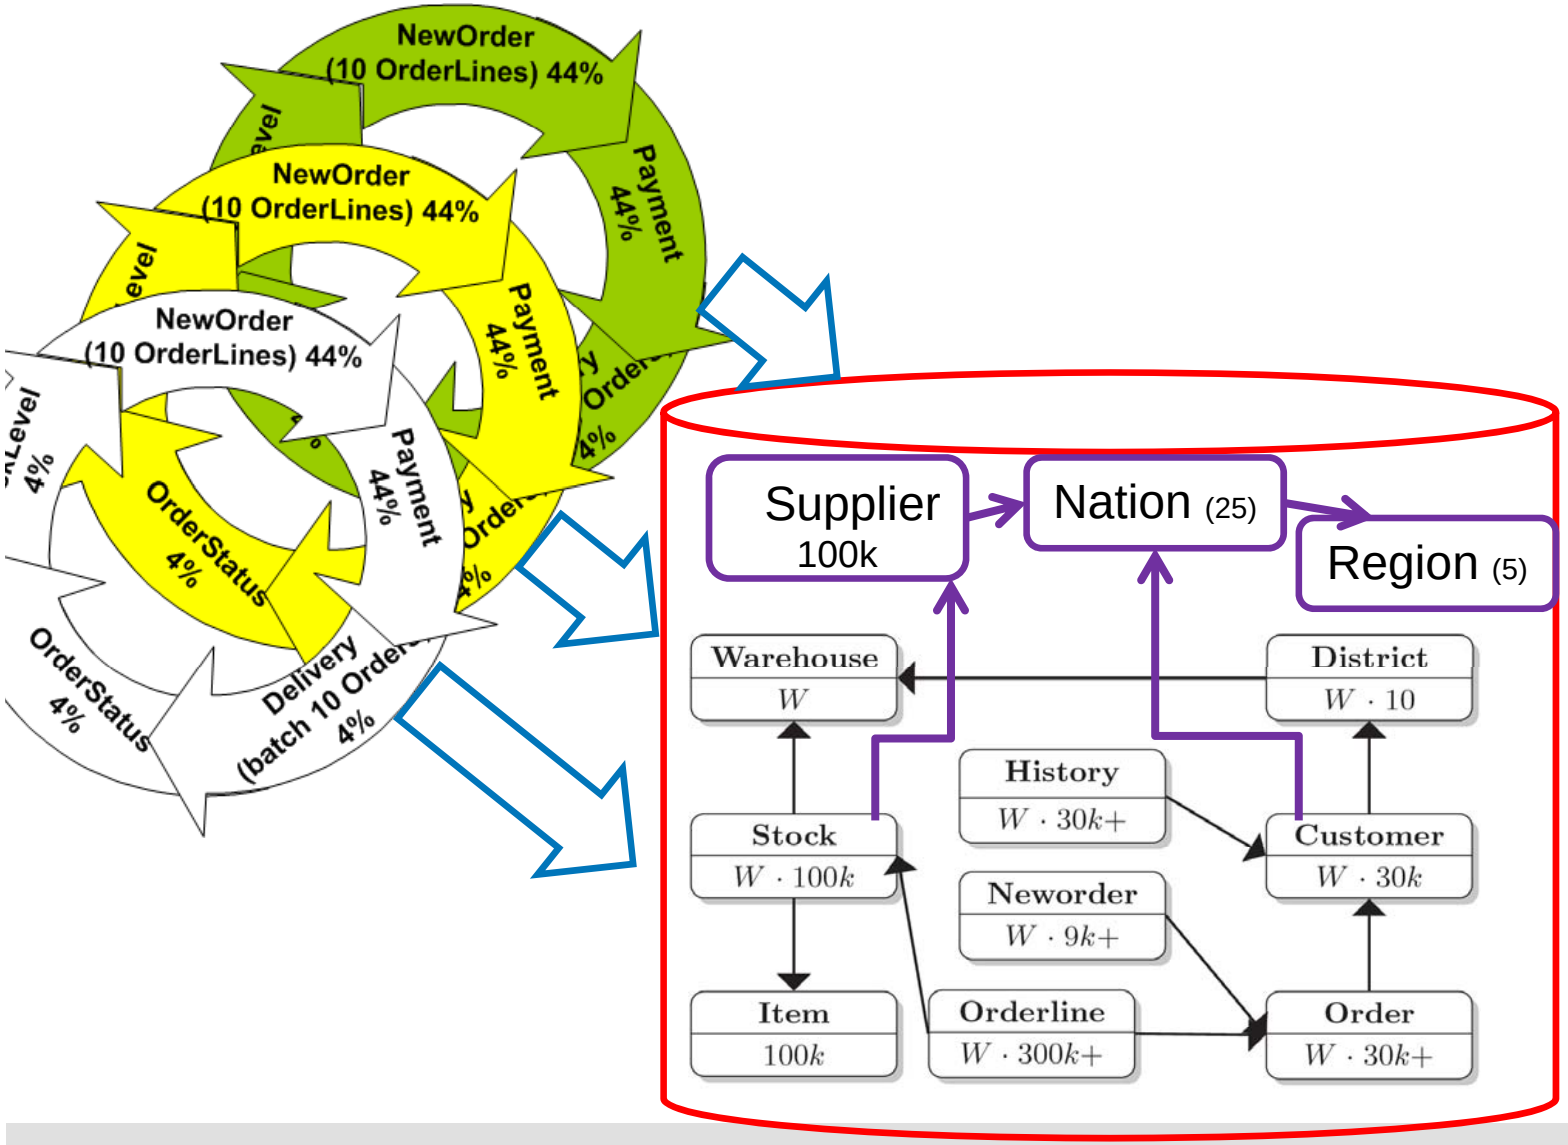

Re-use existing TPC-C-Benchmark Kit

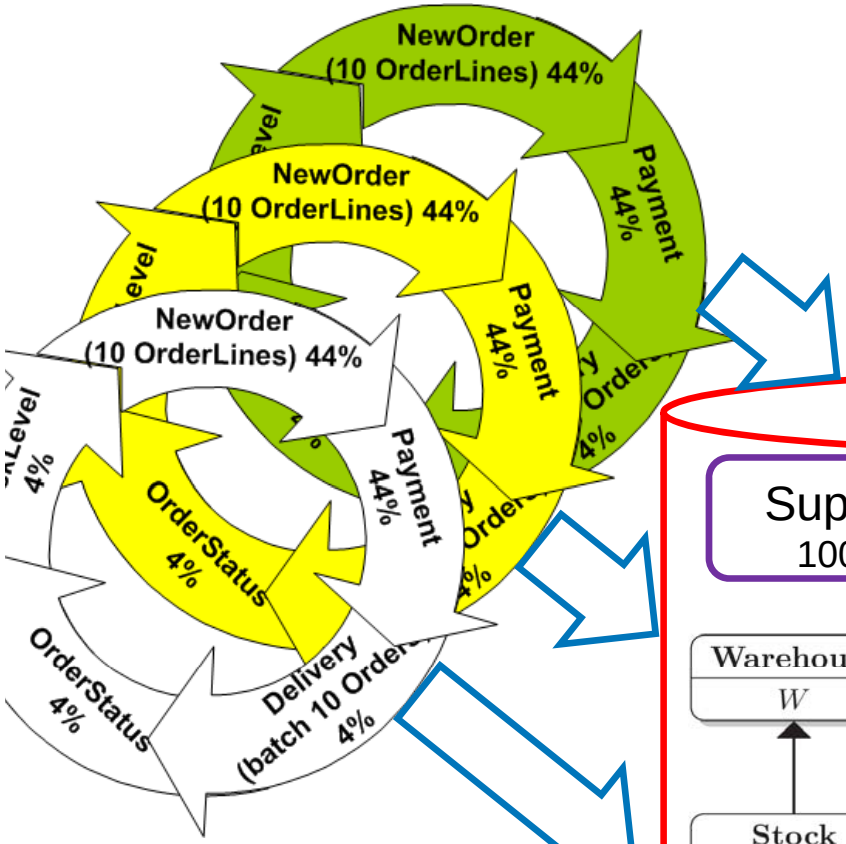

- No updates because new data is generated by OLTP
- Modified TPC-H queries
  - Different schema

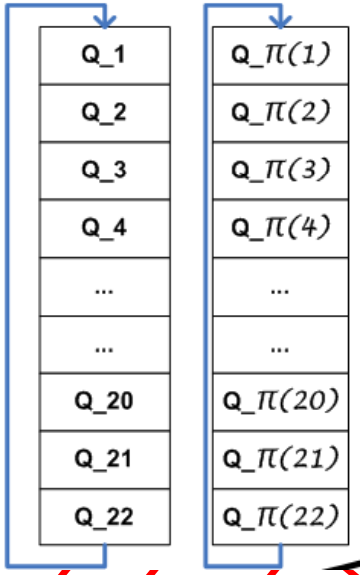

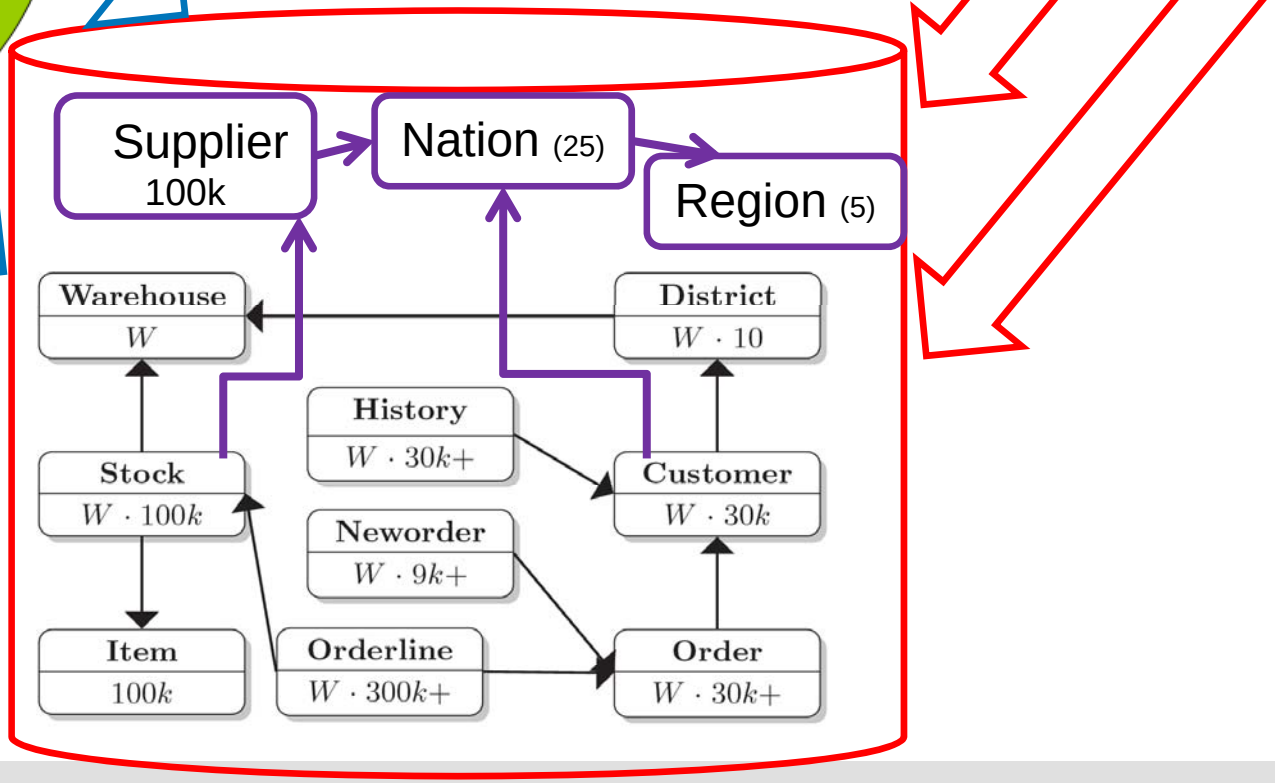

All 5 TPC-C Transactions (no waiting time)

- New-Order
- Payment
- Delivery
- Order-Status
- Stock-Level

All 22 TPC-H Queries

e.g., Query 5 : Intra Country – Revenue by local Suppliers within a Region, per Nation

```
select n_name,
       sum(ol_amount) as revenue
from Nation join Customer on ... join Order on ...
         join Order-Line on ... join Stock on ...
         join Supplier on ... join Region on ...
where su_nationkey = n_nationkey  /* Cu and Su in the *
      and r_name = 'Europe'       /* same N of this R *
      and o_entry_d >= ...
group by n_name
order by revenue desc;
```

Complete Query Suite

- Q1: Generate orderline overview
- Q2: Most important supplier/item-combinations (those that have the lowest stock level for certain parts in a certain region)
- Q3: Unshipped orders with highest value for customers within a certain state
- Q4: Orders that were partially shipped late
- Q5: Revenue volume achieved through local suppliers
- Q6: Revenue generated by orderlines of a certain quantity
- Q7: Bi-directional trade volume between two nations
- Q8: Market share of a given nation for customers of a given region for a given part type

Complete Query Suite

- Q9: Profit made on a given line of parts,broken out by supplier nation and year
- Q10: Customers who received their ordered products late
- Q11: Most important (high order count compared to the sum of all ordercounts) parts supplied by suppiers of a particular nation
- Q12: Determine whether selecting less expensive modes of shipping is negatively affecting the critical-priority orders by causing more parts to be received late by customers
- Q13: Relationships between customers and the size of their orders
- Q14: Market response to a promotion campaign

Complete Query Suite

- Q15: Determines the top supplier
- Q16: Number of suppliers that can supply particular parts
- Q17: Average yearly revenue that would be lost if orders were no longer filled for small quantities of certain parts
- Q18: Rank customers based on their placement of a large quantity order
- Q19: Machine generated data mining (revenue report for disjunctive predicate)
- Q20: Suppliers in a particular nation having selected parts that may be candidates for a promotional offer
- Q21: Suppliers who were not able to ship required parts in a timely manner
- Q22: Geographies with customers who may be likely to make a purchase

Performance and Quality Metrics

Performance

- OLTP Throughput
  - NewOrder Tx per minute
- Query Response Times
  - Geometric Mean
    - One query stream
    - Multiple query streams
- Query Throughput
  - Multiple parallel streams
  - #Queries per hour

Quality

- Isolation Level
  - Serializable for OLTP
  - Except Stock-Level
- Query isolation level
  - Read uncommitted (dirty)
  - Read committed
  - Serializable
  - Snapshot
    - Freshness of the Snapshot
      - In #missed transactions
- Response time guarantees

→ derived from TPC-C

First Results from PostgreSQL to Demonstrate the Reporting (out of the box → no fine-tuning)

	PostgreSQL/CH-BenCHmark configurations							
	1 query session (stream) single threaded OLTP		1 query session (stream) 8 OLTP workers		4 query sessions (streams) 16 OLTP workers		no OLTP	
Q#	OLTP throughput	response times (ms)	OLTP throughput	response times (ms)	OLTP throughput	response times (ms)	1 Q. stream response times (ms)	8 Q. streams response times (ms)
Q1	new order: 5188 *tpmC*	5203	new order: 12550 *tpmC*	6502	new order: 18356 *tpmC*	15214	4466	4346
Q2		1068		1017		2009	986	964
Q3		187		243		470	166	180
Q4		1392		1881		4271	1145	1096
Q5		6418		9593		29292	5548	6900
Q6		1965		2537		5436	1712	1740
Q7		728		948		1466	2548	706
Q8		1568		2098		4583	1417	1442
Q9		584		706		1737	533	521
Q10		7118		8935		22549	6280	6578
Q11		582		624		1233	567	546
Q12		3143		4086		10179	2694	2567
Q13		475		545		1156	525	483
Q14		3966		5346		11716	3479	3843
Q15		11768		15145		26160	9162	9831
Q16		13837		14132		29183	14091	14884
Q17		1489		3792		6837	2183	1312
Q18		30271		38966		96159	26971	27762
Q19		3586		5041		12990	3082	3313
Q20		906		1153		2347	776	744
Q21		1921		2456		5653	1698	1716
Q22		303		361		784	257	239
Geometric mean (ms)		2284		2992		6747	2105	2092
Duration per query set (s)		98		126		292	88	91
Queries per hour (QphH)		804.2 × 1		628.0 × 1		270.8 × 4	895.9 × 1	863.4 × 8

First Results from PostgreSQL: „Powertest“

	PostgreSQL/CH-BenCHmark configurations							
	1 query session (stream) single threaded OLTP		1 query session (stream) 8 OLTP workers		4 query sessions (streams) 16 OLTP workers		n 1 Q. strean	
Q#	OLTP throughput	response times (ms)	OLTP throughput	response times (ms)	OLTP throughput	response times (ms)	response	times (ms)
Q1	new order: 5188 *tpmC*	5203	new order: 12550 *tpmC*	6502	new order: 18356 *tpmC*	15214		446
Q2		1068		1017		2009		98
Q3		187		243		470		16
Q4		1392		1881		4271		114
Q5		6418		9593		29292		554
Q6		1965		2537		5436		171
Q7		728		948		1466		254
Q8		1568		2098		4583		141
Q9		584		706		1737		53
Q10		7118		8935		22549		628
Q11		582		624		1233		56
Q12		3143		4086		10179		269
Q13		475		545		1156		52
Q14		3966		5346		11716		347
Q15		11768		15145		26160		916
Q16		13837		14132		29183		1409
Q17		1489		3792		6837		218
Q18		30271		38966		96159		2697
Q19		3586		5041		12990		308
Q20		906		1153		2347		77
Q21		1921		2456		5653		169
Q22		303		361		784		25
Geometric mean (ms)		2284		2992		6747		210
Duration per query set (s)		98		126		292		8
Queries per hour (QphH)		804.2 × 1		628.0 × 1		270.8 × 4		895.9 ×

First Results from PostgreSQL: „OLTP centric“

	PostgreSQL/CH-BenCHmark configurations					
	1 query session (stream) single threaded OLTP		1 query session (stream) 8 OLTP workers		4 query sessions (streams) 16 OLTP workers		n 1 Q. strean
Q#	OLTP throughput	response times (ms)	OLTP throughput	response times (ms)	OLTP throughput	response times (ms)	response times (ms)
Q1	new order: 5188 *tpmC*	5203	new order: 12550 *tpmC*	6502	new order: 18356 *tpmC*	15214	446
Q2		1068		1017		2009	98
Q3		187		243		470	16
Q4		1392		1881		4271	114
Q5		6418		9593		29292	554
Q6		1965		2537		5436	171
Q7		728		948		1466	254
Q8		1568		2098		4583	141
Q9		584		706		1737	53
Q10		7118		8935		22549	628
Q11		582		624		1233	56
Q12		3143		4086		10179	269
Q13		475		545		1156	52
Q14		3966		5346		11716	347
Q15		11768		15145		26160	916
Q16		13837		14132		29183	1409
Q17		1489		3792		6837	218
Q18		30271		38966		96159	2697
Q19		3586		5041		12990	308
Q20		906		1153		2347	77
Q21		1921		2456		5653	169
Q22		303		361		784	25
Geometric mean (ms)		2284		2992		6747	210
Duration per query set (s)		98		126		292	8
Queries per hour (QphH)		804.2 × 1		628.0 × 1		270.8 × 4	895.9 ×

First Results: „balanced OLTP & OLAP“

Q#	PostgreSQL/CH-BenCHmark configurations: 1 query session (stream), single threaded OLTP — OLTP throughput	1 query session (stream), single threaded OLTP — response times (ms)	1 query session (stream), 8 OLTP workers — OLTP throughput	1 query session (stream), 8 OLTP workers — response times (ms)	4 query sessions (streams), 16 OLTP workers — OLTP throughput	4 query sessions (streams), 16 OLTP workers — response times (ms)	n 1 Q. strean response times (ms)
Q1	new order: 5188 *tpmC*	5203	new order: 12550 *tpmC*	6502	new order: 18356 *tpmC*	15214	446
Q2		1068		1017		2009	98
Q3		187		243		470	16
Q4		1392		1881		4271	114
Q5		6418		9593		29292	554
Q6		1965		2537		5436	171
Q7		728		948		1466	254
Q8		1568		2098		4583	141
Q9		584		706		1737	53
Q10		7118		8935		22549	628
Q11		582		624		1233	56
Q12		3143		4086		10179	269
Q13		475		545		1156	52
Q14		3966		5346		11716	347
Q15		11768		15145		26160	916
Q16		13837		14132		29183	1409
Q17		1489		3792		6837	218
Q18		30271		38966		96159	2697
Q19		3586		5041		12990	308
Q20		906		1153		2347	77
Q21		1921		2456		5653	169
Q22		303		361		784	25
Geometric mean (ms)		2284		2992		6747	210
Duration per query set (s)		98		126		292	8
Queries per hour (QphH)		804.2 × 1		628.0 × 1		270.8 × 4	895.9 ×

First Results: „Queries only“

PostgreSQL/CH-BenCHmark configurations							
ıery session (stream) gle threaded OLTP		1 query session (stream) 8 OLTP workers		4 query sessions (streams) 16 OLTP workers		no OLTP	
LTP ughput	response times (ms)	OLTP throughput	response times (ms)	OLTP throughput	response times (ms)	1 Q. stream response times (ms)	8 Q. streams response times (ms)
new order: 5188 *tpmC*	5203	new order: 12550 *tpmC*	6502	new order: 18356 *tpmC*	15214	4466	4346
	1068		1017		2009	986	964
	187		243		470	166	180
	1392		1881		4271	1145	1096
	6418		9593		29292	5548	6900
	1965		2537		5436	1712	1740
	728		948		1466	2548	706
	1568		2098		4583	1417	1442
	584		706		1737	533	521
	7118		8935		22549	6280	6578
	582		624		1233	567	546
	3143		4086		10179	2694	2567
	475		545		1156	525	483
	3966		5346		11716	3479	3843
	11768		15145		26160	9162	9831
	13837		14132		29183	14091	14884
	1489		3792		6837	2183	1312
	30271		38966		96159	26971	27762
	3586		5041		12990	3082	3313
	906		1153		2347	776	744
	1921		2456		5653	1698	1716
	303		361		784	257	239
n (ms)	2284		2992		6747	2105	2092
ıery set (s)	98		126		292	88	91
ır (QphH)	804.2 × 1		628.0 × 1		270.8 × 4	895.9 × 1	863.4 × 8

Tuning Dimensions

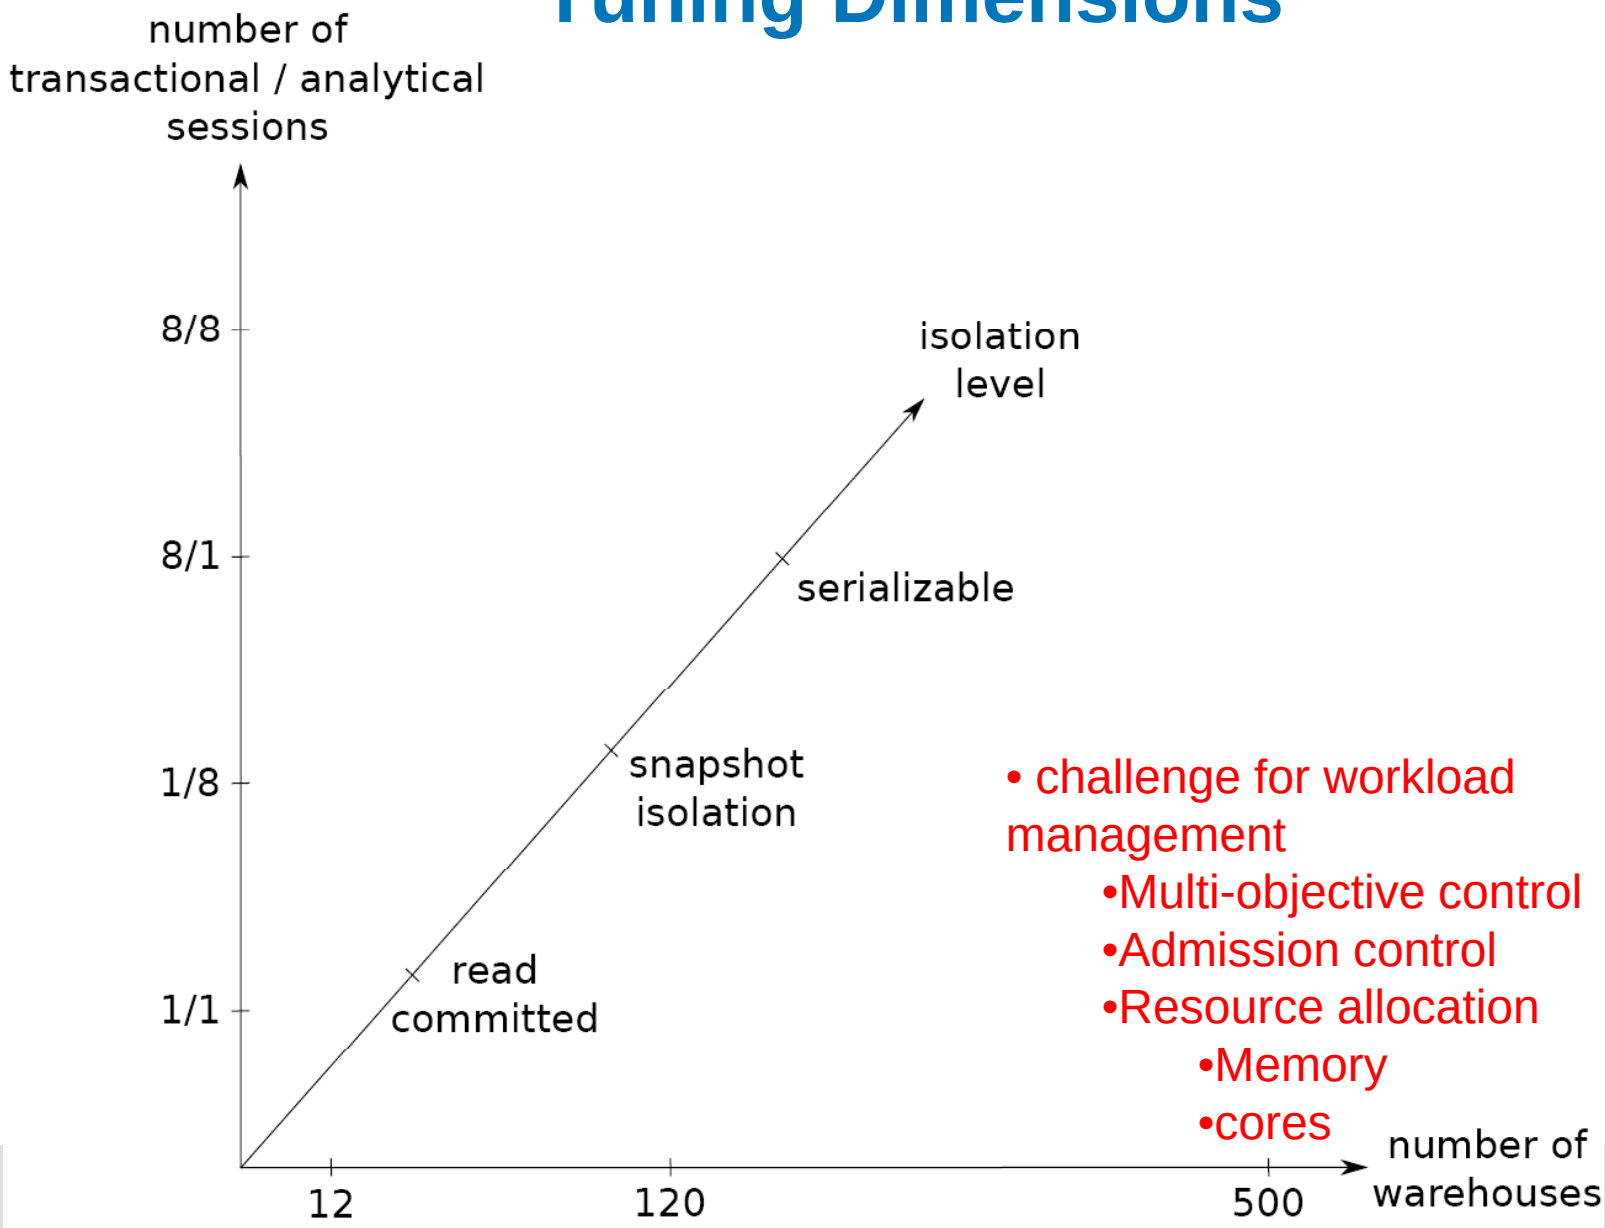

- challenge for workload management
  - Multi-objective control
  - Admission control
  - Resource allocation
    - Memory
    - cores

How to Gain Performance for Mixed Workload Processing: Snapshotting and Main Memory DBMS

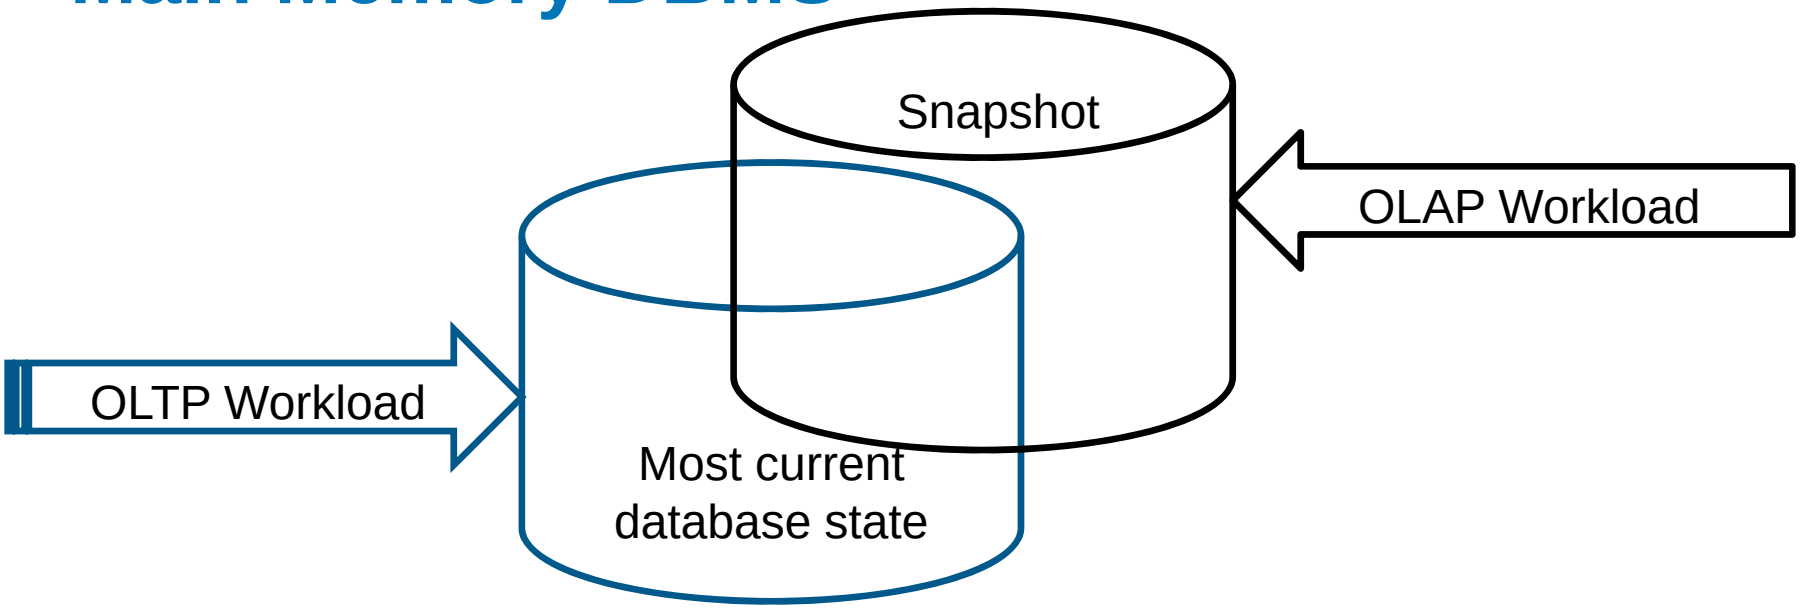

- Versioning: run OLAP on time versions of the data
- Twin block: run OLAP on Tx-consistent snapshot
- Shadowing
  - Tuple level
  - Page level → exploit hardware support for for Virtual Memory Snapshot (HyPer)

Future Work

- Fine-tune (tighten) the benCHmark specification
  - Query Parameters
  - Performance metrics
    - Account for dynamically growing database cardinality
    - Isolation levels
    - Freshness guarantees
- Get TPC.org interested to follow up

  ➔Industry representatives

How it works ... in HyPer

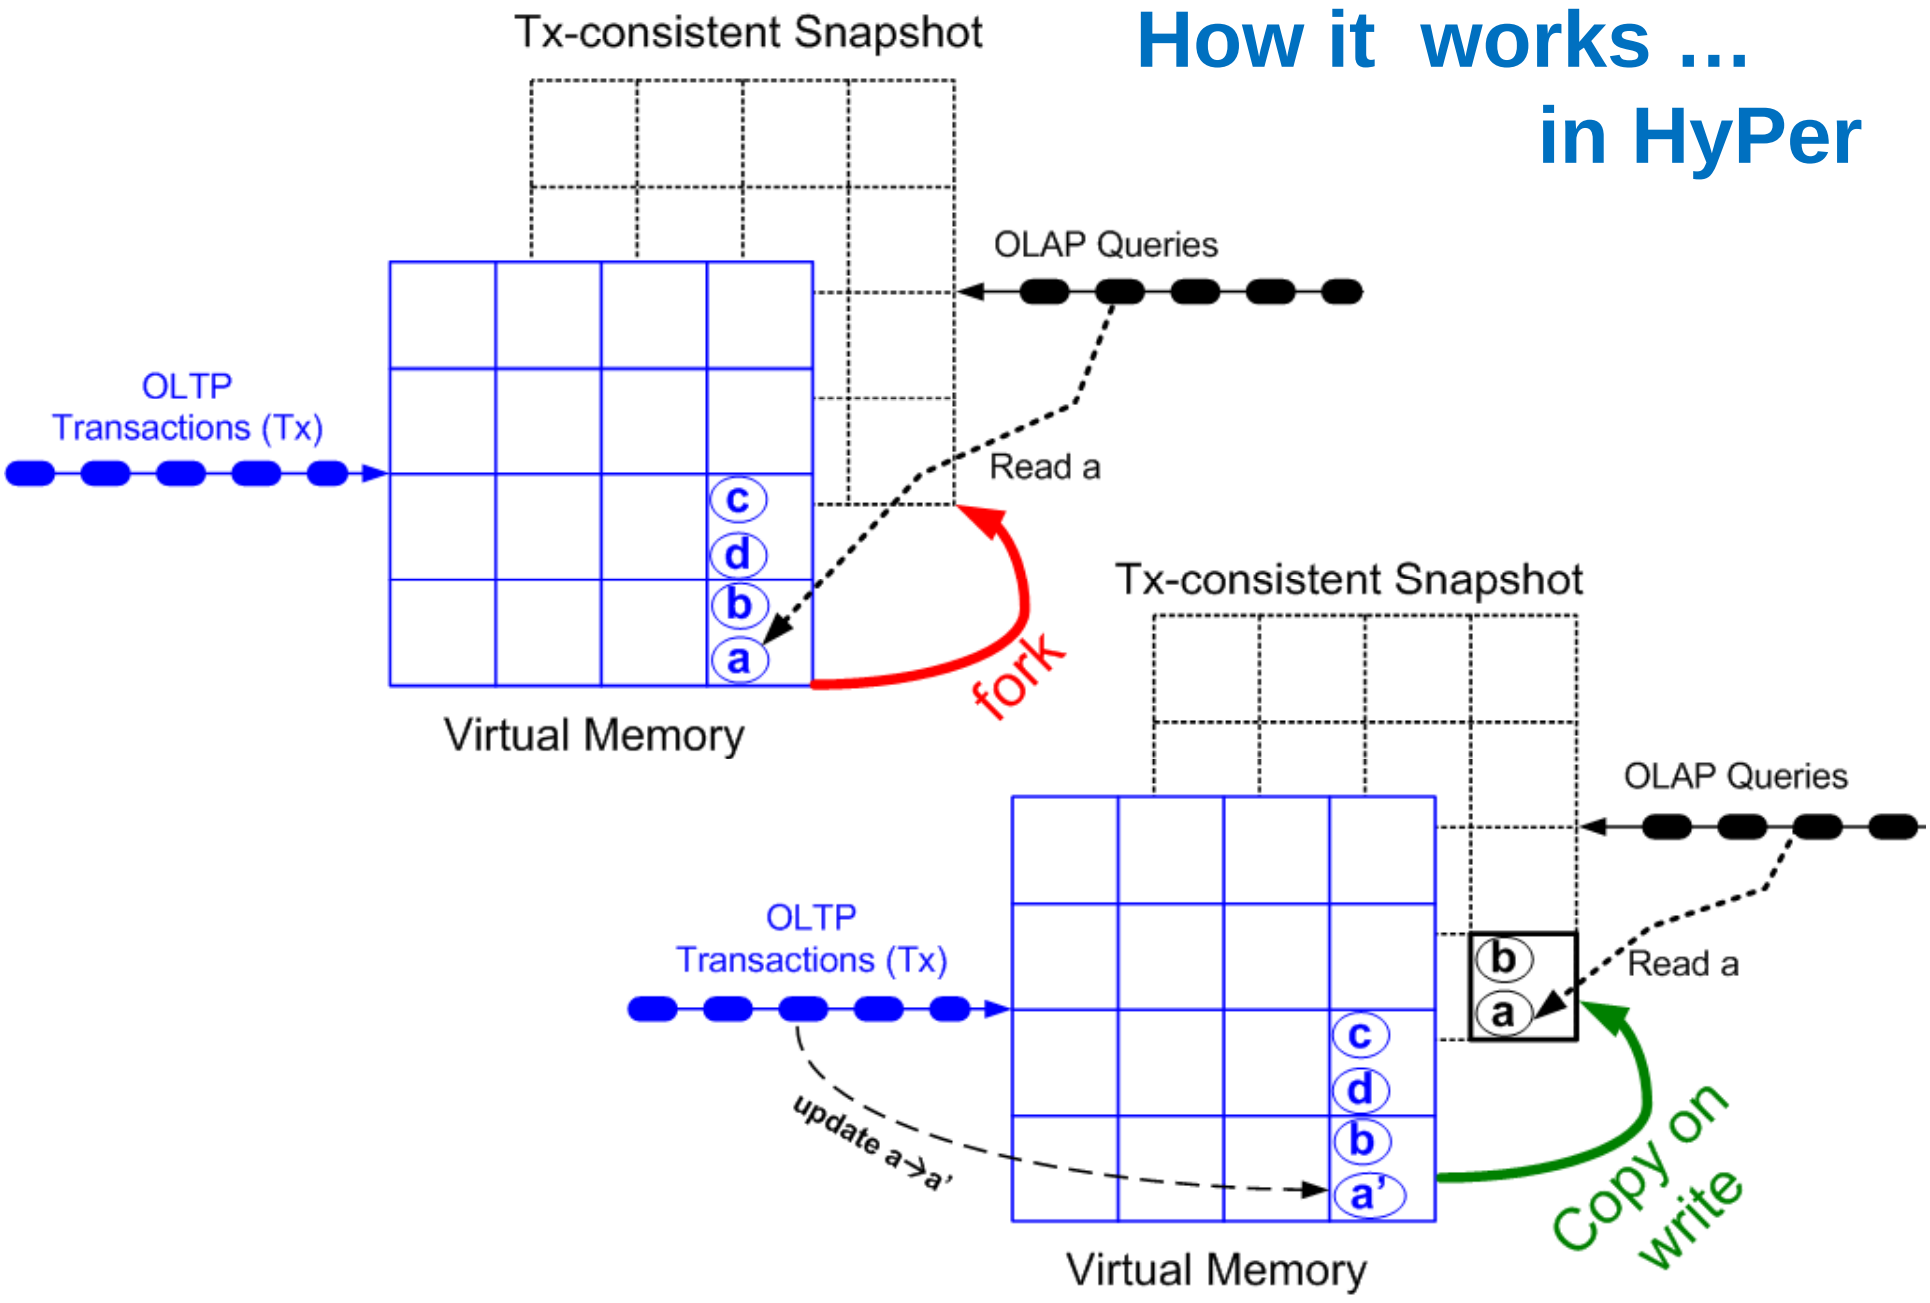

Hardware Supported Data Access and Copy on Update

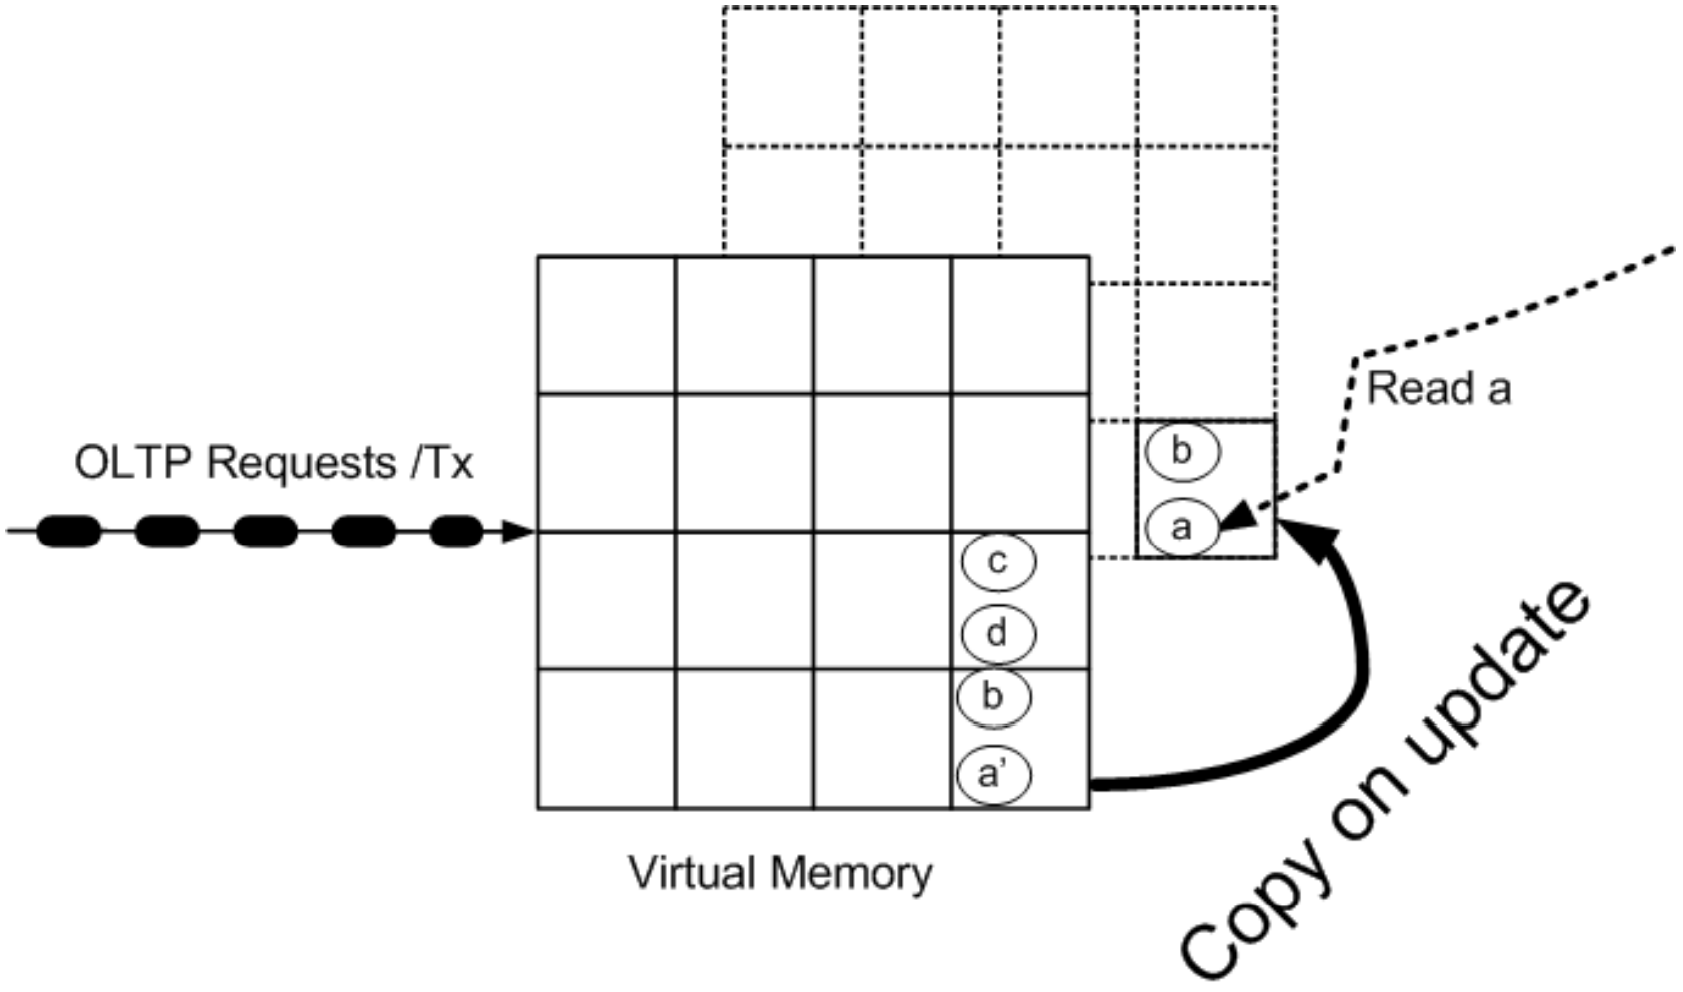

The Best of Both Worlds …