

Successful In-Memory Database Usage – A Structured Analysis

Completed Research Paper

Alexa Scheffler

University of Augsburg

Alexa.Scheffler@wiwi.uni-augsburg.de

Sarah Otyepka

University of Augsburg

Sarah.Otyepka@wiwi.uni-augsburg.de

Abstract

In-memory databases (IMDB) are a raising technology and have the potential to mark the end of computing performance bottlenecks. Our literature review did reveal that to date, no study analyzes the resulting benefits of IMDB usage and little is known about cause and effect of IMDB usage and its economic effects. This paper provides a structured analysis of experiences gained with IMDB usage. Surprisingly, our study showed that the promoted vision of IMDB as enabler for an integrated OLAP/OLTP infrastructure has not been put into practice yet. Currently, IMDB is mainly used for response time improvement in the field of analytical data processing. Furthermore, we observed that IMDB is predominantly used to improve existing business processes, instead of establishing new business models. The findings of our study may serve as basis for the development of hypotheses, validation and advancement of the theory of disruptive innovation and further empirical quantitative studies.

Keywords

In-Memory Databases, Benefits, Real-Time Business Intelligence

Introduction

“This break-through innovation will lead to fundamentally improved business processes, better decision-making, and new performance standards for enterprise applications” says SAP’s co-founder Hasso Plattner about in-memory databases (IMDB) (Plattner and Zeier 2011). IMDB stores data in main-memory instead of physical disks, and therefore gains significant faster response times. IMDB vendors promise real-time insights in company data due to 1,000 times faster response times. At the same time they claim that IMDB enables to simplify data models and to reduce tuning activities. Hence, IT infrastructure is less complex and IT expenses can be reduced (Evans 2011). However, the technology is still in the process of maturation, and hence struggling with teething troubles. Companies have just started to evolve a pool of experience. After a flurry of overenthusiastic press releases – mainly pushed by software vendors – a sound understanding of the technology’s benefits and threats is needed (Read 2013). Since Gartner’s Hype Cycle for Emerging Technologies indicates that IMDB is at the end of the phase “peak of inflated expectations” (Pezzini 2013), proofs of the applicability of IMDB and rational considerations of its valuable use are required (Loos et al. 2011). From a scientific point of view, IMDB might be a future disruptive innovation according to Christensen (1997). We plan to categorize IMDB as a disruptive technology in a next step and survey the development of IMDB over the next years in order to identify the characteristics of the typical niche market for theory validation in a scientific context. In order to do so, we first require a better understanding of the development of IMDB. Therefore, this paper focuses on an exploratory analysis of experiences with IMDB gained so far and its realized benefits.

Apart from scientific interest, economic considerations about the usage of IMDB are relevant to different interest groups: software vendors as well as potential IMDB customers have to weigh the opportunities against threats and are faced with the decision whether to invest in IMDB or not. On the other hand, the interest group of software vendors has to decide if it is necessary to get their products IMDB-ready in order to persist on the market. For companies as potential customers, the introduction of IMDB offers a

lot of benefits like advanced reporting functionality or lower cost for IT governance, but may also bear the risk of an unreliable IT infrastructure or other unforeseeable problems due to a lack of knowledge in the IMDB field.

To date, scientific literature mainly covers technical issues or is concerned with visions of future IMDB use. A structured literature review in major IS journals and conference proceedings revealed only a few papers, that are concerned with the business perspective of IMDB usage. In order to gain more systematic insights to the current status of practical IMDB usage and the resulting benefits, this paper provides an exploratory study of so-called IMDB “success stories”. In this paper “success stories” are defined as reports about the successful introduction and subsequent usage of IMDB. We used these success stories to discover similarities and patterns in existing successful IMDB application scenarios. The results of the following questions are structured in categories and subcategories in order to provide a systematic overview:

- What kind of problems existed before IMDB was introduced?
- What are the main objectives of the IMDB introduction?
- What are the experiences with IMDB implementation projects?
- How did IMDB change IT architecture design?
- What are the realized benefits of the IMDB launch?
- What kind of new business models are enabled by IMDB?

We take into account that such success stories tend to be biased towards positive results and seldom consider failure. Hence, this type of sources can only be used as a first indicator for patterns of IT infrastructure, project experiences and resulting benefits of IMDB use.

The results address both scientists in the field of disruptive IT innovations, value of IT, and benefits management as well as practitioners (future IMDB users, IT consultants, software vendors). The findings may serve as basis for detailed empirical quantitative studies that need to establish hypotheses about cause and effect relationships. Furthermore, the results of this paper can serve as benchmark for companies who seek for potential use cases for IMDB and are confronted with the decision whether to invest in IMDB or not.

The paper is structured as follows: After enlightening the target system and problem context referring to related literature in chapter two, we describe the used methodology of success story analysis in chapter three and present the key results in chapter four. Limitations and a structured proposal for a future research agenda are presented in the last chapter.

Problem Context and Related Work

IMDB belong to the in-memory computing paradigm which is defined as “a computing style, which assumes that the primary data store for applications [...] is the central (or main) memory of the computing environment running these applications” (Pezzini 2013). In contrast to traditional databases that employ a disk storage mechanism, an IMDB stores all data in main memory (Garcia-Molina and Salem 1992). The different properties of main memory compared to hard disk drives lead to significant shorter response times of IMDB compared to disc-based databases (DBDB). For a long time, high hardware cost made main memory, and therefore IMDB, rather a topic for scientific discussion than for commercial production. But the development of constantly decreasing hardware cost – main memory average prices per gigabyte decreased from 6,328,125.00\$ in 1980 to 5.50\$ in 2013 (McCallum 2013) – has made the usage of IMDB economically feasible for large enterprise systems (Acker et al. 2011).

Today, IMDB can be used for operational or analytical purposes. An IMDB for operational purposes supports control and run of business tasks, run by online transaction processing (OLTP) business applications like enterprise resource planning systems (ERP). An analytical IMDB is a DB used for online analytical processing (OLAP) business applications and addresses analytical needs of managerial decision support. It may not be confused with in-memory analytics which is an alternative business intelligence (BI) architecture layer in which detailed data is loaded into main memory to accelerate query processing.

Furthermore, there are hybrid application scenarios of where transactional and analytical systems both use the same IMDB – also referred to as the integration of OLAP and OLTP systems (Pezzini 2013).

In order to gain an overview of existing scientific contributions, we used a structured literature review as proposed by (Webster and Watson 2002). Since this paper aims at the identification typical use case scenarios and benefits of IMDB usage, the literature review focuses on publications that deal with the business perspective of IMDB. Publications that are solely concerned with the technical perspective were excluded from further analysis. Our search strategy is listed in Table 1. The given scientific databases were selected because they maintain journals of the Social Sciences Citation Index. In order to proof scientific reliability only peer-reviewed journals and conferences were taken into account. Potentially relevant publication had to meet the search expression (*“in-memory” OR “main-memory” AND (“data*” OR “comp*”)*) for at least one of the following fields: title, abstract or key words

.Search criteria	Value
Scientific Database	ACM Digital Library, AISeL, EBSCOhost, ProQuest, JSTOR, ScienceDirect, SpringerLink and Wiley InterScience
Publication type	Peer-Reviewed Journals and Conferences
Fields	Title, abstract and key words
Search term	<i> (“in-memory” OR “main-memory”) AND (“data*” OR “comp*”)</i>

Table 1. Search Strategy for Literature Review

The search strategy described above lead to the following results: Many publications in the field of IMDB research are concerned with technical aspects, e.g. (Garcia-Molina and Salem 1992, Faerber et al. 2011), but only few obtain the business perspective like (Wessel et al. 2013, Loos et al. 2012, Loos et al. 2011, Winter et al. 2011, Piller and Hagedorn 2011). Wessel et al. identify three potential consequences of IMDB that are repeatedly mentioned in literature and analyzes its accordance in literature. The identified potential consequences are: a) IMDB helps to realize the vision of real-time BI, b) IMDB helps to solve challenges of Big Data and c) IMDB drives the integration of OLAP and OLTP systems. (Loos et al. 2011) present a discussion of eleven scholars about potentials and consequences of IMDB introduction. They discuss the future integration of OLAP and OLTP systems, new modeling and programming approaches, improved quality of BI and of new IMDB-enables business models in energy management and cloud computing. (Loos et al. 2012) respond to the discussion of (Loos et al. 2011) and provide visions of IMDB use in operational analytics. (Piller and Hagedorn 2011) develop application patterns of IMDB usage that consist of the dimensions “criteria for business processes” and “potential benefits”. Criteria for business processes help to find “promising areas for the usage of IMDB” and are defined as data dynamics, range of variation, number of analysis options, urgency of analysis results and data volume. Business processes that considerably fulfill these criteria can gain the following benefits: massive increase of analysis frequency, higher analysis flexibility, lower data latency, enhancement of data bandwidth and higher level of detail. First analysis of consequences for the business perspective are presented by (vom Brocke et al. 2014, Winter et al. 2011). (Winter et al. 2011) propose several application scenarios for IMDB usage at an international automotive company: quality and fault analysis, calculation of procurement commissions, monthly closing of financial statements. IMDB helped the company to abandon pre-aggregated key figures and materialized view – and therefore, to simplify its IT architecture – and to speed up response times for reports. (vom Brocke et al. 2014) do so at a manufacturing company (namely Hilti). They identified seven promising application scenarios for IMDB at Hilti. 1) Collection of position, status and usage data of tools, 2) capturing and analysis of Hilti-related web content (especially in social media) and further analysis with text mining, 3) advanced knowledge management and 4) introduction of voice control.

Research Design

Since, the review of scientific literature revealed that there are hardly any scientific publications that provide insights into practical usage of IMDB (as for example case study research does), our study is based on reports about the successful usage of IMDB (“success stories”) that were gathered from web sites publishing IT news and industrial publications.

We are aware that these types of sources involve several limitations. Opposed to sophisticated scientific publications that aim to increase existing knowledge and follow a rigor approach (Dubé and Paré 2003), these kind of sources often follow marketing purposes and therefore lack objectivity. This circumstance results in the fact that only positive experiences are reported and problems, failures or critical considerations are out of scope. Another shortcoming is that data origin and approaches for data gathering are often not exactly clear. Furthermore, statements and conclusions are not thoroughly discussed and therefore rather elusive. In order to bring sound knowledge about cause and effect relationships of IMDB usage a comprehensive empirical study is necessary. Considering the fast pace of IMDB development and the increasing pressure to make decisions about strategic information architectures, there is a need for faster answers. In order to provide quick results for IMDB research, we decided to use quickly accessible data of success stories and accept its limitations. We followed a lightweight version of the case survey methodology as described in (Larsson 1993).

Our study which intends to be a first step towards more elaborated investigations comprises a sample of 38 records of IMDB success stories. Each success story has to fulfill the following requirements to be included in the sample data set. At first, an IMDB instance is actually running in the company. Reports that only report future *plans* of using an IMDB were discarded. Secondly, the success story has to provide a certain level of detail, i.e. company background, project information, information about IMDB application in daily business and benefit description.

The success stories were gathered in a systematic approach illustrated in Figure 1. It consists of the three stages *data collection*, *review* and *analysis* that are explained in detail in the following sections.

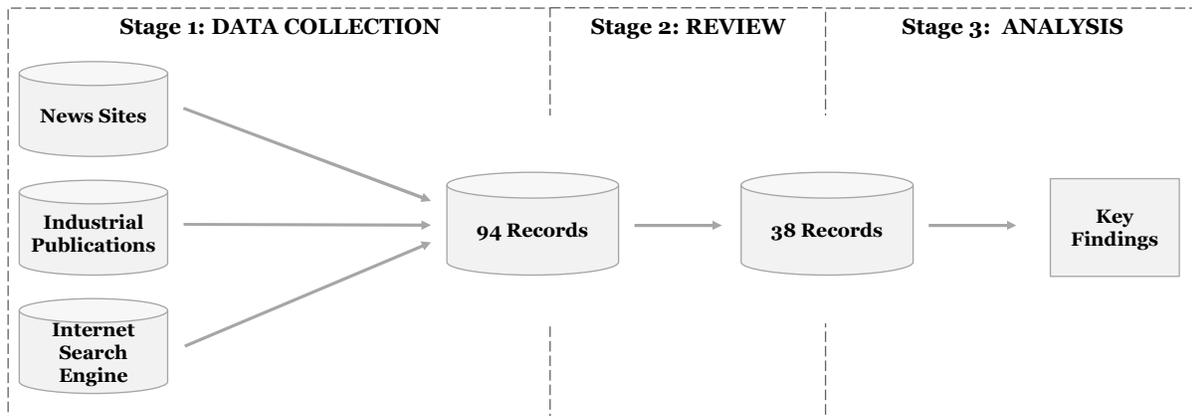


Figure 1: Research Design

Stage 1: Data Collection

According to our systematic literature review there are no scientific publications which analyze the existing usage of IMDB. Hence, we focused on other source types that contain reports about IT scenarios: news sites and industrial publications. For the initial data collection we used a forward and a backward search strategy. The forward search comprises a systematic review of news sites and industrial publications. **News sites** comprise cioinsight.com, computerweekly.com, techrepublic.com, techtargget.com, information-management.com. Since IMDB vendors run intense marketing activities in Germany (e.g. for SAP HANA), and therefore persuaded numerous German companies to act as testimonial customers, we additionally included IT news in German language (these are isreport.de, zdnet.de, managet.de).

The second source type **industrial publications** comprises experience reports and case studies of IMDB vendors (SAP, IBM, Oracle), and IT research institutes (Business Application Research Center (BARC), Forrester Research, Gartner Inc., The Data Warehousing Institute (TDWI)). We included the leading IMDB vendors and well-known research institutes in the search.

The forward search resulted in a few reports that mentioned company names, but did not fulfill the requirements of a success story. We collected these company names for a backward search and used an internet search engine to find more detailed reports about its IMDB usage.

The two approaches (forward and backward search) resulted in 94 records that are collectively stored in a sample data set. Each record includes document data (author, title, source, source type, download date, publication date) and case description (company name, industry, number of employee IDMB product).

Stage 2: Reviews

In order to contribute to increased objectivity, the sample data set was controlled by a peer-reviewer. We carefully excluded those records that do not meet the requirements stated above. Consequently, the review produced 38 reports for the sample data set. Other reports that did not answer the questions were discarded.

Stage 3: Analysis

The aim of the third stage is to find categories and subcategories and hence, answer the initial questions in chapter 1. To do so, we used the deductive coding technique (Saldana 2013). Each report was reviewed in parallel by two researchers and coded. The codes were discussed by both researchers and consolidated to proper categories. The resulting structure and the findings are presented in the next chapter.

Results

The first section proposes a structured description of IT architecture prior to and post IMDB introduction. These characteristics help to understand technological aspects that drive companies to introduce an IMDB and the involved changes in the IT architecture. The second section describes findings concerning the data sources that provide the companies with data. The third section analyzes the project experiences and in the last section, we present the resulting benefits of IMDB usage.

Following a deductive approach, we do not claim the characteristics to be mutually exclusive and collectively exhaustive. Since not all records contain information on every topic or characteristic, we only count them in if there is a positive statement. This means, we present restrictive numbers that are based on the success stories with positive statements. Consequently, a success story can count in none, one or more than one characteristic of each topic.

IT architecture

The motivation to implement an IMDB varies from company to company. It ranges from rapidly increasing amounts of data, to the need for timelier answers to business questions, to overcharged legacy-systems, to the need for a first analytical system.

An analysis of the situation prior to IMDB implementation shows that only few companies (11%) did not run analytical systems prior to the project. Looking for strategies to introduce an analytical database, they decided in favor of IMDB and thereby avoided complex data modelling activities. This shows that surprisingly, IMDB is also an attractive alternative for companies without BI experience. All companies that did not run a data warehouse (DW) prior to IMDB introduction, are small- and medium sized. They introduce an IMDB-based DW as their first analytical database.

We observe that bigger companies begin with a pilot project to evaluate IMDB usage. Most of the pilot projects introduced IMDB as sidecar to the existing database (32%). 53% of records report that follow-up projects are planned or already ongoing. These project scopes range from moderate projects (like expanding additional regions or additional subsidiaries), to more broadened projects (like full shift of existing ERP systems on IMDB). The detailed results are illustrated in Table 2.

Topic	Characteristic/Value	% of total sample
OLAP systems prior to IMDB	DW (non-IMDB-based)	42%
	No dedicated OLAP system	11%
	Insufficient information	47%
IT Architecture post IMDB launch	IMDB as database “sidecar”	32%
	Full DW on IMDB	21%
	ERP and DW on IMDB (integrated OLAP and OLTP)	3%
	Full ERP on IMDB	0%
	Insufficient information	44%

Table 2. IT architecture

Project experiences

After analyzing the selected records in terms of project characteristics, we found that the project scope equally ranges from IMDB-supported ERP systems, to shift of existing DW to IMDB, to complete relaunch of ERP systems and DW on an IMDB platform. Most of the records describe the implementation as pilot projects with limited and clearly defined scope, which were realized with on-demand applications.

The examined projects have short project duration. In total, 29% of records give an indication of their implementation period (i.e. time for implementation, testing and roll-out without time for requirements analysis and product evaluation). 13% finished the projects in less than two months, 11% in less than four months. At the same time, most analyzed projects had a very limited scope, which might related. The detailed results are illustrated in Table 3.

Topic	Characteristic/Value	% of total sample	
Project Scope	Existing ERP system and DW prior to project	(Parts of) Existing DW is migrated on an IMDB	55%
		(Parts of) Existing ERP system is migrated on an IMDB	16%
	No existing DW but existing ERP system prior to project	Launch of DW for selected business cases on IMDB	16%
		Launch of full DW in IMDB	13%
Project duration	Less than 2 Months	13%	
	2-4 months	11%	
	Insufficient information	76%	

Table 3. Project experiences

Benefits

The analysis of the success stories revealed five major benefits, shown in Table 4 and discussed in the following: Productivity improvement, usability improvement, higher scalability, improved information provisioning and cost reduction.

Topic	Characteristic/Value	Multiple selection	
Productivity improvement	Faster response times of queries	100%	
	Increasing number of parallel user queries	13%	
	More effective data compression (than DBDB)	16%	
	Effects on the generation of financial statements	Faster generation of financial statements	37%
		Therefrom: Faster consolidation of international companies and internal service charges	13%
		Therefrom: Adjustment Postings can be revised and corrected (iterative decision making)	13%
Usability improvement	Ability of report design is not limited to IT experts (→Self-service-BI)	82%	
Higher scalability	Ease of handling growing amounts of data	47%	
	Ease of handling growing numbers of users	Increasing number of users due to added user groups	11%
		Increasing number of users within existing user groups	8%
Improved information provisioning	New analysis options	Inventory management/ dynamic demand/supply adjustment	24%
		Customer needs/cross selling potentials	16%
		(Dynamic) Pricing	16%
Cost reduction	Lower cost of IT administration/operation/ data storage	21%	

Table 4. Business benefits

Productivity improvement

All records show, that IMDB seems to enable significant faster response times for analytical queries. Inventory reports, for example, had a response time of 22 minutes and turned out to be incomplete, whereas after IMDB introduction, it takes 17 seconds to generate an accurate report. 16% of the cases observed IMDB to work with more effective data compression algorithms, which lead to a reduction of the database volume by the factor two to four, e.g. from prior 550 Gigabyte to about 130 Gigabyte.

Another example for significantly faster response times is generating financial statements. In several cases, financial month-end-closings require less than three days with IMDB, whereas it took seven days or even more without IMDB. 37% of the records name the faster generation of financial reports one of the major benefits. 13% of them additionally found company consolidation to be much easier and another 13% found the system to be fast enough to revise and correct adjustment postings. This enables iterative decision making which is seen as a big benefit.

Usability improvement

Since, reports do not have to be pre-generated, users have the ability to define reports on their own and therefore, perform ad-hoc queries.

Higher scalability

Due to the ability to execute more queries in parallel, more users can access data to support decision making. This might lead to entirely new application scenarios like advanced sports statistics provided online for general public shortly after the NFL matches in the USA.

Improved information provisioning

Besides getting query results faster, most companies also are able to improve their information provisioning for decision making by running new types of analyses with IMDB. These new types of analyses comprise 37% profitability analyses (e.g. elasticity of market, demand or price) and 29% scenario analyses (what-if analyses) to be the most named.

Cost reduction

Surprisingly, many success stories (21%) suggest lower IT costs after introducing IMDB. The main reasons are elimination of database tuning activities, lower number of standard reports (due to more possibilities of self-service BI), and decreasing storage demand (due to more effective compression algorithms).

Summary

In order to summarize the key findings, we refer to the questions posed at the beginning of this article.

What kind of problems existed before IMDB was introduced? Companies had to deal with constantly growing amounts of data, data sources, users, as well as increased analytical reporting requirements. Reporting was difficult and about to become impossible for complex international companies. Existing systems were overcharged, and the analytics systems were unable to keep up with extended requirements. Companies that were not using a dedicated analytical database, decided to introduce their DW on IMDB to avoid complex data modelling activities. They are all medium sized companies who were in need of analytical systems and deliberately decided for IMDB to set up a future-proof IT architecture.

What are the main objectives of the IMDB introduction? Named requirements were faster response times or ad hoc reporting, easy integration into the existing IT-landscape, support of real-time analytics, and advanced analytical functions.

What are the experiences with IMDB implementation projects? Although, none of the records gives specific numbers about project cost, information about implementation times are available and allow anticipating project sizes. The implementation times range from 3 weeks to 4 months. One of the IMDB vendors offers packages of IMDB project implementation for a fixed price that makes it easier to calculate project costs. One repeatedly mentioned project experience is the problem caused by the lack of technical experts at the moment.

How did IMDB change IT architecture design? The analyzed reports do not suggest that the vision of integrated OLAP and OLTP systems has been realized, yet. Only one of the companies described such a project scope, but it did not run a DW prior to the IMDB introduction. Currently, most companies add an IMDB to their existing IT architecture and run it as a sidecar to DBDB.

What are the realized benefits of the IMDB launch? Short response times and immediate availability of new data enables the companies to iterative decision making, which is considered to be a big advantage. 42% of the companies are now able to run analyses that were not available without IMDB due to technical restrictions. Additionally, in 45% of the cases analyses might have been available, but were not used due to unacceptable response times. At 47% of the companies, it was possible to get information for obligatory reporting without IMDB, e.g. financial statements, but the IMDB makes the process much easier and quicker. 42% of the records stated that they use completely new types of analyses that had not been available with former DBDB.

What kind of new business models are enabled by IMDB? Many companies chose business cases within the finance department for their pilot project, e.g. applications for revenue management, financial planning, and budgeting or business consolidation systems. Since financial reports are obligatory, and liable to compliance and legal restrictions, companies see the benefits of the investment in this business case. Other application scenarios that seem attractive for IMDB pilots are highly volatile business areas where IMDB can enable new analysis options that cannot be provided by existing DBDB. One repeatedly application scenario is inventory management. Since IMDB enables dynamic calculation of demand, inventory costs can be reduced. Another application scenario is customer relationship management which

is able to identify cross-selling potentials to raise revenue. For example, the gaming company Big Point uses IMDB to calculate the right time to invite customers of free game-versions to upgrade to charged versions. Only for one company IMDB is a key component to realize a new business model. WeissBeerger provides solutions for monitoring beverage consumption for bars, breweries and beverage distributors. It offers flow sensors for taps, so bars can introduce a “self-pour” business model. In combination with their in-memory based product Alcohol Analytics dynamic pricing is possible. Since, this report is the only hint for new business models, hence it can be concluded that the development of new IMDB-enabled business models is still at an early stage.

Conclusion

In order to gain an overview of current IMDB usage, this study proposes a structured analysis of news site and industrial reports about successful IMDB projects - so called “success stories”. This type of data source has to be chosen carefully, because it is often influenced by vendors. Therefore, it is subject to several limitations, such as uncertainty about objectivity, and lack of profound research. Another limitation is the fact that only successful implementation projects are regarded. An interesting insight would be to know why potential customers have decided against IMDB or why IMDB implementation projects failed. Since the IT sector, in contrast to other industries – like e.g. aviation industry – has not a well-established culture of failure research, only success stories lie within the scope of our possibilities. Therefore, it is possible to propose at least preliminary answers to urgent questions in the fast developing field of IMDB. Despite these limitations, our study may serve as a modest contribution to current discussions in scientific literature: In contrast to the wide-spread opinion in scientific literature that IMDB will lead to the integration of OLAP and OLTP systems (Wessel et al. 2013, Plattner and Zeier 2011), we have only found one record that gives evidence to this development. Another opinion of IMDB potentials is that it opens up new business models (Loos et al. 2012). Since our data set shows only one record that confirms this assumption (real-time beer consumption statistics that helps brewers to help them discover the immediate impact of campaigns), it seems the potential for new IMDB-enabled business models has not been put into practice yet. Our study may only indicate future directions of IMDB evaluation. Besides aiming at an overview of current IMDB usage, we see the results of this paper as a first step of better understanding IMDB development. This is necessary to classify IMDB as potential disruptive technology in a next step. Using IMDB in this early stage of maturity as an application scenario for disruptive technologies allows us to work on theory validation, e.g. in terms of characteristics of the niche market which allows technology development or in terms of predictive use of the theory. However, in order to develop sound theories about cause and effect relationships, it is indispensable to conduct a large-scaled study to derive verifiable hypotheses.

REFERENCES

- Acker, O., Groene, F., and Blockus, A. 2011. "In-Memory Analytics - Strategies for Real-Time CRM," *Journal of Database Marketing & Customer Strategy Management* (18:2), pp. 129-136.
- Christensen, C. M. (1997). *The Innovator's Dilemma*, Boston, MA: Harvard Business School Press.
- Dubé, L., and Paré, G. 2003. "Rigor in Information Systems Positivist Case Research: Current Practice, Trends, and Recommendations," *MIS Quarterly* (27:4), pp. 597-636.
- Evans, B. 2011. SAP HANA and the Game-Changing Power of Speed. Available at: <http://www.forbes.com/sites/sap/2011/11/21/sap-hana-and-the-game-changing-power-of-speed/>, accessed 02/2014.
- Faerber, F., Sang, K. C., Primsch, J., Bornhoevd, C., Sigg, S., and Lehner, W. 2011. "SAP HANA Database - Data Management for Modern Business Applications," *ACM SIGMOD Records* (40:4), pp. 45-51.
- Garcia-Molina, H., and K. Salem. 1992. "Main Memory Database Systems: An Overview," *IEEE Transactions on Knowledge and Data Engineering* (1992:6), pp. 509-516.
- Larsson, R. 1993. „Case Survey Methodology: Quantitative Analysis Patterns Across Case Studies,“ *Academy of Management Journal* (13:6), pp. 1515-1546.

- Loos, P., Lechtenböcker, J., Vossen, G., Zeier, A., Krüger, J., Müller, J., Lehner, W., Kossmann, D., Fabian, B., Günther, O., and Winter, R. 2011. "In-Memory Databases in Business Information Systems," *Business & Information Systems Engineering* (3:6), pp. 383-390.
- Loos, P., Strohmeier, S., Piller, G., and Schütte, R. 2012. "Comments on "in-Memory Databases in Business Information Systems"," *Business & Information Systems Engineering* (4:4), pp. 209-213.
- McCallum, J. C. 2013 Average Historic Price of RAM. Available at: <http://www.statisticbrain.com/average-historic-price-of-ram> , accessed 02/2013.
- Pezzini, M. 2013. Hype Cycle for In-Memory Computing Technology, 2013. Available at: <https://www.gartner.com/doc/2563015/hype-cycle-inmemory-computing-technology>, accessed 02/2014.
- Piller, G., and Hagedorn, J. 2011. "Business Benefits and Application Capabilities Enabled by In-Memory Data Management," in Proceedings des Workshops Innovative Unternehmensanwendungen mit In-Memory Data Management, Mainz, Germany, pp. 45-56.
- Plattner, H., and Zeier, A. 2011. In-Memory Data Management: An Inflection Point for Enterprise Applications, Springer, Heidelberg.
- Read, K. 2013. "Is "in-Memory" always the Right Choice?" *Business Intelligence Journal* (18:1), pp. 46-49.
- Saldana, J. 2013. The Coding Manual for Qualitative Researchers, Sage, London.
- vom Brocke, J., Debertoli, J., Müller, O., and Reuter, N. 2014. "How in-Memory Technology can Create Business Value: Insights from the Hilti Case," *Communications of the Association for Information Systems* (34:7), pp. 151-168.
- Webster, J., and Watson, R. T.. 2002. "Analyzing the Past to Prepare for the Future: Writing a Literature Review," *MIS Quarterly* (26:2), pp. xiii-xxiii.
- Wessel, P., Köffer, S., and Becker, J. 2013. "Auswirkungen von In-Memory Datenmanagement auf Geschäftsprozesse im Business Intelligence," in Wirtschaftsinformatik Proceedings 2013, Leipzig, Germany, pp. 1781-1795.
- Winter, R., Bischoff, S., and Wortmann, F. 2011. "Revolution Or Evolution? Reflections on In-Memory Appliances from an Enterprise Information Logistics Perspective," in Proceedings des Workshops Innovative Unternehmensanwendungen mit In-Memory Data Management, Mainz, Germany, pp. 23-34.