SAP HANA – The Evolution of an In-Memory DBMS from Pure OLAP Processing Towards Mixed Workloads

Norman May (SAP SE), Alexander Böhm (SAP SE), Wolfgang Lehner (TU Dresden)

March 9, 2017

INTERNAL

Run Simple

Stonebraker: One Size does not Fit All!

- Specialized Engines Beat General-Purpose DBMS by Factors
- Proofpoints
  - StreamBase – Event stream processing
  - Vertica – Columnar Storage for OLAP
  - VoltDB – Lock-free, highly effficient OLTP processing
  - SciDB – Array-processing for multidimensional datasets
- ACM Turing Award 2014
  - Many significant contributions to the DBMS field
  - Successful both in acadmia and commercially

Vision: Enterprise Operational Analytics Data Management Systems

- Vision: One Combined DBMS for OLTP and OLAP Workloads
- SIGMOD 2009 Keynote: “Information at your fingertips”
- Foundation: In-Memory Columnar Database
- Goals: Flexible Reporting, Lower TCO, No ETL

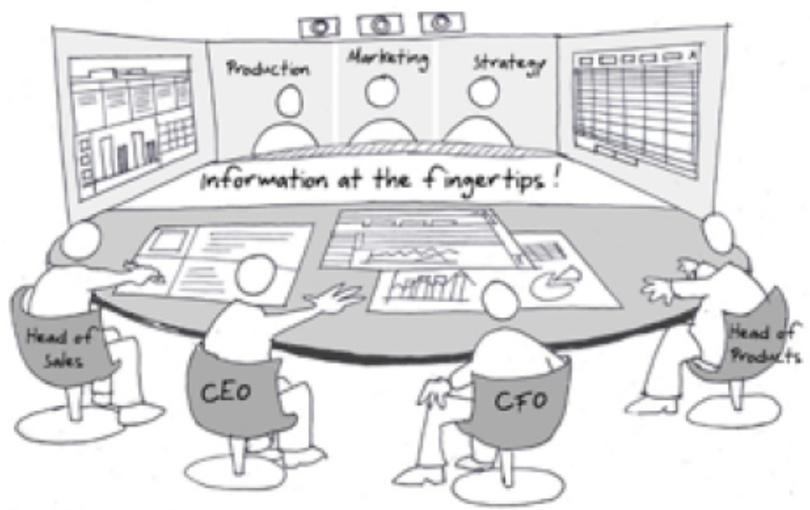

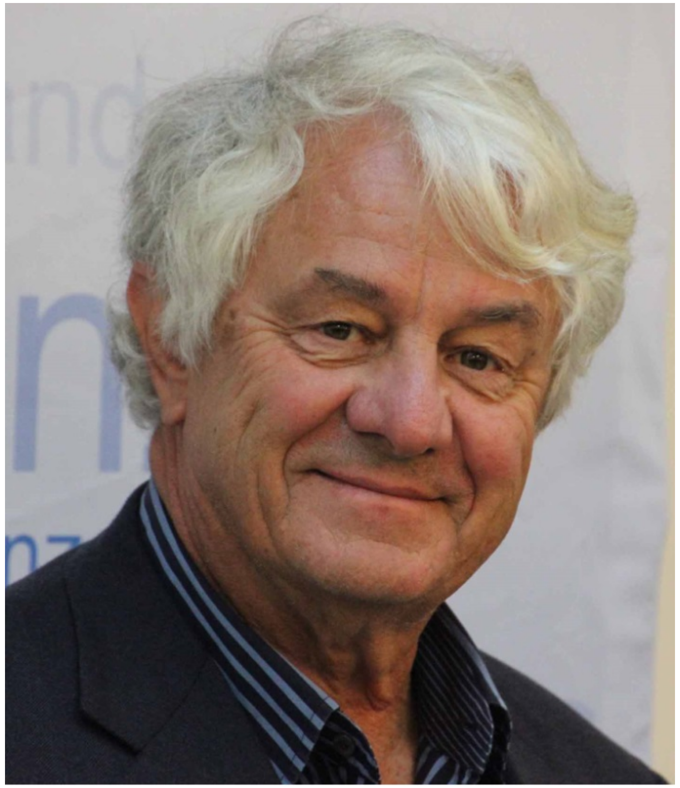

Agenda

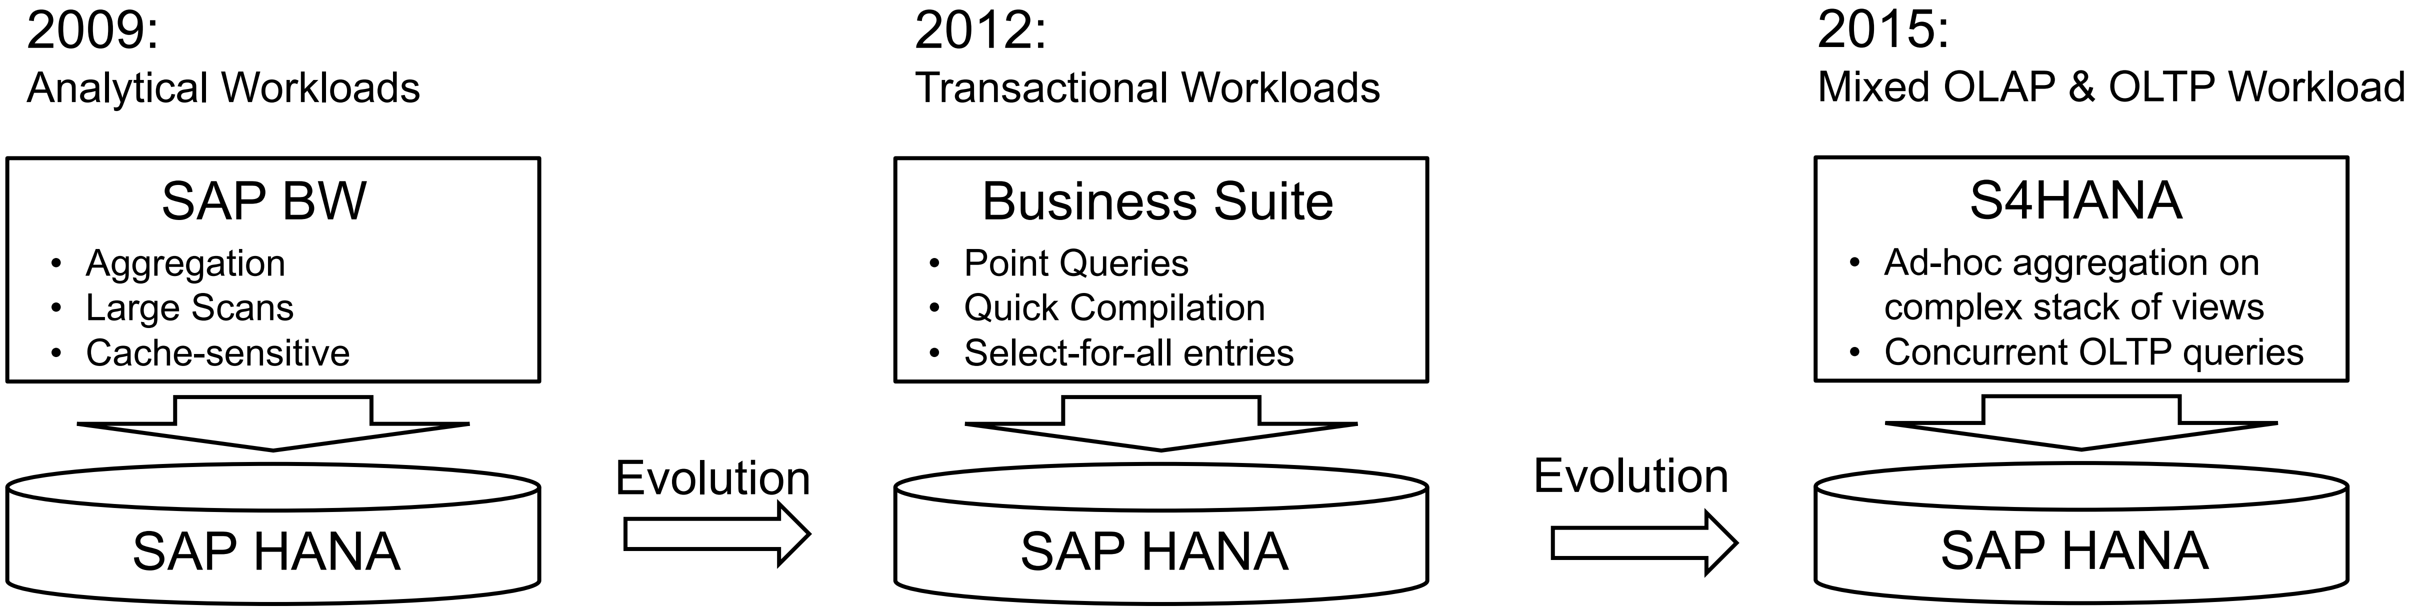

2009:
Analytical Workloads
SAP BW
• Aggregation
• Large Scans
• Cache-sensitive
SAP HANA
Evolution
2012:
Transactional Workloads
Business Suite
• Point Queries
• Quick Compilation
• Select-for-all entries
SAP HANA
Evolution
2015:
Mixed OLAP & OLTP Workload
S4HANA
• Ad-hoc aggregation on complex stack of views
• Concurrent OLTP queries
SAP HANA

SAP HANA for Analytical Scenarios

Parallelization at all levels

- Inter-query
- Intra-query & inter-operator
- Intra-operator
- Instruction-level via SIMD

Scan-friendly columnar data layout

- Read-optimized compressed storage for main fragment
- Update-friendly delta fragment
- Focus was on bulk-loading

Specialized Engines, e.g.

- Text engine
- Planning engine

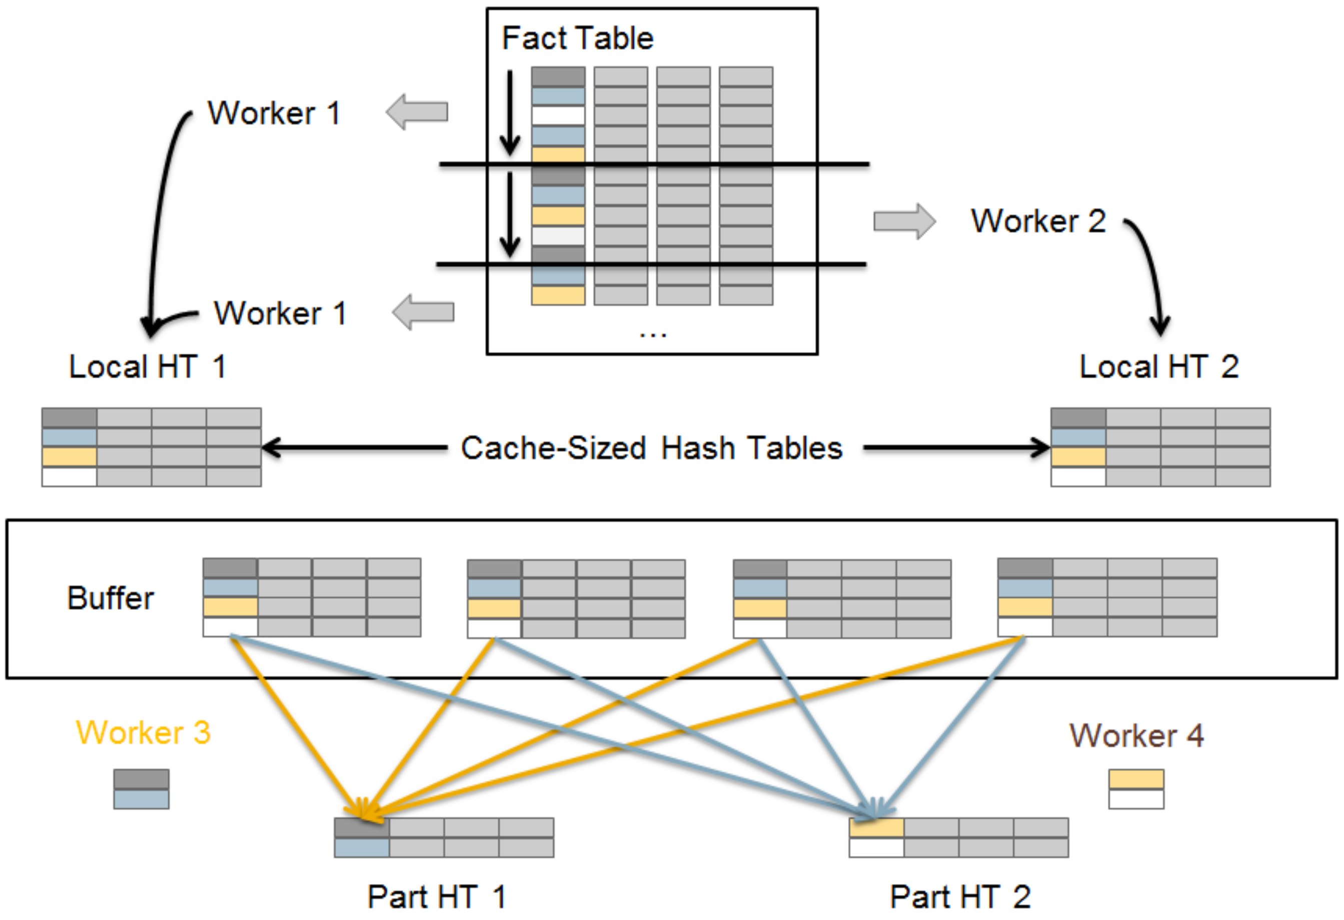

OLTP Optimizations in SAP HANA

Challenge: Run enterprise-scale OLTP workloads on a column store

Cannot afford costs of duplicate data storage (row and columnar)

Need to scale to more than 50’000 queries/second

Efficient support for short-running queries (sub-millisecond range)

Efficient support of update statements

➔ **Requirements very different from initial design space in analytics**

OLTP Optimizations in SAP HANA

Plan cache and OLTP fast path in query optimizer

OLTP-friendly MVCC snapshot handling

Efficient table latching and usage of advanced synchronization primitives, e.g. using hardware transactional memory

Application server / DBMS co-design for efficient data transfer

The Importance of the Plan Cache

Prepared statements

Parameterized statements are optimized at first execution

Plan cache hit ratios >= 99% are critical

Statements are not even parsed

DDLs invalidate plans

Periodic reoptimization to handle updates

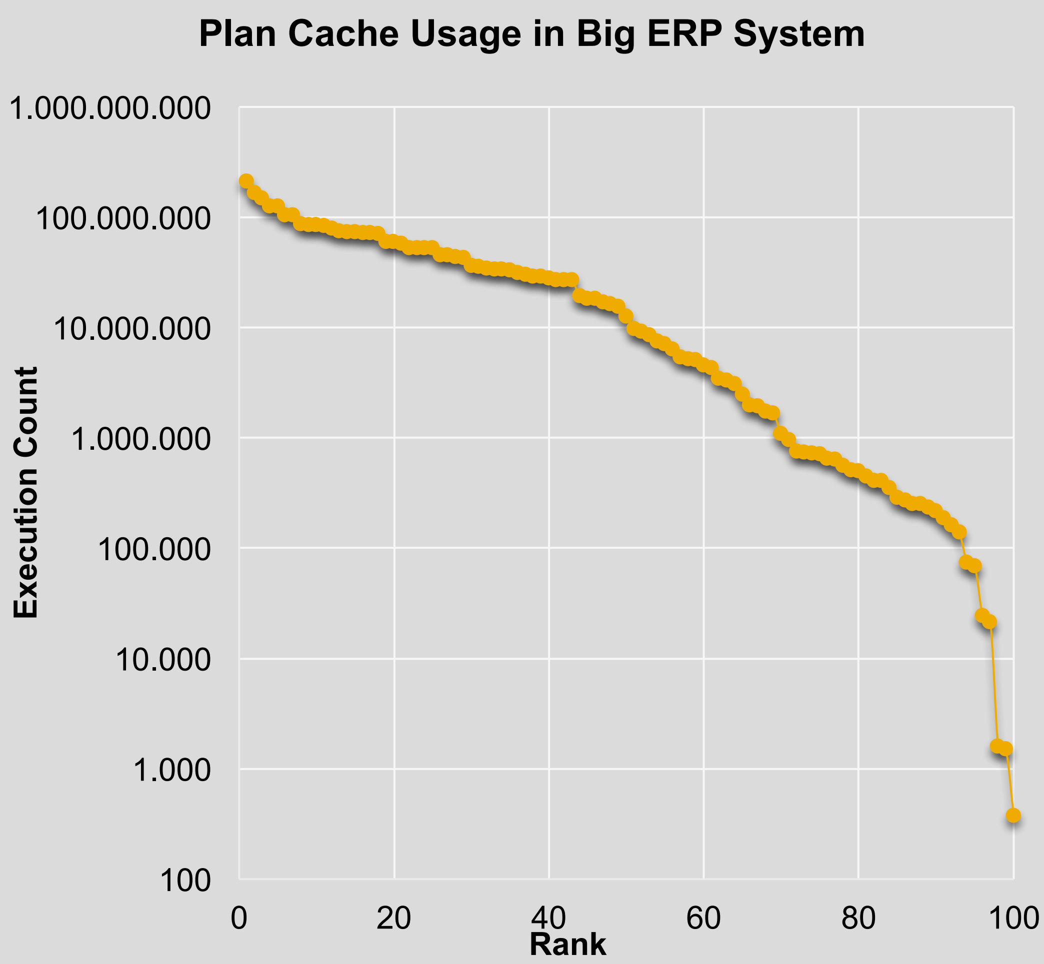

OLTP Fast Path in Query Optimizer

Simple queries, e.g. primary key access single table

Bypass most phases of query optimizer

Single-threaded execution to avoid context switches and to minimize cache misses

Parallelization normally already done in application server

OLTP Performance Improvement over Time

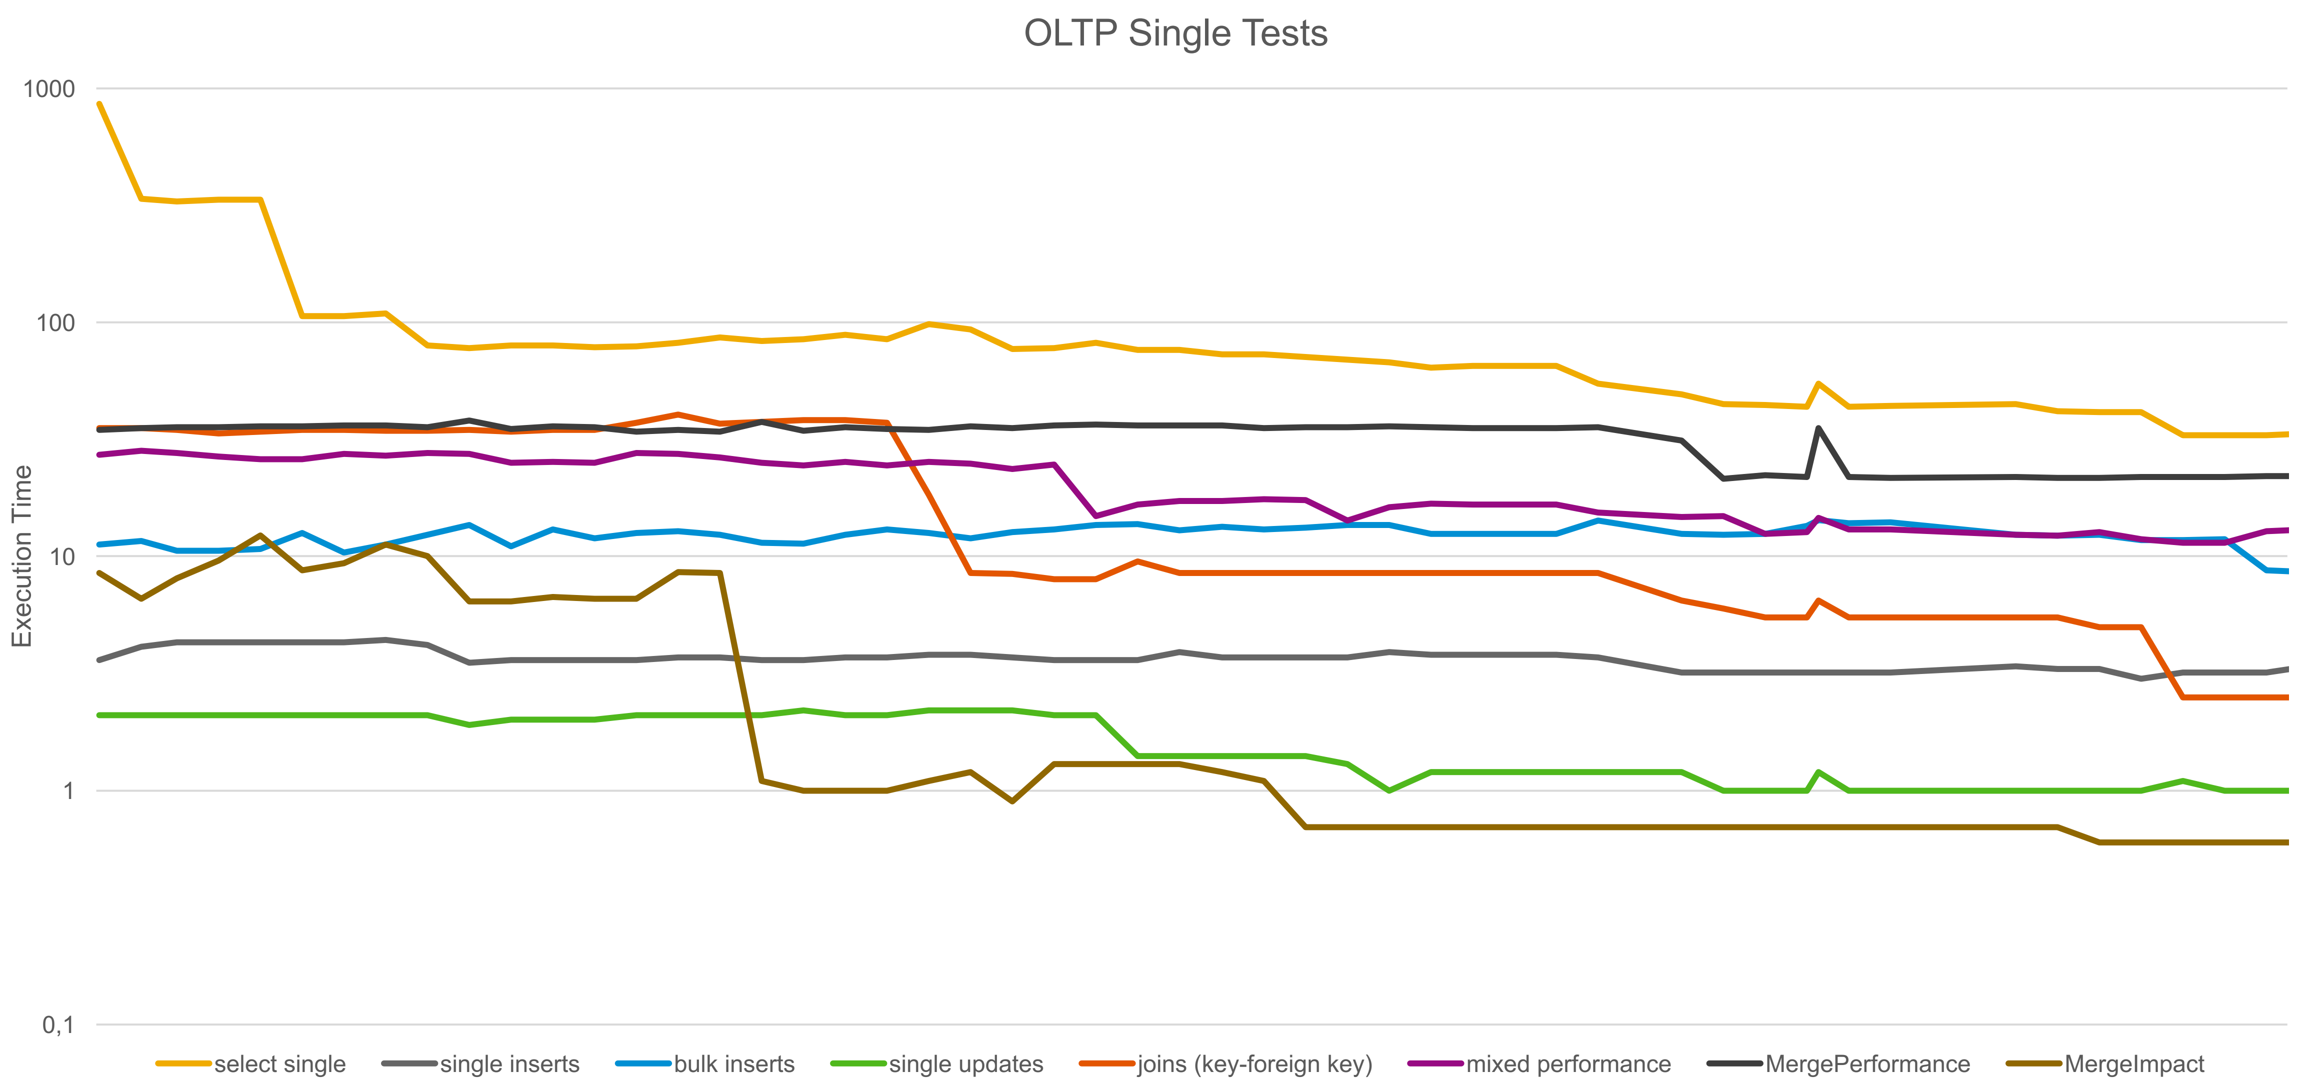

The Vision: Mixed OLAP and OLTP Workloads

Enables novel applications like

- Digital board room or
- Year- / quarter-end closing

result in

- Very complex analytical queries on transactional data
- Analytical query may consume a lot of resources

Challenge: Resource Management of Mixed Workloads

OLAP queries consume CPUs, e.g. 8 sec on 120 CPU cores for retail scenario

Delays on OLTP workloads might lead to queuing effects and must be avoided

Solution: Reporting queries can yield to OLTP queries, but must not be delayed substantially

Tailored resource allocation for different workloads

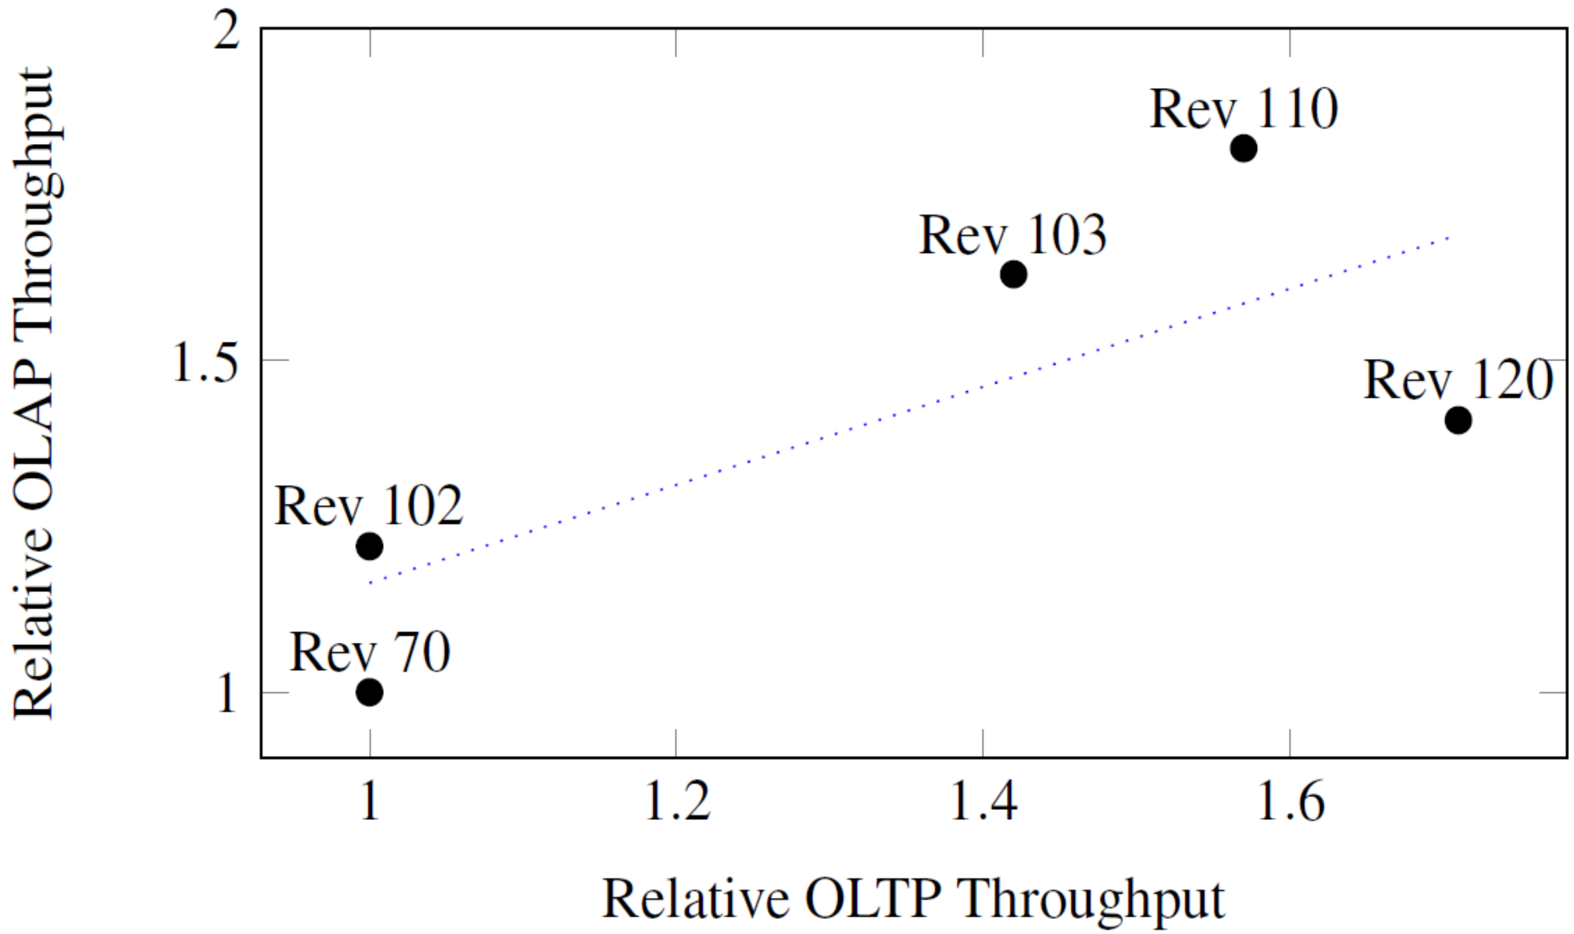

Challenge: High Complexity of Analytics on Normalized Database Schemas

Complex hierarchy of views to expose the normalized database schema for reporting

Solution: result caches for complex views

Research challenge: Analytical model of query complexity

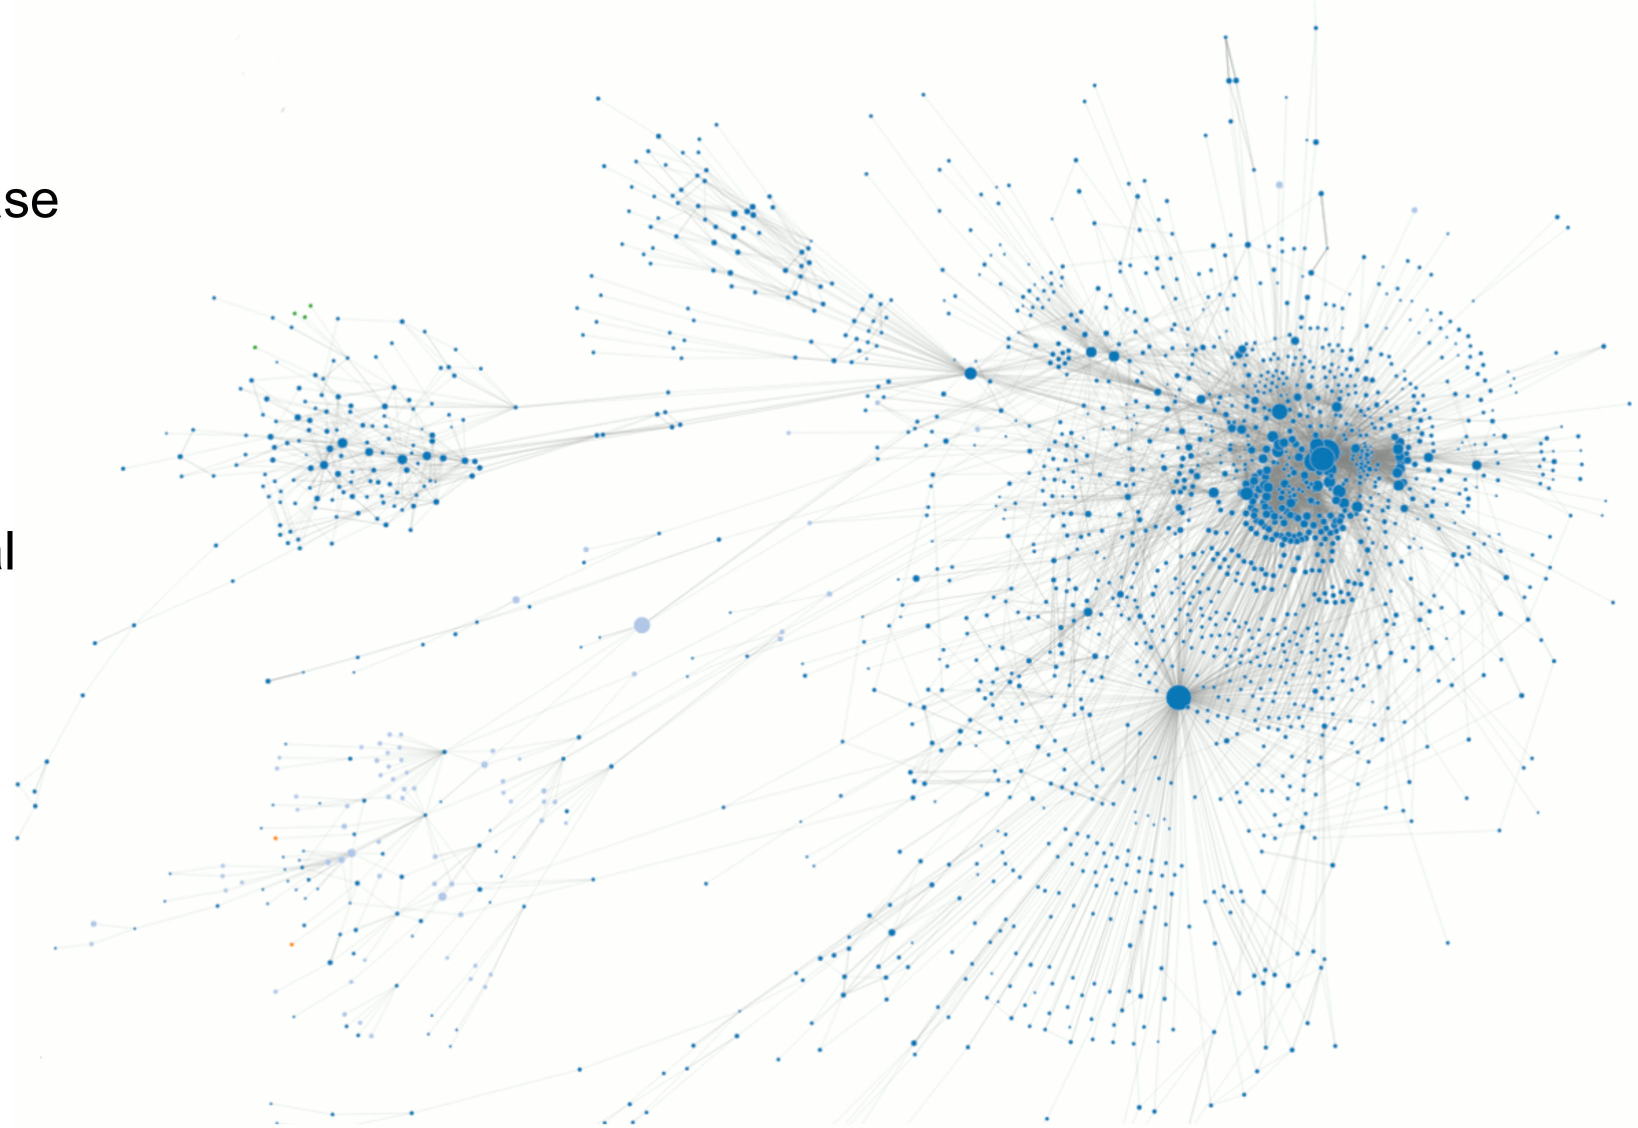

Conclusion: Many Open Questions

Robustness for mixed OLAP and OLTP workload:

- Predicable and low response times
- Dynamic resource assignment
- Good resource utilization

Analytics will become more demanding

- Integrated reporting in OLTP applications
- In-database machine learning

Growing data sizes

- Warm and cold data
- Data aging
- Byte-addressable NVM vs. page-oriented disk

Thank you.

Contact information:

Norman May
SAP SE
norman.may@sap.com

Run Simple