Business Benefits And Application Capabilities Enabled By In-Memory Data Management

Gunther Piller[1], Jürgen Hagedorn[2]

[1]University of Applied Sciences Mainz
Lucy-Hillebrand-Straße 2
55128 Mainz
gunther.piller@fh-mainz.de

[2]SAP AG
Dietmar-Hopp-Allee 16
69190 Walldorf
juergen.hagedorn@sap.com

Abstract: Current developments in the area of in-memory data management can significantly change the way business applications will be used in the future. Thus, it will be possible to store huge volumes of single documents directly in main memory for high-speed processing. To introduce and evolve in-memory data management successfully, it is necessary to understand which types of applications benefit most from this new technology. To address this question, we develop typical application patterns. They help to identify promising domains for this innovative technology. We also introduce parameters which support a systematic assessment of corresponding benefits. Our approach is illustrated with examples from different industries.

1. Introduction

In-memory data management (IMDM) is a strongly technology driven innovation: Powerful multi-core processors, availability of massive capacity in main memory and new developments in data organization allow storage of large data volumes in main memory for high-speed processing [PZ11]. From a business perspective, the possibility to access actual and historic data from the entire value chain of an enterprise for interactive real-time analyses has been an idealistic goal for a long time. As information technology comes closer to this goal, it supports not only better decisions but also enables new, innovative business processes (see e.g. [KO07, SAB03, Gar02]).

In this respect in-memory data management is a big step forward. Expectations on its potential for business information systems are high [Gar11]. First applications have been already delivered by software vendors or were implemented as pilot solutions [PZ11, SS11]. To successfully introduce and advance in-memory concepts in enterprises, it is necessary to understand which application scenarios could benefit most from this new technology. This is the focus of our discussion. We sketch how one can identify business domains and processes where IMDM could add significant value. In particular, we focus on the question: Which kind of analyses and decision processes fit best to the capabilities of in-memory technology? Besides an investigation of the improvement potential in existing use cases, it is important to find out how IMDM may enable new business processes. These questions seem to be of interest for software vendors,

consulting companies, as well as enterprises, who consider the use of in-memory solutions on their own.

The usage of IMDM in business application systems is apparently still a relatively new field. A comprehensive or even final assessment is not yet possible. Therefore, this contribution focuses mainly on the introduction of current ideas and trends, as initially discussed in [PH11]. As compared to earlier more technically oriented publications, see e.g. [PZ11], this article concentrates entirely on application aspects for this new technology.

2. Capabilities of In-Memory Data Management

In practice the response times for elaborate computations in current ERP or CRM systems can still amount to several minutes. Consequently, interactive analyses, planning runs or simulations are in many cases not possible. The separation of transaction and analytical processing is one way to deal with this problem in traditional information system architectures: Data, which are usually distributed across several transactional applications, are transferred into analytic systems using cyclical batch processes. As a consequence data available in data warehouse systems are not up-to-date but rather of historic nature. A combined analysis of actual and historic data is, therefore, often difficult (an overview about traditional analytic information systems can be found in e.g. [CG10, MM09]).

IMDM can help in many cases to overcome or reduce the sketched restrictions. In the following we briefly summarize the essential characteristics of this new technology. A detailed discussion can be found in [PZ11].

- **Short response times:** Huge data volumes for analyses, simulations or planning runs can be processed very quickly due to short access and read times, as well as high computation rates. In current performance tests IMDM could easily handle 10,000 queries per hour against 1.3 terabytes of data, returning results within seconds. Here commercially available hardware with 32 cores, 0.5 terabytes of memory and a RAID 5 disk system was used [SI11].
  Essential for this progress are: Increase in clock-speed, the use of multi-core processors, high bandwidth between CPU and main memory through a clever utilization of caches, keeping data exclusively in main memory, as well as column-oriented data storage.

- **Unified transactional and analytic data processing:** The separation of transactional and analytic data processing in traditional application architectures can be overcome with IMDM. All data, including records from latest documents, can be stored in main memory for both usage types. Delays due to cyclical batches for conventional data transfers into data warehouses are absent as already indicated. In

addition the redundant storage of data in operational and analytic systems becomes obsolete.
Crucial for these improvements is the huge, still strongly increasing capacity of main memory. Currently 2 terabyte are often used in practice. Future blade server architectures will allow storage volumes in the petabyte range. Furthermore, with column-oriented data storage efficient compression algorithms are feasible.

- **Analysis of line item data:** Aggregation for the acceleration of response times is no longer necessary. Analyses and planning runs can be based on original data records. Current restrictions due to pre-defined aggregates are obsolete.

The hardware trends sketched above are likely to continue with similar dynamics in the foreseeable future (Moore's Law). Therefore, one can anticipate that adequate powerful computers for IMDM will be available at reasonable costs.
These advantages could give the impression that IMDM may replace current data warehousing concepts completely. However, in our opinion IMDM can improve traditional data warehousing mainly in the areas of response time optimization and data storage efficiency. Multidimensional and relational OLAP implementations and corresponding specific data models may thus become less important. Data warehouse functionality around extraction, cleansing, harmonization and integration of heterogeneous data sets will still continue to be relevant.

Practical experience shows that data warehousing and IMDM can be viewed as complementary approaches. As an example, SAP has been using in-memory technology already for several years to accelerate OLAP-based reporting in their product "Business Warehouse". The next version of the Business Warehouse will optionally be delivered entirely on top of SAP's in-memory platform called HANA (High Performance Analytical Appliance). Main advantages are the reduction of operating costs through a consolidation of data base technologies, as well as an acceleration and simplification of reporting. Here the omission of aggregates and a simplification of data models are important. As a consequence the transfer of data into the Business Warehouse becomes significantly more efficient. In this sense data warehousing can be viewed - in anticipation of Section 3 - as a useful "horizontal" area of application for IMDM.

3. Typical Application Areas for IMDM

In this section we characterize business processes and application categories with respect to capabilities and potential of IMDM. We describe typical application patterns where the usage of IMDM provides additional benefits.

3.1 Selection of Business Processes

As already emphasized, the concept of IMDM is still in a quite early stage of development. A comprehensive assessment is, therefore, difficult. Nevertheless, we formulate general criteria for the classification of application areas for IMDM. In this way we point out future domains that can presumably leverage in-memory technology. We deductively derive meaningful application areas from the described technical capabilities of IMDM. In addition we inductively obtain common characteristics from already existing in-memory applications.

In this sense, promising areas for the usage of IMDM can be identified with the help of the following criteria for business processes:

- **Data dynamics:** How often and predictably do data from a business process change?
- **Range of variation:** To what extend do key figures typically change and how severe is the corresponding influence on corporate success?
- **Number of analysis options:** How many alternatives or working hypotheses shall ideally be compared?
- **Urgency of analysis results:** To what extend is it urgent to obtain analysis results, e.g. to meet given deadlines, carry out subsequent and dependent process steps or to exploit results while still meeting with specialists and managers?
- **Complexity of analysis:** Is an analysis comparatively complex due to underlying algorithms or data structures?
- **Data volume:** Are huge data volumes processed in the considered business activities?

Apparently, the use of IMDM is promising in areas with a high level of data dynamics and variation, many analysis options, urgently needed results, high complexity and huge data volumes. Therefore, the listed characteristics help to judge, whether a business process may benefit from IMDM. Of course, in real use cases all conditions will rarely exist in a single use case.

In business processes, which qualify based on the criteria described above, one can in principle achieve the following effects by using IMDM:

- **Massive increase in analysis frequency** through a drastic reduction of response times
- **Higher analysis flexibility** as predefined reporting hierarchies become obsolete and a multitude of options can be explored
- **Lower data latency**, up to and including evaluations in real-time
- **Enhancement of data bandwidth**, e.g. historic and current data can be analyzed together
- **Higher level of detail** through an access to single documents and line item data

Of course, an assessment of the overall benefit depends on the business impact resulting from changes of individual parameters in a specific use case.

3.2 Application Patterns

Based on the considerations in Section 3.1 one can identify a list of "ideal-typical" application patterns for the use of IMDM. These are explained in more detail below. A summary is shown in Figure 1.

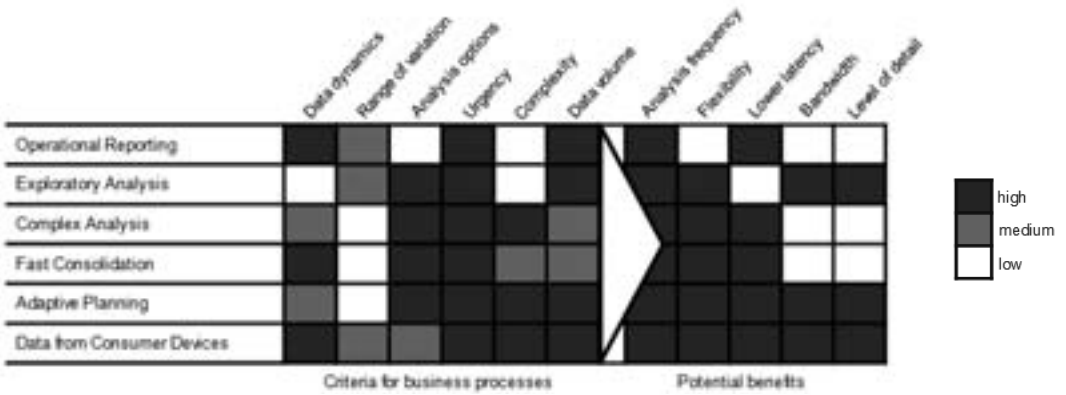

Figure 1: Typical application patterns for the use of IMDM: Business process characteristics and potential benefits.

3.2.1 Operational Reporting

Day-to-day business requires fast and well informed decision making. For this purpose it is often critical to be aware of recent changes in operational transactions and to understand their business context. Here early alerts based on the analysis of many single documents and individual data records are of particular interest.

Based on the criteria of Section 3.1 the pattern "Operational Reporting" can be classified as follows:

- High level of dynamics and variation for data due to their operational character
- Great urgency, as critical situations have to be identified immediately, to rapidly take countermeasures
- In many cases very large volumes of data, as single transactions have to be considered

The potential benefits of IMDM result primarily through:

- Reduced data latency, so that operational decisions can be taken in due consideration of the most recent facts
- Increased frequency of analysis, up to continuous ongoing evaluation. Reports can then be used anytime in operational processes without causing delays.

A striking example for the application of this pattern is the monitoring of sales and inventory data in the retail industry to avoid out-of-stock situations. Corresponding data volumes are typically very high: For leading retailers 100 million data records per day are common. Huge variations of sales figures can be observed in particular during promotions where they reach up to 300%. At present analyses are carried out in larger time intervals. As calculations usually last several hours, batch processing is used. As a consequence, data records, on which new reports are based, are mostly outdated. This often leads to a poor alignment of orders and deliveries. With IMDM most current analysis results can be made available on an ongoing basis. This allows a significant reduction of out-of-stock risks. Replenishment quantities, which are commonly delivered up to five times per day, can then be adjusted to most current sales forecasts.

Quality monitoring is another case for future, IMDM-based operational reporting. For example, in the semiconductor industry wafer fabrication facilities place a great deal of importance on the increase of yields. During production a large amount of material data and equipment parameters are captured continuously. Information from several hundred process steps have to be evaluated and brought into context. The sooner one can identify patterns with a high probability for defects, the quicker countermeasures can be triggered to reduce damaging yield excursions. Besides the operational monitoring of production data, systematic development and improvement of algorithms for early failure forecast and diagnosis is also important here (see e.g. CHH07, Mer11]).

Altogether, it seems that IMDM opens up a considerably more effective operational utilization of information already available today: For example, companies are interested in always up-to-date reports on the payment behavior of customers, e.g. measured in "Days-Sales-Outstanding". If available, they allow rapid interventions in operational processes, like shipments of goods or sales activities. This example clearly shows how the usage of operational data (here for incoming payments) can be extended beyond traditional areas (in this case: dunning) to directly control operational activities.

Another natural domain for the use of in-memory technology seems to be the management of complaints. Here it is obviously critical that a current, complete and consistent picture of a specific customer problem is available without major delay for all involved persons in a company. In this case IMDM can serve as the "single source of truth". In addition it is important to put complaints about deficiencies into a statistic and historic context. Then, in specific customer situations, appropriate actions can be taken, e.g. through accommodating concessions. With the help of the analytic capabilities of IMDM systems, the required analyses could be realized relatively quickly and easily.

3.2.2 Exploratory Analysis of Mass Data

In many cases predefined standard reports are not sufficient for an appropriate assessment of business situations. Often exploratory investigations with a free choice of selectable analysis criteria are required.

Profitability analysis can serve as an almost classical use case: In investigations of contribution margins data are usually aggregated in advance to improve response times. Profitability fluctuations can then be evaluated only for predefined product or market segments. In this case calculations of contribution margins down to the level of customer-product-combinations are not possible. Course grained data also limit detailed evaluations of deviations in contribution margins related to changes in quantity, price, costs, discounts or structural variations. IMDM allows investigations without any predefined aggregation. An exploratory and sophisticated assessment of influencing factors – e.g. in the context of a gross margin flow analysis [Lin88] – is then possible. In this context it is worth to mention that IMDM also offers great potential for the efficient implementation of statistical analysis methods, as well as data mining algorithms. Partially automated top-down-navigation is a corresponding example. Here one aims at a detailed understanding of main causes for aggregated deviations from planed, target or historic figures.

Customer segmentation is another common example for an exploratory analysis of mass data. The idea is to identify customer segments which can be targeted in marketing campaigns successfully and efficiently – often measured through conversion rates. For this purpose huge amounts of customer and market data have to be explored with respect to flexibly chosen combinations of characteristic values. This extremely interactive type of investigation is possible only if response times are short.

The following criteria from Section 3.1 seem obvious for this application pattern:
- Large variations of key figures
- Many analysis options
- High urgency of results
- Large data volumes

Potential improvements through the use of IMDM are:
- Increased analysis frequency
- Increase of analysis flexibility, e.g. following variable reporting structures
- Extension of data bandwidth
- Greatest level of detail

3.2.3 Complex Analysis

There is still a whole series of reports in today's ERP systems which cannot be used interactively due to their long response times. In practice they are therefore often run in batch mode. Examples are material availability checks, complex pricing, cost allocation and dunning runs. It seems natural to use IMDM for a drastic reduction of response times for these analyses and embed them better into operative processes.

Typical for this pattern are:
- High level of data dynamics
- Many decision options with significantly different business consequences

- High urgency to continue operational processes quickly and seamlessly
- Partially very complex analyses
- High data volume

The following improvements can be obtained with IMDM:
- Increased analysis frequency
- Increase of analysis flexibility; different strategies can be explored, e.g. for the allocation of material stock to customer orders
- Reduced data latency; reports carried out in an ad hoc way accurately reflect a momentary business situation

Within the mentioned examples IMDM can be used in practice in the following scenarios:

Available-to-Promise checks provide available quantities of requested products and corresponding delivery dates (see e.g. [KMZ+09]). The sequence of planned deliveries can be changed according to the priority of customers or customer orders. The consequence of changes in the sequence of production or shipment with respect to key figures, like revenue, penalties or customer satisfaction, can be evaluated through corresponding simulations. IMDM allows you to commit to delivery dates immediately, e.g. in conversations with customers. In addition different strategies for the allocation of material can be played through and compared on short notice.

The use case "flexible pricing in sales" seems to be quite similar. In many industries prices depend on a complex set of rules – e.g. in consumer industries. In ERP or CRM systems this is typically realized through condition techniques [HH10]. In particular for sales representatives, who access business applications through mobile solutions, response times in the order of a few seconds – as possible with IMDM – are desirable. In negotiations with customers different options for prices and discounts can then be investigated directly.

3.2.4 Fast Consolidation

The term "Fast Consolidation" summarizes all use cases, where documents created in the past have to be consolidated according to given rules up to a fixed deadline, or on request. Typical are attempts to accelerate the financial closing for the filing of monthly, quarterly or annual reports – often referred to as "Fast Close" (see e.g. [Bra09, Sch06]). Amongst other steps it is necessary here, to reconcile internal payables as quickly as possible. In large enterprises this often involves the evaluation of several million invoice documents. For many companies it is also important to create consolidated interim reports with most current data for different reporting levels on short notice.

Altogether this field of application is characterized through:
- Huge data dynamics due to documents created on short notice
- Different analysis options, e.g. for the allocation of costs

- High urgency through deadline constraints, e.g. as a consequence of legal requirements or stakeholder demand
- High level of complexity though a comprehensive body of rules and regulations or elaborate calculation methods
- Large data volumes

IMDM can lead to the following improvements:
- Increase of the number of consolidation runs
- Increased flexibility; different scenarios for consolidation can be compared
- Low data latency; short term postings can be accounted for immediately

One can imagine that the use of IMDM will massively influence the way consolidations are carried out: Batch procedures - which still dominate today - will evolve towards interactive and iterative practices. In this way different calculation schemes can be run through and compared. In addition one could work with original documents instead of replicated data. Problems with data consistency and efforts for corresponding adjustments would be reduced and eventually become obsolete.

A similar use case is the consolidation of up-to-date sales figures, like expected or closed deals: In many industries the closing of deals is emphasized towards the end of quarters - sometimes driven by budget cycles. As a consequence there are many late transactions - just before quarter closing. To allow for short-term interventions of sales management teams and to support the preparation of quarterly reports, it is necessary to provide consolidated results with most current data as fast as possible.

3.2.5 Adaptive Planning

The use of IMDM for planning scenarios has been discussed by Sinzig and Sharma in [SS11]. For completeness we provide a short classification with respect to the criteria from 3.1.

Planning processes often show the following characteristics:
- Partly high level of data dynamics, e.g. in cases where the latest customer requests need to be included
- High number of planning options which potentially have to be considered
- High level of urgency to trigger subsequent business activities
- Complex planning algorithms
- Large data volumes, in particular for company-wide planning runs

Potential benefits from IMDM arise from:
- Plenty of planning runs with different variables
- Reduced latency of data, so that planning results can be adjusted better to momentary circumstances

- Increase of data bandwidth, to enable company wide, integrated planning - avoiding "planning-silos"
- High level of detail for planning objects

A classical example can be found in production planning [KMZ+09, SK05]. Planning runs using MRPII or APS usually need many hours. Therefore, necessary computations are typically carried out in batch mode and repeated in intervals of several days only. To save on computing power, planning runs are often restricted to master scheduled items. The consequences are inaccurate results with partly outdated information. Due to reduced computing time, planning runs can be completed with IMDM in short time intervals with always up-to-date starting parameters. Then realistic and exact results on the level of materials can be expected. This enables flexible reactions to changes in market conditions – e.g. short term fluctuations in customer demand.

3.2.6 Analysis of Data from Consumer Devices

For providers of traditional business software, like SAP, the use of IMDM seems to open up a new domain: In-memory technology enables the processing of extreme volumes of highly dynamic consumer data from embedded systems, as they appear e.g. in domestic homes.

"Smart Meter Analytics" can serve as a first use case for which a solution is currently in development at SAP. Consumer data from electric meters located in households and businesses are central to this application. These data are brought together using intelligent electric supply networks - so called smart grids. With IMDM they are then analyzed in nearly real time. On the one hand, this allows consumer-oriented solutions, which e.g. benchmark the energy consumption of individual households and serve consumers with suggestions for improvements. On the other hand, it is planned to provide components which support utility companies in their efforts to dynamically adjust capacities and prices to actual consumption profiles.

According to the classification from Section 3.1, large potential for the usage of IMDM can be expected:
- Apparently extremely high dynamics of consumer data
- Potentially large range of variations, e.g. for electricity consumption
- Many alternatives, e.g. different pricing policies for electricity
- High urgency, e.g. to immediately react to particular consumption patterns of consumers
- High complexity of analyses, e.g. for simulations
- Extremely large data volumes

Expected benefits are:
- Massive increase of analysis frequency, continuous data access for consumers
- Significant rise of analysis flexibility: The consequences of different measures, e.g. for pricing, can directly be simulated.

- Reduction in data latency: All analyses are based on the most recent consumer behavior.
- Maximal bandwidth of data: In an ideal case all actual and historic data are available.
- High level of detail: For example, one can utilize data for electricity consumption, which are typically measured in time intervals of a few minutes.

4. Summary

This article highlights actual and future key areas for the use of IMDM in business information systems. Our discussion should be seen as a current snapshot rather than a final assessment. This is mainly due to the fact that the use of IMDM in business applications is still in an early stage of development. Therefore, a detailed evaluation of particular application scenarios is left to future investigations.

In our opinion, it already becomes apparent that in-memory technology will have a massive impact on the redesign of existing information systems and the preferred way of their use within business processes. Applications, which today can only be operated in batch mode due to extensive computing times, will - through IMDM – evolve towards interactive, exploratory, ad hoc usable solutions.

Furthermore, IMDM opens up a completely new type of business applications: Solutions which involve consumers seem to be of particular interest here. The aim is to leverage the corresponding huge data volumes in “consumer apps”, as well as applications for business purposes.

References

[CHH07] Chien, C.; Hsu, S.; Hsu C.: Enhancing Competitive Advantages and Operational Excellence for High-Tech Industry through Data Mining and Digital Management. In: Liao, W.; Triantaphyllou, E. (Eds.): Recent Advances in Data Mining of Enterprise Data: Algorithms and Applications. World Scientific, 2007.

[Bra09] Bragg, S.: Fast Close. John Wiley & Sons, 2009.

[CG10] Chamoni, P.; Gluchowski, P. (Eds.): Analytische Informationssysteme. Springer, 2010.

[Gar02] Gartner: Now is the Time for Real-Time Enterprise. 2002. http://www.gartner.com/pages/story.php.id.2632.s.8.jsp. Retrieved June 15th 2011.

[Gar11] Gartner: Gartner Identifies Nine Key Data Warehousing Trends for the CIO in 2011 and 2012. Gartner Press Release, February 2011. http://www.gartner.com/it/page.jsp?id=1542914. Retrieved June 15th 2011.

[HH10] Hirn, M.; Herhuth, W.: Preisfindung und Konditionstechnik in SAP ERP. SAP PRESS, 2010.

[KMZ+09] Knolmayer, G.; Mertens, P.; Zeier, A.; Dickersbach, J.: Supply Chain Management Based on SAP Systems. Springer, 2009.

[KO07] Kagermann, H.; Oesterle, H.: Geschäftsmodelle 2010. FAZ, 2007.

[Lin88] Link, J.: Verbreitung und Einsatzformen der Deckungsbeitrags-Flussrechnung in der Industrie. Die Betriebswirtschaft, Nr. 6, 1988.
[Mer11] Mertens, P.: Integrierte Informationsverarbeitung I, Gabler, 2011.
[MM09] Mertens, P.; Meier, M.: Integrierte Informationsverarbeitung II, Gabler, 2009.
[PH11] Piller, G.; Hagedorn, J.: Einsatzpotenziale für In-Memory Data Management in betrieblichen Anwendungssstemen. Wirtschaftsinformatik und Management, Nr. 5, 2011.
[PZ11] Plattner, H.; Zeier, A.: In-Memory Data Management. Springer, 2011.
[SAB03] Scheer, A.; Abolhassan, F.; Bosch, W. (Eds.): Real-Time Enterprise. Springer, 2003.
[Sch06] Schulte, O.: Fast-Close-Abschlüsse und Schadenrückstellungen nach HGB, IAS/IFRS und US-GAAP. Deutscher Universitätsverlag, 2006.
[SI11] SAP Inside: First HANA Performance Results Released. March 2011. http://www.insidesap.com.au/_blog/News/post/First_HANA_performance_results_released. Retrieved June 15th 2011.
[SK05] Stadtler, H.; Kilger, C.: Supply Chain Management and Advanced Planning. Springer, 2005.
[SS11] Sinzig, W.; Sharma, K.: In-Memory-Technologie: Verbesserung bei Planung, Simulation und Entscheidungsunterstützung. Wirtschaftsinformatik und Management, Nr. 2, 2011.