

In-memory Databases in Business Information Systems

DOI 10.1007/s12599-011-0188-y

The Authors

Prof. Dr. Peter Loos (✉)

IWi at DFKI
Saarland University
Campus D3 2
66123 Saarbrücken
Germany
loos@iwi.uni-sb.de

Jens Lechtenböcker

Gottfried Vossen

Alexander Zeier

Jens Krüger

Jürgen Müller

Wolfgang Lehner

Donald Kossmann

Benjamin Fabian

Oliver Günther

Robert Winter

lechten@helios.uni-muenster.de

vossen@uni-muenster.de

Published online: 2011-11-05

This article is also available in German in print and via <http://www.wirtschaftsinformatik.de>: Loos P, Lechtenböcker J, Vossen G, Zeier A, Krüger J, Müller J, Lehner W, Kossmann D, Fabian B, Günther O, Winter R (2011) In-Memory-Datenmanagement in betrieblichen Anwendungssystemen. WIRTSCHAFTSINFORMATIK. doi: 10.1007/s11576-011-0296-9.

© Gabler Verlag 2011

1 Introduction

In-memory databases are developed to keep the entire data in main memory. Compared to traditional database systems, read access is now much faster since no I/O access to a hard drive is required. In terms of write access, mechanisms are available which provide data persistence and thus secure transactions. In-memory databases have been available for a while and have proven to be suitable for particular use cases. With in-

creasing storage density of DRAM modules, hardware systems capable of storing very large amounts of data have become affordable. In this context the question arises whether in-memory databases are suitable for business information system applications. Hasso Plattner, who developed the HANA in-memory database, is a trailblazer for this approach. He sees a lot of potential for novel concepts concerning the development of business information systems. One example is to conduct transactions and analytics in parallel and on the same database, i.e. a division into operational database systems and data warehouse systems is no longer necessary (Plattner and Zeier 2011). However, there are also voices against this approach. Larry Ellison described the idea of business information systems based on in-memory database as “wacko,” without actually making a case for his statement (cf. Bube 2010). Stonebraker (2011) sees a future for in-memory databases for business information systems but considers the division of OLTP and OLAP applications as reasonable.

Therefore, this discussion deals with the question of whether in-memory databases as a basic data management technology can sustainably influence the conception and development of business information system or will remain a niche application. The contributors were invited to address the following research questions (among others):

- What are the potentials of in-memory databases for business information systems?
- What are the consequences for OLTP and OLAP applications?
- Will there be novel application concepts for business information systems?

The following researchers accepted the invitation (in alphabetic order):

- Dr. Benjamin Fabian and Prof. Dr. Oliver Günther, Humboldt-Universität zu Berlin
- Prof. Dr. Donald Kossmann, ETH Zürich
- Dr. Jens Lechtenböcker and Prof. Dr. Gottfried Vossen, Münster University

- Prof. Dr. Wolfgang Lehner, TU Dresden

- Prof. Dr. Robert Winter, St. Gallen University

- Dr. Alexander Zeier with Jens Krüger and Jürgen Müller, Potsdam University

Lechtenböcker and Vossen discuss the development and the state of the art of in-memory and column-store technology. In their evaluation they stress the potentials of in-memory technology for energy management (cf. Loos et al. 2011) and Cloud Computing.

Zeier et al. argue that the main advantage of modern business information systems is their ability to integrate transactional and analytical processing. They see a general trend towards this mixed processing mode (referred to as OLXP). In-memory technology supports this integration and will render the architectural separation of transactional systems and management information systems unnecessary in the future. The new database technology also greatly facilitates the integration of simulation and optimization techniques into business information systems.

Lehner assumes that the revolutionary development of system technology will have a great impact on future structuring, modeling, and programming techniques for business information systems. One consequence will be a general shift from control-flow-driven to data-flow-driven architectures. It is also likely that the requirement for ubiquitously available data will be abandoned and a “need-to-know” principle will establish itself in certain areas.

Kossman identifies two phases in which in-memory technology will influence business information systems. The first phase is a simplification phase which is caused by a separation of data and application layers of information systems. In a second phase, however, complexity will increase since the optimization of memory hierarchies, such as the interplay between memory and cache, will also have consequences for application developers.

Fabian and Günther stress that in-memory databases have already proven

their potential in certain practical applications and they see a great potential for these systems in business intelligence software. They also point out that higher costs for hardware and software procurement as well as for re-engineering of organizational structures are to be expected. It will also be necessary to develop special licensing models.

Winter states that the potential integration of OLTP and OLAP databases by means of in-memory technology will have far reaching consequences for the architecture and the processes of the entire business information logistics. If the promises are met, Winter believes that the next generation of research questions for information systems will develop after a phase of data orientation and a phase of process orientation into a phase of analysis and decision orientation.

Despite open questions regarding maturity and feasibility the authors agree on the potential of in-memory technology databases and expect far reaching consequences for business information systems.

If you would like to comment on this topic or another article of the journal *Business & Information Systems Engineering*, please send your contribution (max. 2 DIN A4 pages) to the editor-in-chief, Prof. Hans Ulrich Buhl, University of Augsburg, Hans-Ulrich.Buhl@wiwi.uni-augsburg.de.

Prof. Dr. Peter Loos
IWi at DFKI
Saarland University, Saarbrücken

2 In-Memory Database Technology: Challenges and Opportunities

2.1 Introduction

Many people consider in-memory database technology as a major breakthrough for the future development of business information systems, among them Plattner and Zeier (2011) who, in their new book, claim that “we have now reached a new inflection point ... in-memory computing is changing the way businesses are run.” Others, including Ousterhout et al. (2009), simply see the necessity of moving from disk storage to RAM storage to be able to scale to the needs of present-day Web applications as well as yet unknown upcoming high-performance applications.

In a nutshell, an in-memory database system, previously also termed a main-memory database system, is a database system that primarily relies on primary (instead of secondary) memory for data storage. Since database systems are traditionally based on disk storage and only keep small portions of the database in a buffer in main memory, the in-memory approach immediately eliminates the I/O bottleneck, i.e., the bottleneck typically existing between a (fast) RAM and a (slow) disk. As a consequence, a considerable performance gain can be expected. However, the obvious price to pay is the need to rethink most techniques a database system comes with, in particular regarding transactional guarantees, data organization and access, and query optimization. Indeed, transaction recovery (needed for an atomicity guarantee) is commonly based on keeping a log on disk; data organization relies on index structures which allow fast access to disk pages; query optimization typically tries to reduce the I/O overhead that would result from executing a user-defined query as is. It should be mentioned that the idea of basing a database system on main instead of disk memory is not new; this issue has been investigated in the mid-1980s already (Eich 1987a, 1987b), but then has been abandoned and forgotten for some time.

While other recent developments, such as column stores, have addressed some of the issues mentioned, in-memory database systems attempt to cover them all at once. The premise is that potential performance gains are enormous, and with multi-core processing capabilities as well as main memory sizes in giga- to terabytes, the vision of “real-time computing” for both OLTP (online transaction processing) and OLAP (online analytical processing) as well as for BI (business intelligence) applications may finally become a reality, even if possible obstacles, such as database systems, are considered.

2.2 A Glimpse of Current Systems

The area of in-memory database systems is currently characterized by a lot of R&D activity, with some systems such as solidDB (IBM), TimesTen (Oracle), and SAP HANA already being commercially available. At the heart of this activity is a redesign of the classical database system architecture which has long been advocated by Bernstein et al. (1998) and which

is finally materializing now. We can only give a glimpse of this here.

One of the oldest projects in the field is MonetDB (Boncz et al. 2009), developed at CWI in Amsterdam, which is essentially a column store based on the idea that row-based storage is inferior to column-based storage in many data warehouse and OLAP applications. Consequently, MonetDB features a storage model that is based on vertical fragmentation; in addition, it comes with a CPU-tuned query execution engine and a modular software architecture. When it comes to transaction processing, in-memory systems typically work with snapshots for recording database states, complemented by measures such as replication or the use of non-volatile RAM. The HyPer system (Kemper and Neumann 2011), a main-memory database system under development at the Technical University of Munich, puts an emphasis on transaction processing. HyPer can make transactional guarantees and executes multiple OLAP queries on the same, arbitrarily current and consistent snapshot. It relies on hardware-supported page shadowing that is controlled by the memory management unit of the underlying processor and it processes transactions sequentially on individual database partitions so that locks are no longer needed.

While in-memory systems will most likely not revolutionize end-user software applications, they are considerably different from traditional systems and require a plethora of novel algorithmic approaches. Therefore, it can be expected that they will indeed have a considerable effect on the evolution of standard business software since their performance will make applications such as real-time analytics or on-demand and ad-hoc business intelligence become reality.

2.3 Outlook

With technology becoming more and more reliable in certain areas and with computer hardware being one of them (another being airplanes, which nowadays can cover long distances on just two engines), it is to be expected that in-memory technology, in contrast to what happened in the 1980s, will now spread widely during the next couple of years. As large amounts of main memory (DRAM) typically account for a significant and load-independent amount of energy usage, we expect in-memory

database power consumption to benefit particularly from research into DRAM power management (Hur and Lin 2008). Roughly, the challenge is to take advantage of low power modes of modern DRAMs, e.g., via throttling and scheduling of memory commands to fewer physical devices.

With controlled power consumption, the proliferation of in-memory databases will concur (and ultimately be combined) with cloud technology. Although in-memory database systems are typically not “cloud-aware,” they are not designed with cloud computing in mind, and are in a sense even complementary to the cloud, their combination will make real-time computing in any area not only fast, but also affordable.

Dr. Jens Lechtenbörger
Prof. Dr. Gottfried Vossen
University of Münster

3 Potential of In-Memory Technologies

3.1 Introduction

Considered individually, the technologies used for “in-memory data management” are mainly not new – for example, in-memory databases or column-oriented storages have already been considered in the 1980s. However, the hardware available today enriched with innovative software developments enables whole new areas of application.

Modern operational application systems draw much of their added value from the latency-free integration of data and processes from transactional and analytical areas. Accordingly, in-memory technology offers a data management specialized for enterprise application systems to implement these requirements efficiently.

Enterprise application systems that are based on this type of data management can simultaneously employ the conventional single entity processing as well as set processing with analytical functions of data.

3.2 Impact on the Further Development of Business Management Software

The interaction of developments in hardware and software now enable the flexible ac-

cess of all enterprise data in a matter of sub-seconds. Thus, not only the dream of controlling departments has come true, but also the enterprise application software can be re-thought afresh. Such a radical break may only occur once a decade. In this context, many applications will be newly designed in a way that they can be operated by mobile devices. The trend towards applications with mixed workloads (OLXP) will intensify further, meaning that operational and transactional applications (OLTP) provide more and more analytical functionality and that analytical applications (OLAP) access transactional systems and their data.

Based on insights we have made over the past five years in the area of in-memory data management, we assume that nearly all components of enterprise management software can be raised to a qualitatively new level. This offers a great potential but also implies the effort to adapt existing software to a new form of data management. Currently emerging applications should be adjusted to in-memory databases. Applications already existing should be prioritized and then gradually rewritten. In addition, there can be a re-examination to integrate existing concepts of enterprise application software. Simulation and optimization methods, which so far – if at all – have been carried out in special systems, deserve particular consideration.

Now it is potentially possible to integrate data as well as computing intensive applications into enterprise application software close to the database, for example, in the form of stored procedures. This not only creates new opportunities, but in our opinion also new obligations: enterprise application software and standard software in particular are not exactly known for their user-centricity. It is now the time to change this. In private use, only those applications prevail that are desirable; i.e., those that are wanted and appreciated by the user. If an application does not fulfill its expectations, it is immediately deleted (from the mobile device). Why should this not apply to enterprise applications?

3.3 Consequences for Operational and Analytical Applications

Operational information systems, which form the basis of decision making

in enterprises, traditionally analyze a large data volume and retrieve the data from especially prepared data warehouse cubes. However, what we have seen in practice is that the decision makers prefer to analyze the company’s data in a flexible way in order to enable an informed decision. Since this is often not possible despite the large number of existing data, decisions are made intuitively. This is exactly why we are against a systematic separation of transactional systems and separate management information systems. Rather, planning and controlling systems should work with the same data as do transactional systems. The way towards this aim is already outlined (Plattner and Zeier 2011, pp. 205 ff.).

3.4 What Will Practice Look Like in 5 Years Time?

The change caused by in-memory data management is comparable with the switch from Sony’s walkman to Apple’s iPod. This of course does not happen overnight. Obviously, applications have to be rewritten. What is more important, however, is that users as well as providers move away from the prevailing image of the operational enterprise software, how this software is designed and what is currently not possible. Thinking outside the box is required! This may also be a starting point for Business and Information Systems Engineering which would be able to unify economic requirements and the new technological possibilities.

It will be exciting to see how different software development enterprises will respond to this new technical opportunity. Agile user-centric processes should cause enterprise application software to have a new face in the next five years: mobile, lightweight, flexible, and user-centric.

Dr. Alexander Zeier
Jens Krüger
Jürgen Müller
Hasso-Plattner-Institut für
Softwaresystemtechnik, Potsdam

4 “In-Memory Data Management” – Evolution or Revolution?

Over the past few years, a barely noticeable change has occurred in the field

¹Derived from Jim Gray’s statement that “Tape is Dead, Disk is Tape, Flash is Disk, RAM Locality is King,” 2006. (http://research.microsoft.com/en-us/um/people/gray/talks/Flash_is_Good.ppt).

of computer architecture. Driven by the idea that “disk is tape, RAM is king,”¹ storage hierarchies have shifted so that large main-memory systems can be used economically, while at the same time, a significant number of computing units (trend towards mega cores!) have become ubiquitous in modern systems. The question on the effects thereof (revolution or evolution?) must be considered based on differentiated aspects. One fact can be noted for sure already: this development requires a completely novel approach on the system level. For example, data structures and base algorithms must be adjusted so that those highly available parallel systems can be used to the full extent – and despite using in-memory technology, it must be possible to guarantee persistence. Thus, are we talking about evolution or revolution here? On the system technology level, it’s clearly the latter!

On the level of business applications, we find a different picture. Here, the transition to in-memory-based technologies will result in a feeling of vertigo if IT-based processes suddenly do not take minutes or even hours to complete but can be executed much faster. A sense of real-time will emerge if decision templates based on large amounts of data can be provisioned to mobile devices within seconds. But, as the proverb goes: appetite comes with eating! Business applications will exploit those novel capabilities and evolve into new application categories. As early as in 1997, the ACM Symposium on “The Next 50 Years of Computing” noted that software behaves in a gaseous way. The saying “software is a gas” was born as a part of “Nathan’s four laws of software.” Nowadays, it is also generally acceptable to say that “data is a gas.” As soon as more data can be stored and processed, this current void will be filled as well. As of now, those software and data gases have not yet spread out from the perspective of operational applications, but this will change soon!

For instance, the application field of simulations will be used to a much larger extent by being able to forecast detailed market and customer behavior within the context of extensive simulations. If users requesting such data- and computation-intensive application scenarios approach these systems with the same expectations they have for traditional scenarios, the field will be doomed. So, how can we confront this trend? Two measures shall be explained as representative examples:

- (a) ... from control flow to data flow: The structure of a variety of business applications is based on control flows, i.e., the modeler or programmer predefines the required processing steps and their order. The necessary data is forwarded to subsequent steps as “ballast.” In order to efficiently exploit in-memory data management technologies, however, we must transform the modeling and implementation of business processes and associated applications into a design concept that is based on data flows. Data is treated as a “first-class citizen” and determines the processing semantics. Prototypical implementations based on SAP’s in-memory technology have shown that the data-flow-centric application modeling and programming yields runtime improvements of several orders of magnitude when employed on an identical platform.
- (b) ... from the principle of ubiquity to the principle of “need to know”: Current interaction patterns with IT systems are characterized by the fact that all information is available at any time, in the most up-to-date form, and with utmost consistency. With regard to the possibility of real-time data processing, this principle of ubiquity may be replaced by the “need-to-know” principle for certain operational application fields. Very often, the real-time data analysis will even overshoot. It may suffice to receive the results at the “right” time, not at “any” time. Thus, if the need-to-know principle is consistently pursued from within the application, this may represent a key method for how to reach equilibrium between cost and benefit in the economical sense when operating in-memory infrastructures.

Revolution or evolution? Most certainly, we cannot answer this question with an “exclusive or.” In order to exploit the full power of in-memory technology, business applications will also have to turn towards alternative structuring, modeling, and programming methods. In particular, however, we must combine the application knowledge with the system architecture. The only veritable way to leverage the potential of in-memory platforms is found in strong interdisciplinary collaborations, and only with the respective application knowledge can the

technology be exploited for further developments.

Prof. Dr.-Ing. Wolfgang Lehner
Director of the
Institute of System Architecture
at the Faculty of Computer Science
Technische Universität Dresden

5 Simplicity is the Name of the Game

5.1 Introduction

Computers are constantly becoming faster and cheaper. Main-memory sizes are growing and processing power is increasing at amazing rates. In the year 2011, a machine with 40 cores and 1 TB of main memory is well-affordable even for a medium-sized business.

Things are getting faster, cheaper, and smaller, but why bother? These trends make working with IT more convenient and more affordable. Furthermore, they enable the development of more sophisticated applications that result in a higher degree of automation (for enterprises) or more fun (for home entertainment). But why should they change the way we build information systems?

This position paper argues that we have reached a cross-over point and that the availability of excessive main memory will indeed change the way that modern information systems will be built in future. The paper speculates that there will be two phases: In the first phase, there will be a revolution resulting in completely new designs for modern information systems. In the second phase, this new generation will evolve in similar ways as today’s information systems have been evolving.

5.2 Phase I: Simplification

In a nutshell, a modern information system has two layers: (a) a database system (e.g., MySQL or Oracle) that takes care of all data management and (b) the application logic (e.g., a CRM or ERP system). The first hypothesis is that in the short term, there will be a rewrite of both layers with dramatically simplified designs and architectures. This hypothesis is based on the observation that a great deal of the complexity of modern information systems (at both layers) is caused by techniques to improve performance. Many of

the techniques embedded in state-of-the-art information systems are obsolete if all the data can be kept in main memory. The following lists some examples how information systems could be simplified:

- **Decision Support:** Many data warehouses will disappear because complex decision support queries can be served directly from the transactional (in-memory) database. This observation has been made by Hasso Plattner (2009) in a recent keynote.
- **Tuning:** Complex physical database designs with indexes and materialized views become obsolete because it becomes affordable to simply scan the data all the time. This observation has been exploited in systems such as Blink (Vijayshankar et al. 2008) and (Unterbrunner et al. 2009). These systems dramatically reduce the complexity of database administration.
- **Storage Hierarchy:** Complex techniques to buffer disk-resident data and synchronize concurrent transactions become obsolete. This observation has been exploited in systems such as VoltDB (Stonebraker et al. 2007) and HyPer (Kemper and Neumann 2011).

These simplifications are not applicable to all information systems. There are certain kinds of data (e.g., Web logs, social networks) that grow at a faster rate than the size of main memory. For such information systems, simplification will be triggered by other hardware and system software trends (e.g., Hadoop on large clusters). Studying such systems is beyond the scope of this position paper, but the general goal of simplification is indeed relevant for all information systems.

5.3 Phase II: Main Memory is the New Disk

History has taught us that performance will continue to be important. It has been critical for all generations of IT systems, no matter how powerful the computers have been. Once the party of Phase I is over and the low hanging fruit has been harvested, organizations will start optimizing the new breed of simplified information systems in order to reduce cost; e.g., purchase less hardware and reduce energy consumption. Furthermore, applications never stop becoming more demanding and every technology shift has also enabled a completely new breed of applications and business models that require new optimization techniques.

Optimizations will be applied first to the database systems, the lowest layer of an information system. Optimizations at this layer have the biggest impact because they can be applied to many applications. Furthermore, the database system vendors have an army of engineers that have been trained to implement such optimizations; these engineers will do just that, after they got the first versions of the new generation of database systems running. These engineers will soon realize that there is a (new) deep storage hierarchy with several layers of processor caches within a computer: Main memory is the lowest layer of that hierarchy, just as disks are the lowest layer in most database systems today. Furthermore, these engineers will learn that a modern machine is a distributed system with many computing nodes and a complex network. So, they will learn to apply distributed database optimization techniques.

For database administrators and application developers, life will also become more complicated, albeit later. At some point, the database systems will have a new set of knobs that allow database administrators to tailor the new breed of main-memory database systems to the specific workloads of an application. As improved performance and reduced cost will continue to be crucial in order to stay competitive, organizations will have sufficient incentives to engage into this arms race; the same kind of arms race we are experiencing today, just with a different generation of arms.

So, after a phase of dramatic simplification, there will be a phase of increasing complexity. The hope is that complexity will increase at a slow rate and that the next cross-over point that triggers dramatic simplification will be reached before we are back to the complexity of today's information systems.

Prof. Dr. Donald Kossmann
ETH Zürich

6 In-Memory Data Management for Business Intelligence

In-memory databases have proven their high practical value for fast transactional processing in real-time applications, such as stock exchange, network routing, and telecommunications services. However, it is still a subject of discussion if the same applies for more analytical enterprise applications services such as business intelligence (BI).

In our opinion, there are good arguments for adopting in-memory approaches in enterprise BI applications. First, online analytical processing (OLAP), a cornerstone of BI and other enterprise applications, is much faster with in-memory storage for user-driven, flexible “ad hoc” queries compared to classical disk-based stores. This saves costs and means lower latency for complex queries. These speed and usability improvements could lead to a new level of quality of in-depth business intelligence, going way beyond simple reporting. It will facilitate data mining and visualization in applications, such as efficiency assessment and reengineering of business processes, but also for enhancing product and service lifecycle management.

Furthermore, the hardware necessary for in-memory data management has become more easily affordable: the price of RAM is, with only few fluctuations, decreasing over time, as are the costs for blade servers that offer scalable parallel processing. There are recent results in database theory on improving the storage and processing efficiency of column-oriented data stores. Such stores could also work in memory without construction of aggregates and explicit indexing, avoiding the corresponding overhead and costs. Finally, column stores are especially well suited for the parallel processing of multidimensional OLAP queries, which often involve aggregations along selected attributes.

On the other hand, before investing into in-memory BI, certain drawbacks need to be considered and mitigated. First and foremost, most existing applications will have to be at least partly rewritten – current ERP and BI solutions will not be able to work in an in-memory environment without major revisions. Companies will have to reengineer their current BI processes and adapt their legacy software, including self-developed database scripts, to the new environment. These revisions cost time and money. Second, there is the base cost of investing into new hardware such as blade servers and enough RAM. However, depending on the lifecycle of existing infrastructure, we do not expect such one-time costs for hardware investments to be critical. Third, software license and support costs for in-memory databases could constitute a major long-term cost factor and possible adoption barrier. Based on current practices, such costs depend on total

memory size, processor count, and number of users. Since the former two dimensions are highly correlated in larger blade systems, database vendors should avoid billing their potential customers twice. License and support costs could also become pivotal if companies consider hybrid strategies by combining existing classical databases and data warehouses with new in-memory stores for OLAP.

Moreover, for consistency and durability and in order to cope with the volatile nature of RAM, backup and emergency recovery strategies (such as transaction logging on hard disk) need to be implemented and supported by acquiring additional, permanent storage hardware. Such backup systems (including second-level backups) and recovery processes are even more crucial than with classical database systems and need to be tested on a regular basis, resulting in personnel costs and possible system downtime.

In spite of these caveats, in our opinion the in-memory approach constitutes a highly promising road for future business intelligence applications, and we expect it to thrive and become commonplace within the next few years.

Dr. Benjamin Fabian
Prof. Oliver Günther, Ph.D.
Humboldt-Universität zu Berlin

7 In-memory Appliances – a Substantial Progress in Information Logistics?

Information logistics supplies actors (humans and systems/machines) with the right information at the right time and in the right quality. While business analytics focuses on specific (decision) tasks, information logistics has a more holistic perspective. Thereby, it covers planning and management of data sourcing, transformation of data into information, and the provisioning of information – be it within a business unit, an enterprise or a business network (Winter et al. 2008).

During the past 20 years, multilayered IT architectures have been established for the technical implementation of information logistics. These architectures were intentionally built on top of existing transactional processing systems. Key layers are data extraction and transformation, integrated storage of all relevant data (“single version of truth” in

the sense of data warehouses), domain-specific processing and dissemination of data (in the sense of data marts) as well as various data analysis tools. From an organizational point of view, a mix of more centralized operating and consulting services (BI Competence Center) and more decentralized, usage-related support services has often been established (Winter and Klesse 2009).

The technical and organizational development of information logistics may be pictured in maturity models which ascribe an advanced state of development to many organizations (Lahrmann et al. 2011a, 2011b). Nevertheless, there are by no means all key requirements met and all major problems solved:

1. Today, data extraction, transformation, preparation, integration, provisioning, and analysis cannot be performed in real time in larger organizations, even when using powerful organizational concepts and IT tools. Underlying inhibitors are rising data volumes, increasing needs for data integration, and ever more diverse data analysis possibilities. For efficiency reasons, analyses have to be optimized carefully and updates are often merged together into batches. Iterative, interactive analytics of live data, as required by many decision makers, are therefore only inadequately supported (Finucane et al. 2010).
2. The many integration and processing steps often caused by specialized tools and infrastructures lead to high complexity and high operating and development costs (Eckerson 2009). In order to achieve substantial complexity encapsulation and cost reductions, the extraction, transformation, preparation, integration, provisioning, and analysis functions have to be brought together in integrated IT systems. However, with existing technologies this seems to be possible only to a very limited extent, as numerous processing steps are indispensable simple because of performance reasons (e.g., pre-computed aggregates in data marts).
3. Due to the need for optimizations in today's IT landscapes, only certain specific analysis functions and paths can be specified and supported efficiently (Finucane et al. 2010). However, in increasingly dynamic decision situations, decision makers ask more and more for solutions that enable them to analyze any desired business object from any perspective, e.g.,

on the basis of individual and spontaneous merge of data. Today, these needs cannot be addressed.

A component that could better meet the speed, integration, and flexibility requirements as outlined above would be a substantial progress in information logistics. This is exactly what in-memory appliances promise (Plattner 2009):

1. Data updates are basically propagated incrementally and in real time. The asynchronous actualization of persistent aggregates (precomputed interim results) can be omitted. Thus, iterative, interactive analyses can be supported much better: The right data is available at the right quality much faster.
2. Transactional and decision-related data is managed in an integrative manner. The need for data redundancy is eliminated and possible inconsistencies are thereby avoided. Processing stages can be reduced and the architectural complexity can be cut down, if the corresponding technology is scalable.
3. Analysis paths are not static so that data can be integrated in any form. Thus, an individual and spontaneous fusion of data would be possible.

In-memory appliances could lead to a very significant advance in information logistics even if only some of these promises were fulfilled. Such architecture components would combine the consistency and efficiency advantages of state-of-the-art extraction, integration, processing, and storage technology (ETL, data warehouse, data mart) with the flexibility and speed advantages of modern analysis technology (e.g., Hadoop). If they were sufficiently scalable for large organizations and could simultaneously be used by hundreds or thousands of decision makers, an immense commercial success would be foreseeable.

The early stage of commercial and pilot projects, however, does not allow a serious evaluation of the potential of in-memory-appliances. Apparently, it became possible not only to address scaling problems of existing in-memory solutions, but also to develop a vision to get over the separation between the management of transactional and decision-related data – even though the pilot projects do not go very far in this respect.

However, the outlined vision should be questioned critically. Information logistics infrastructures fulfill two central functions today: First, they integrate data from different systems. This is currently

enabled primarily by data warehouses. Second, information logistics infrastructures facilitate the immediate provisioning and analysis of information. This is currently enabled mainly by data marts. In-memory solutions can apparently take over the provisioning tasks of data marts and thereby avoid data redundancy and make an entire processing step obsolete. Yet, in heterogeneous system landscapes the integration tasks have to be implemented outside the in-memory-appliance and seem therefore not replaceable in the near future.

If in the course of further development of in-memory-technology transaction processing and decision support would be enabled on the basis of one single core, far reaching consequences for our domain as an integration discipline are foreseeable: After a long period of transactional data integration in the 1970ies and 1980ies, and another long period of process integration in the 1990ies, a whole new fundamental orientation could follow. Business analyses and decisions as well as their flexibility and quality could come to the fore. This development would foster decentralization in regards to (enterprise wide) data, (business unit specific) business processes, and (task-specific) decisions.

On the one hand, the current discussion of the potentials and challenges of in-memory appliances fits to the increasingly important role of analytics, being reflected especially in the U.S. under the label of “Big Data.” On the other hand, it is questionable whether the discussion should be about tools and IT solutions as we had it already with the rise of ERP systems on the basis of integrated transactional databases (e.g., Scheer 1976) or in the context of supporting processes by ERP systems (e.g., Scheer 1995). Although technical and IT solution driven changes appear more effective and more purposeful than cumbersome, hardly predictable, business-driven innovations with corresponding IT support, such transformations often fail exactly because of their focus on technology (Lahrman et al. 2011a, 2011b).

If the discussion should not start with the technical solution (i.e. in-memory appliances), where is it supposed to start? We should specify modern and customized use scenarios of decision based action in organizations that systematically use the potentials of more and more heterogeneous and comprehensive data. Such reference scenarios are, because of

being business driven, a much better starting point for a successful transformation than IT solutions. Our discipline can make an important contribution to this field through reference models and method construction. Software companies can then leverage these contributions in order to develop and hopefully also to establish innovative business applications.

Prof. Dr. Robert Winter
Institute of Information Management
University of St. Gallen

References

to: Section 1

- Bube L (2010) In Memory Datenbanken. Interview-Schlammloch: Plattner vs Ellison. <http://www.crn.de/panorama/artikel-83521.html>. Accessed 2011-09-01
- Loos P, Nebel W, Marx Gómez J, Hasan H, Watson RT, von Brocke J, Seidel S, Recker J (2011) Green IT: a matter of business and information systems engineering? BISE 3(4):245–252
- Plattner H, Zeier A (2011) In-memory data management – an inflection point for enterprise applications. Springer, Heidelberg
- Stonebraker M (2011) My top 10 assertions about data warehouses. Communications of the ACM 54(5):10–11

to: Section 2

- Bernstein PA et al (1998) The Asilomar report on database research. ACM SIGMOD Record 27(4):74–80
- Boncz PA, Manegold S Kersten ML, (2009) Database architecture evolution: mammals flourished long before dinosaurs became extinct. PVLDB 2(2):1648–1653
- Eich MH (1987a) A classification and comparison of main memory database recovery techniques. In: Proc 3rd IEEE international conference on data engineering, pp 332–339
- Eich MH (1987b) MARS: The design of a main memory database machine. In: Proc 5th international workshop on database machines, pp 325–338
- Hur I, Lin C (2008) A comprehensive approach to DRAM power management. In: Proc 14th international conference on high-performance computer architecture (HPCA), pp 305–316
- Kemper A, Neumann T (2011) HyPer: a hybrid OLTP&OLAP main memory database system based on virtual memory snapshots. In: Abiteboul S, Böhm K, Koch C, Tan K-L (eds) Proc 27th international conference on data engineering (ICDE 2011), Hannover, pp 195–206
- Ousterhout J et al (2009) The case for RAM-Clouds: scalable high-performance storage entirely in DRAM. ACM SIGOPS Operating Systems Review 43(4):92–105
- Plattner H, Zeier A (2011) In-memory data management: an inflection point for enterprise applications. Springer, Berlin

to: Section 3

- Plattner H, Zeier A (2011) In-memory data management: an inflection point for enterprise applications. Springer, Berlin

to: Section 5

- Plattner H (2009) A common database approach for OLTP and OLAP using an in-memory column database. In: Cetintemel U, Zdonik SB, Kossmann D, Tatbul N (eds) Proc ACM SIGMOD international conference on management of data (SIGMOD 2009), Providence, USA, pp 1–2
- Vijayshankar R, Swart G, Qiao L, Reiss F, Djalani V, Kossmann D, Narang I, Sidle R (2008) Constant-time query processing. In: Proc 24th international conference on data engineering (ICDE 2008), Cancun, Mexico, pp 60–69
- Unterbrunner P, Giannakis G, Alonso G, Fauser D, Kossmann D (2009) Predictable performance for unpredictable workloads. PVLDB 2(1):706–717
- Kemper A, Neumann T (2011) HyPer: a hybrid OLTP & OLAP main memory database system based on virtual memory snapshots. In: Abiteboul S, Böhm K, Koch C, Tan K-L (eds) Proc 27th international conference on data engineering (ICDE 2011), Hannover, pp 195–206
- Stonebraker M, Madden S, Abadi DJ, Harizopoulos S, Hachem N, Helland P (2007) The end of an architectural era (it’s time for a complete rewrite). In: Koch C, Gehrke J, Garofalakis MN, Srivastava D, Aberer K, Deshpande A, Florescu D, Chan CY, Ganti V, Kanne C-C, Kias W, Neuhold EJ (eds) Proc 33rd international conference on very large data bases (VLDB 2007), Vienna, Austria, pp 1150–1160

to: Section 7

- Eckerson W (2009) TDWI BI benchmark report. TDWI, Chatsworth
- Finucane B, Bange C, Mack M, Keller P, Vierkorn S (2010) The BI-survey 9. Business application research center, BARC GmbH
- Lahrman G, Marx F, Mettler T, Winter R, Wortmann F (2011a) Inductive design of maturity models – applying the Rasch algorithm for design science research. In: Jain H, Sinha A, Vitharana P (eds) Service-oriented perspectives in design science research. Lecture notes in computer science, vol 6629. Springer, Heidelberg, pp 176–191
- Lahrman G, Winter R, Uhl A (2011b) Transformationsmanagement – Aktueller Stand der Praxis und Potenziale. Wirtschaftsinformatik & Management 3(3):62–67
- Plattner H (2009) A common database approach for OLTP and OLAP using an in-memory column database. In: Cetintemel U, Zdonik SB, Kossmann D, Tatbul N (eds) Proc ACM SIGMOD international conference on management of data (SIGMOD 2009), Providence, pp 1–7
- Scheer AW (1976) Produktionsplanung auf der Grundlage einer Datenbank des Fertigungsbereichs. Oldenbourg, München
- Scheer AW (1995) Wirtschaftsinformatik. Springer, Heidelberg
- Winter R, Klesse M (2009) Organizing data warehouse service providers. Business Intelligence Journal 14(3):31–39
- Winter R, Schmaltz M, Dinter B, Bucher T (2008) Das St. Galler Konzept der Informationslogistik. In: Dinter B, Winter R (eds) Integrierte Informationslogistik. Springer, Heidelberg, pp 1–16