HPI Hasso Plattner Institut

IT Systems Engineering | Universität Potsdam

In-Memory Data Management and Enterprise Applications

Dr. Alexander Zeier

Deputy Chair of „Enterprise Platform and Integration Concepts" of Prof. Hasso Plattner & Visiting Professor at MIT, Cambridge,

Hasso Plattner Institute

GI/ITG-Fachgruppe der Betriebssysteme
Herbsttreffen am 10. November, 2011

Enterprise Platform and Integration Concepts Group

2

Prof. Dr. h.c. Hasso Plattner and Dr. Alexander Zeier

- Research focuses on the technical aspects of enterprise software and design of complex applications
  - In-Memory Data Management for Enterprise Applications
  - Human-Centered Software Design and Engineering
  - Maintenance and Evolution of Service-Oriented Enterprise Software
  - Integration of RFID Technology in Enterprise Platforms
- Industry cooperations with
  - SAP, Siemens, Audi, EADS
- Research cooperations with
  - Stanford, MIT, Berkeley

Partner of Stanford Center for Design Research | Partner of MIT in Supply Chain Innovation and CSAIL | Partner at UC Berkeley RAD / AMP Lab

---

Agenda

3

- **In-Memory Data Management**
- In-Memory Applications
- In-Memory Databases and Co-Processing

In-Memory Technology Enables Combining OLTP and OLAP in Real-Time

4

- Data-centric architecture: In-Memory database serves as **single source** of truth for ERP data
- Architecture based on 4 distinct pillars
  - Multi-Core computing
  - In-Memory
  - Column and Row Store
  - Insert-Only
- Enables informed management decisions based on up-to-the-moment data through real-time combination of
  - Transactional applications
  - Analytical applications

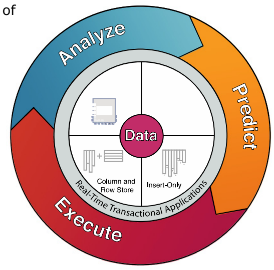

Enterprise Performance In-Memory Circle (EPIC)

… and the Combination with Software Innovations Enables its Full Potential

5

HW Technology Innovations

Multi-Core Architecture (8 x 10core CPU per blade)

Massive parallel scaling across many blades

One blade ~$50.000 = 1 Enterprise Class Server

64bit address space – 2TB in current servers

100GB/s data throughput

Dramatic decline in price/performance

SW Technology Innovation

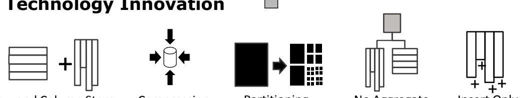

Row and Column Store | Compression | Partitioning | No Aggregate Tables | Insert Only

---

Enterprise Applications Have a Specific Database Footprint

6

- **Today's Enterprise Applications**
  - Complex processes
  - Increased data set (but real-world events driven)
  - Separated into transactional (OLTP) and analytical (OLAP) applications
- **Enterprise Data Management**
  - Wide schemas
  - Sparse data with limited domain
  - Workload consists of complex analytical queries
  - Workload characteristics
    - Set processing
    - Read access
    - Insert operations instead of updates

Many Columns are not Used Even Once in Enterprise Standard Software

7

Results from analyzing financials

Distinct values in accounting document headers

(99 attributes)

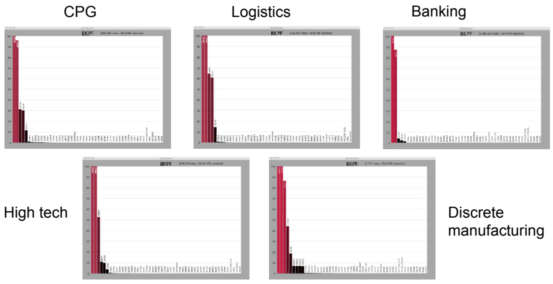

Enterprise Workloads are Read-Mostly and Often Process Ranges of Data

8

- Workload in enterprise applications constitutes of
  - Mainly read queries (OLTP 83%, OLAP 94%)
  - Many queries access large sets of data

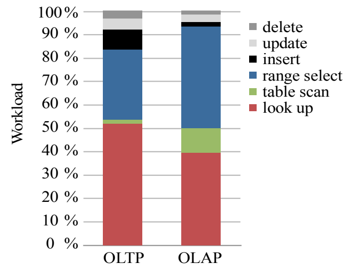

Two Different Principles of Physical Data Storage: Row vs. Column Store

9

Document Number	Document Date	Sold-To Party	Order Value	Status	Sales Organization	...
95769214	2009-10-01	584	10.24	CLOSED	Germany Frankfurt	...
95769215	2009-10-01	1215	124.35	CLOSED	Germany Berlin	...
95779216	2009-10-21	584	47.11	OPEN	Germany Berlin	...
95779217	2009-10-21	454	21.20	OPEN	Germany Frankfurt	...

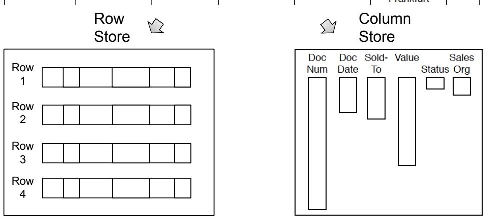

Row-, Column-, and Text-Store

10

Row Store

Small tables
Frequent updates
Materialized aggregates

Column Store

Large tables
Rare updates
Dynamic aggregates

Text

Crawler
Join structured & unstructured data

Transactional Data

Direct access to tuples
Blade-local transactions
Status updates
Active / passive

Historical Data

Sequential access
No updates

Enterprise Data Stored in Columns Can be Compressed by a Factor of up to 50

11

Document Number	Document Date	Sold-To Party	Order Value
95769214	2009-10-01	584	10.24
95769215	2009-10-01	1215	124.35
95779216	2009-10-21	584	47.11
95779217	2009-10-21	454	21.20

Dictionaries

Document Number

0	95769214
1	95769215
2	95779216
3	95779217

Order Value

0	10.24
1	21.20
2	47.11
3	124.35

Document Date

0	2009-10-01
1	2009-10-21

Sold-To Party

0	454
1	584
2	1215

Document Number	Document Date	Sold-To Party	Order Value
0	0	1	0
1	0	2	3
2	1	1	2
3	1	0	1

- Typical compression factor for enterprise software 10
- In financial applications up to 50

New Processor Architectures Support Parallelization of Database Operations

12

- **Set processing** – scan is dominant pattern in enterprise apps
- **Sequential** scans allow best bandwidth utilization between CPU cores and memory
- **Independence** of tuples within columns allows easy partitioning and therefore parallel processing (see Hennessy [1])
- Increased memory bandwidth in current and next generation CPUs allows even **faster memory scans**. Current Nehalem architecture allows multiple memory channels, with an increased combined bandwidth.
- No more materialized views and aggregates: everything is calculated **on-the-fly**

[1] John L. Hennessy, David A. Patterson:
Computer Architecture: A Quantitative Approach

Business Function Libraries Encapsulate Application Logic on Database Level

13

- New enterprise data management requires **rethinking** of how application logic is written
- The objective is to bring application logic closer to the storage layer using stored procedures / business library functions
- Stored procedures can be shared among applications
  - Identify common application logic and implement as stored procedures
  - Reduce application code by leveraging stored procedures

Summary – Numerous Innovations Enable the Full Potential of In-Memory Computing

14

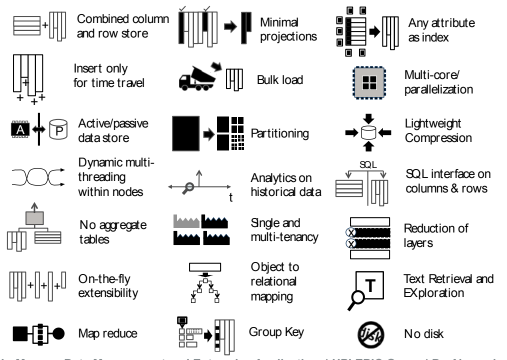

Agenda

15

- In-Memory Data Management
- **In-Memory Applications**
  - **HANA Oncolyzer**
  - ATP
  - Demand Planning
  - GORFID
- In-Memory Databases and Co-Processing

HANA Oncolyzer

16

- International research initiative for exchanging relevant tumor data to improve treatment
- In-Memory Technology as key-enabler for real-time analysis of tumor data in seconds instead of hours
- In-Memory enables join of third-party data to improve analysis results
- Information available at your fingertips: In-Memory Technology on mobile devices (iPad)
- Interdisciplinary cooperation between medical doctors, researchers, and software engineers

Agenda

17

- In-Memory Data Management
- **In-Memory Applications**
  - HANA Oncolyzer
  - **ATP**
  - Demand Planning
  - GORFID
- In-Memory Databases and Co-Processing

In-Memory Technology has been applied to the Available-to-Promise Check

18

- Goal
  Analyze and validate the potential of **in-memory** and **highly parallel data processing** for Available-to-Promise (ATP)
- Challenges
  - Dynamic aggregation
  - Instant rescheduling in minutes vs. nightly batch runs
  - Real-time and historical analytics
- Outcome
  - Real-time ATP checks
  - Ad-hoc rescheduling
  - No materialized aggregates

In-Memory Available-to-Promise

19

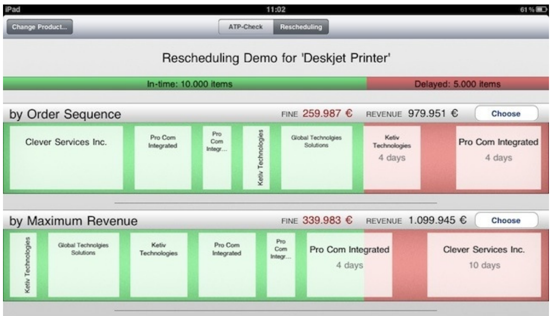

Agenda

20

- In-Memory Data Management
- **In-Memory Applications**
  - HANA Oncolyzer
  - ATP
  - **Demand Planning**
  - GORFID
- In-Memory Databases and Co-Processing

Demand Planning

21

- Flexible analysis of demand planning data
- Zooming to choose granularity
- Filter by certain products or customers
- Browse through time spans
- Combination of location-based geo data with planning data in an in-memory database
- External factors such as the temperature, or the level of cloudiness can be overlaid to incorporate them in planning decisions

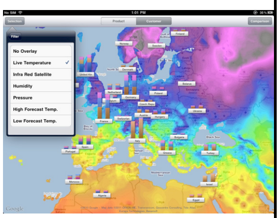

Agenda

22

- In-Memory Data Management
- **In-Memory Applications**
  - HANA Oncolyzer
  - ATP
  - Demand Planning
  - **GORFID**
- In-Memory Databases and Co-Processing

GORFID

23

- HANA for Streaming Data Processing
- Use Case: In-Memory RFID Data Management
- Evaluation of SAP OER
- Prototypical implementation of:
  - RFID Read Event Repository on HANA
  - Discovery Service on HANA (**10 Billion data records** with ca. 3 seconds response time)
  - Frontends for iPhone, iPad2

Key Findings:

- HANA is suited for streaming data (using **bulk inserts**)
- Analytics on streaming data is now possible

Agenda

24

- In-Memory Data Management
- In-Memory Applications
- **In-Memory Databases and Co-Processing**

Why “Co-Processing”

25

- CPUs are seem to be optimal with regards to query execution on sequential data streams
  - Multi-core systems allow high parallelism, but how to handle long running tasks?
- In memory databases use compression to reduce memory footprint and increase scan performance
  - Compression requires updating the compressed data, this has to take minimal time and resource utilization
    - Merge Process: Merge two dictionaries by identifying duplicates in both and generate a mapping from the old to the new dictionary. Apply mapping to all values stored in the column.
  - Idea: Offloading recompression to GPU

GPU Motivation / Trade-Offs

26

- GPUs offer up to **two orders** of magnitude more cores than a CPU, possible speed-up through parallelization
- However, data needs to be **transferred** over PCI-Express from main memory to the co-processor
  - Requirement: improved implementation needs to be faster, more scalable and cover the **amortized** cost for data transfer

GPU Merge Initial Findings

27

- GPU Dictionary Merge
  - Merging Dictionaries is mostly duplicate removal
  - **27 x** performance improvement on duplicate removal
  - **40%** improvement on merging two dictionaries
- Trade-offs
  - Limited bandwidth for data transfer
  - Limited amount of global memory (10s of GB) compared to systems main memory (1000s of GB)
- Future Directions
  - **Combine** CPU / GPU Merge
  - Hide latencies for data transfers

All Findings are Summarized in the Book In-Memory Data Management

28

This book is the culmination of five years worth of in-memory research

- PART I – An Inflection Point for Enterprise Applications
  - Overview of our vision of how in-memory technology will change enterprise applications
- PART II – SanssouciDB – A Single Source of Truth through In-Memory
  - Technical foundations of in-memory data management
  - In-depth description of how we intend to realize our vision
- PART III – How In-Memory Changes the Game
  - Resulting implications on the development and capabilities of enterprise applications

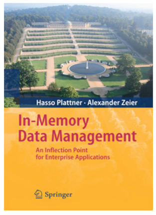

-> Central Parts of SanssouciDB productized as SAP Product HANA - Available since June 2011

29

Thank you for your interest!

Dr. Alexander Zeier

Deputy Chair „Enterprise Platform and Integration Concepts“ of Prof. Hasso Plattner
Hasso Plattner Institute at University Potsdam
& Visiting Professor at MIT, Cambridge

http://epic.hpi.uni-potsdam.de
zeier@hpi.uni-potsdam.de