



Massachusetts
Institute of
Technology



SAP HANA: In-Memory Data Management for Enterprise Applications

Dr. Alexander Zeier

Massachusetts Institute of Technology (MIT)
Visiting Professor

March 23rd 2012

SAP Academic Conference Americas, San Antonio



The Vision by Prof. Hasso Plattner


- Transactional (OLTP) and analytical (OLAP) data processing has to be on **one system again**
- Enterprise applications have to reflect latest developments in:
 - Hardware, such as:
 - Multi-core processors
 - Huge Main Memory
 - Data management, such as:
 - Column-oriented storage
 - Light-weight compression





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

- 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)




In-Memory Data Management

Advances in Hardware

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

Parallel scaling across 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




Advances in Software



Row and
Column Store




Compression




Partitioning




No Aggregate
Tables




Insert Only




On-the-fly
extensibility




Combined column and row store



Minimal projections




Any attribute as index




Insert only for time travel




Bulk load




Multi-core/parallelization




Active/passive data store




Partitioning




Lightweight Compression




Dynamic multi-threading within nodes




Analytics on historical data




SQL interface on columns & rows




No aggregate tables




Single and multi-tenancy




Reduction of layers




On-the-fly extensibility




Object to relational mapping



Text Retrieval and EXploration



Map reduce



Group Key



No disk






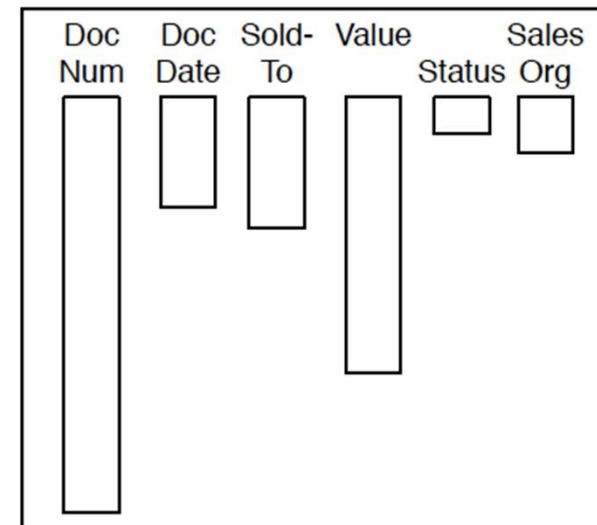
Two Different Principles of Physical Data Storage: Row vs. Column Store

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 Store 



 Column Store






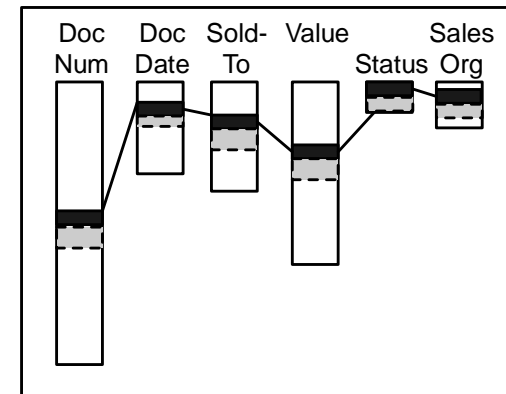
Accessing Enterprise Data

```
SELECT *  
FROM Sales Orders  
WHERE Document Number = '95779216'  
(OLTP-style query)
```



Row Store



Column Store



```
SELECT SUM (Order Value)  
FROM Sales Orders  
WHERE Document Date > 2009-01-20  
(OLAP-style query)
```



Dictionary Compression

- ❑ Reduces I/O operations to main memory (bottleneck)
- ❑ Operations directly on compressed data

8

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		Order Value	
0	95769214	0	10.24
1	95769215	1	21.20
2	95779216	2	47.11
3	95779217	3	124.35

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



Document Date		Sold-To Party	
0	2009-10-01	0	454
1	2009-10-21	1	584

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

8



Table Characteristics

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

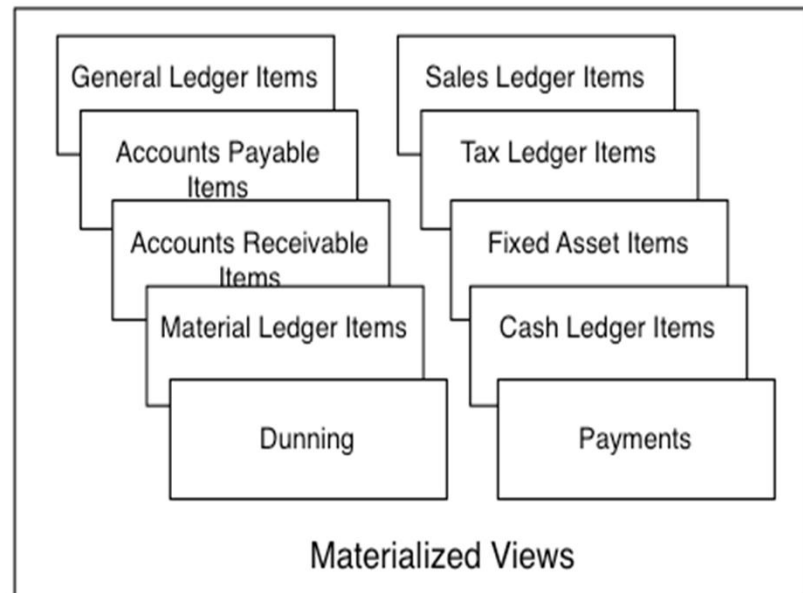
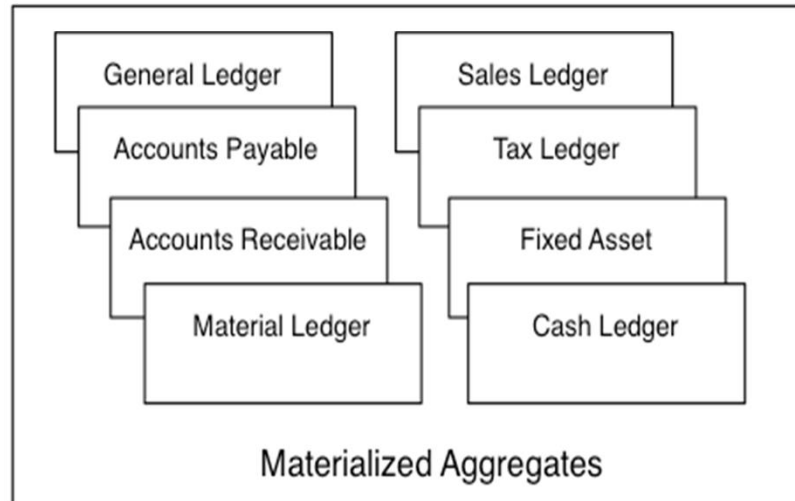

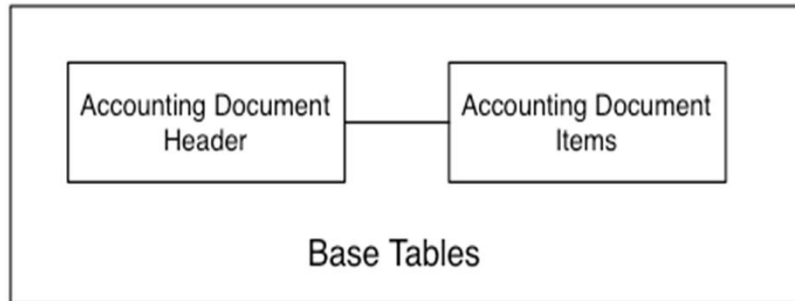
Sequential access
No updates



Innovative In-Memory / HANA Applications



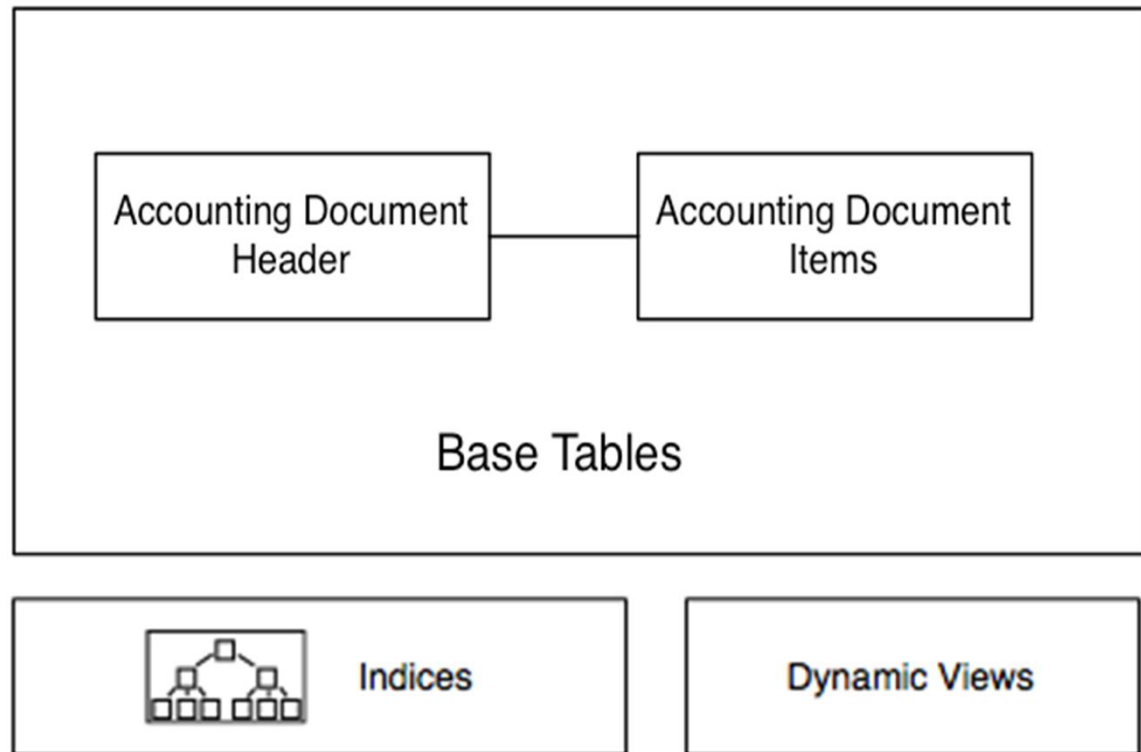
Nowadays Financials





Simplified Financials (Target)

Only base tables,
algorithms, and some
indices





Customer Study 1: Dunning Run in < 1s?

- Dunning run determines all open and due invoices
- Customer defined queries on 250M records
- Current system: 20 min
- New logic: **1.5 sec**
 - In-memory column store
 - Parallelized stored procedures
 - Simplified Financials



Dunning Application

The screenshot shows a mobile application interface for debtors and dunning items. The top bar displays the status 'Telekom.de', signal strength, and battery level. The main content is divided into two sections: a list of debtors on the left and a detailed view of a debtor on the right.

Debtors List:

Company Name	Outstanding	Lost Interest
American Axle & Manufacturing Holdings, Inc.	\$1,270,764	\$41,183
Fidelity National Corporation	\$386,625	\$12,535
Ault, Inc.	\$319,156	\$10,054
Powerwave Technologies, Inc.	\$308,262	\$9,821
Mystic Financial, Inc.	\$225,217	\$7,293
Wavecom S.A.	\$219,259	\$7,079
Value City Department Stores, Inc.		
Total Amount		
Total Outstanding	\$6,742,899.21	
Total Lost Interest		\$217,939.47

Customer Details (9 of 10): Johnson Controls, Inc.

Address: Emerson Street 720
77215 Houston TX

Contact: Arminda Lank
- CFO -
(555) 325-4909179

Outstanding: \$202,804.44 | Lost Interest: \$6,552.68

Top Dunning Items:

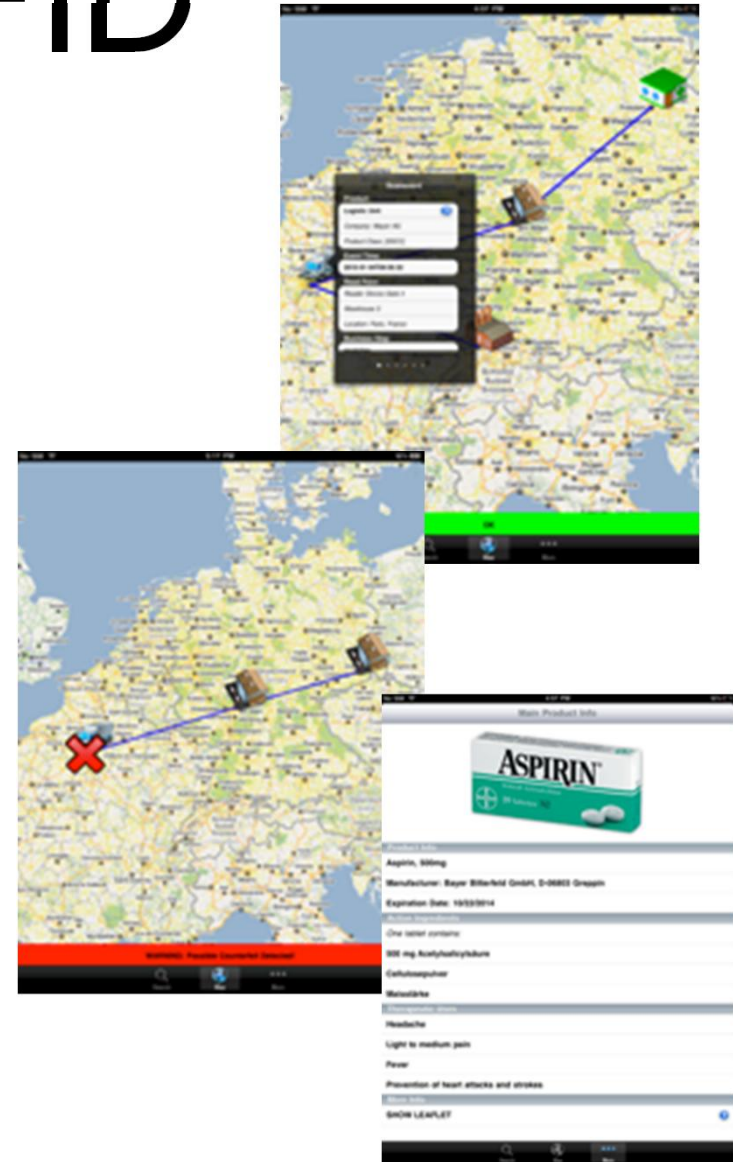
Due Date	Days Overdue	Amount	Lost Interest
2009-09-19	240	\$3,465.60	\$113.94
2009-09-19	240	\$3,184.72	\$104.70
2009-09-19	240	\$1,478.40	\$48.60
2009-09-20	239	\$3,806.06	\$124.61
2009-09-20	239	\$3,592.68	\$117.62
2009-09-20	239	\$1,478.40	\$48.40
2009-09-21	238	\$12,026.35	\$392.09



GORFID

- Tracing pharmaceutical packages in Europe
- 15 bn packages / 35 bn read events per year

Prototype with
12 billions
records with
response time:
23 ms





HANA Oncolyzer

- Medical doctors have all patient data at hand to apply personalized medicine
- Medical researchers perform real-time analysis to define cohorts for clinical studies
- International research initiative for exchanging relevant tumor data started at World Health Summit 2011 in Berlin
- In-Memory Technology as
 - key-enabler for real-time analysis
 - provider for information at your fingertips (iPad)





HANA Oncolyzer - combining Structured and Unstructured Data





HANA Oncolyzer was presented on CeBIT 2012 to Germany's Chancellor Angela Merkel as SAP's Innovation 2012





First Results of Customers using SAP HANA

- 1,000x Faster: Many (Dunning, Aging, ...)
- 10,000x Faster: NongFu Spring, Essar Group, SAP IT, Cornell, Charmer Sunbelt
- 100,000X Faster: YodoBashi, MKI

OR

- 24+ Hours to 3.8S: Food and Beverage / Distribution - Logistics
- 15+ Hours To 4.8S: Project Management / Services, Profitability, Performance
- 30 Days to 28S: Manufacturing – Order to Cash
- 3 Days to 2s: Retail / Insurance – Incentives





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

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 – 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

-> Book launched at Cebit 2011, SAP Product HANA is available since June 2011.

-> New extended Book Edition “In-Memory Data Management - Technology and Applications ” focusing on Application Development will be available for Sapphire May 2012.





In-Memory/HANA Drives Worldwide Innovation Book Launch at CeBIT 2011 with Vice-President of the *European Commission* Neelie Kroes





Massachusetts
Institute of
Technology



SAP and HPI win the German Innovation Award 2012 for SAP HANA!

This year's winners were announced on March 16, 2012 in Munich, Germany.

Please feel free to contact me:

Dr. Alexander Zeier

Massachusetts Institute of Technology (MIT)

Visiting Professor

Executive Director MIT Forum for SC Innovation

Email: zeier@mit.edu

Website with list of over 150 Publications: <http://zeier.mit.edu>