Communications of the Association for Information Systems

Volume 34 Article 7

1-2014

How In-memory Technology Can Create Business Value: Insights from the Hilti Case

Jan vom Brocke
University of Liechtenstein, jan.vom.brocke@uni.li

Stefan Debortoli
University of Liechtenstein

Oliver Müller
University of Liechtenstein

Nadine Reuter
University of Liechtenstein

Recommended Citation
vom Brocke, Jan; Debortoli, Stefan; Müller, Oliver; and Reuter, Nadine (2014) "How In-memory Technology Can Create Business Value: Insights from the Hilti Case," *Communications of the Association for Information Systems*: Vol. 34 , Article 7.
DOI: 10.17705/1CAIS.03407
Available at: https://aisel.aisnet.org/cais/vol34/iss1/7

Communications of the Association for Information Systems

CAIS

How In-memory Technology Can Create Business Value: Insights from the Hilti Case

Jan vom Brocke

University of Liechtenstein

Jan.vom.Brocke@uni.li

Stefan Debortoli

University of Liechtenstein

Oliver Müller

University of Liechtenstein

Nadine Reuter

University of Liechtenstein

Abstract:

With in-memory technology, all data and applications are kept in the computer's main memory to avoid expensive mechanical hard-drive I/O access, reduce latency times, and increase the ability to process large volumes of data or complex data. In this "innovation and novel concepts" article, we discuss how in-memory technology may create business value. Based on our experiences in collaborating with the Hilti Corporation, one of the first adopters of SAP's in-memory technology appliance (SAP HANA), we describe and discuss illustrative application scenarios that are made possible through the increased computing power offered by in-memory technology. Based on these scenarios, we identify principles of value creation through in-memory technology: the first-order effects of reduced latency times and increased ability to process large volumes of complex data (big data processing) that lead to the second-order effects of advanced business analytics and the convergence of OLTP and OLAP that themselves lead to business value through improved organizational performance.

Keywords: in-memory technology, IT business value, Big Data, OLTP, OLAP, advanced business analytics, real-time analytics, business intelligence, sentiment analysis

Volume 34, Article 7, pp. 151-168, January 2014

I. INTRODUCTION

In-memory technology is a major technological innovation in the field of enterprise systems. It is expected that 30 percent of all enterprises will have one or more critical applications running on in-memory database systems in five years [Dutta and Bilbao-Osorio, 2012]. Gartner [2011] identifies in-memory computing as one of the top-ten strategic technologies and describes it as a technology "with the potential for significant impact on the enterprise in the next three years" (para. 2). Plattner and Zeier [2011] go so far as to state that in-memory technology "marks an inflection point for enterprise applications" (p. 209) and that this innovation "will lead to fundamentally improved business processes, better decision making, and new performance standards for enterprise applications" (p. 2). However, although the technologies behind in-memory computing may be mature, analysts assert that there is a shortage of the kind of expertise that is required to realize business value from in-memory technology [Manyika et al., 2011]. Against this backdrop, the purpose of this article is to provide first insights into business value creation through in-memory technology.

The technological novelty of in-memory computing lies in the management of the data and applications that are kept in the computer's main memory in order to avoid expensive mechanical hard-drive I/O access [Chaudhuri, Dayal, and Narasayya, 2011]. This approach results in a significant increase of information processing capacity. For example, SAP reports increases in transaction speed by factors up to 100,000 for its in-memory computing appliance, HANA (High-performance Analytic Appliance). This increase in speed can be compared to improving the flight time from New York to Paris to 0.28 seconds or that of circling the world twice to less than three seconds [SAP, 2012a].

Apart from technical considerations, real-world application scenarios that demonstrate the tangible business value of in-memory technology are scarce. Increased transaction-processing speed alone might not pay off, considering the large investments needed to implement the new technology. We set up a collaborative research project with the Hilti Corporation, a leading supplier of products, systems, and solutions for construction professionals, to investigate how in-memory technology can create business value and whether in-memory technology represents a disruptive innovation in enterprise computing. In particular, we seek to identify concrete application scenarios using the context of Hilti in order to derive general principles for leveraging the potential of in-memory technology. As a result, we present an "innovation and novel concepts" article that discusses first observations that could foster the discourse on in-memory value creation in our community.

The remainder of this article is structured as follows: The next section provides the research background of this study. In particular, we describe the technical background of in-memory technology and review the literature on the technology-push-versus-needs-pull discussion. Then we describe the overall approach of our project, followed by a presentation of selected application scenarios, including product development, marketing, sales, and customer service. Building on these findings, we derive and discuss general principles concerning how to leverage in-memory technology in order to generate business value. We close with a conclusion and suggestions for future research.

II. RESEARCH BACKGROUND

Technology Push versus Need Pull

In innovation literature, one typically distinguishes between the concept of technology push (TP) and that of need pull (NP) [Langrish, 1972; Voss, 1984; Zmud, 1984]. Proponents of the technology push paradigm argue that technological developments are brought about by science and that new technological advances are the key drivers of innovation [Chidamber and Kon, 1994], while the pull argument suggests that customer need is the origin of technological developments, so markets, users, and applications are the key drivers of innovation [Chidamber and Kon, 1994]. Some studies (e.g., Langrish, 1972) have concluded that technology push and need pull coexist, but that the need pull model is more prevalent. Consequently, adoption of a new technology may be motivated by the recognition of a promising new technology (TP), a perceived performance gap in existing technologies (NP), or both [Chau and Tam, 2000].

We argue that in-memory technology is driven by both technology push and need pull. The need pull force is motivated by the emergence of "Big Data" and the desire to glean intelligence from it in order to gain competitive advantage [McAfee and Brynjolfsson, 2012]. New technologies are needed to collect vast amounts of complex data and analyze it within an economical timeframe, but in-memory technology is not required, as most of the Big Data processing today is achieved by leveraging massive distributed parallel computing across large-scale clusters of machines. One of the most prominent platforms for Big Data processing is Hadoop, which implements the fundamental programming model MapReduce [Dean and Ghemawat, 2008]. MapReduce is a batch query processor that is specifically designed for analyzing whole data sets in an ad hoc fashion [White, 2009], but not for transaction systems and real-time processing [Woods, 2012]. In contrast, in-memory technology addresses both of these aspects of data analysis in addition to providing high-performance analytical capabilities.

The technology-push force originates from the massive investments of top enterprise software vendors, such as SAP and Oracle. SAP, for instance, made the strategic decision to replace the traditional relational databases used to build the basis for their enterprise systems (e.g., ERP, CRM, DWH) with its self-developed in-memory database SAP HANA. SAP also actively markets its in-memory products over channels that include the SAP HANA Competitions [Nerney, 2013], the One Hundred Thousand Club [SAP, 2012a], and the SAP HANA Use Case Database [SAP, 2012b].

Still, proven in-memory technology business cases are lacking, in part because of how corporate users think about IT innovation. Darwin’s model of evolution as a slow stream of small mutations has heavily influenced our image of change processes [Gersick, 1991], and this pervasive model fits well many of today’s IT projects that focus on incremental improvements (e.g., changing to a new software release, improving data quality, implementing new reports). More radical, revolutionary IT innovations are rarely found on CIOs’ agendas [Lyytinen and Newman, 2008], which is natural enough, considering the complexity and riskiness of such projects. (At least this phenomenon was clearly at work in our case company, where IT users just did not dare to express some demands, such as real-time analytics, because they assumed that they were technically unrealistic.) In-memory technology breaks with existing paradigms and promises advances that are beyond the imagination of most IT professionals. Martin Petry, Hilti’s CIO, made the following analogy to overcome existing mental constraints and stimulate out-of-the-box thinking: “Imagine you can fly from Zurich to New York in less than a second. How would you plan your weekend with your spouse? You might actually consider visiting two to three continents as opposed to going shopping in the village close by.”

In-memory Technology

The technological foundations of in-memory computing were developed in the mid-1980s [Eich, 1987], but only recent developments in the area of computer hardware have made the use of these technologies economically feasible. As a result of these developments—primarily increases in main memory sizes and in computing power at affordable prices—many enterprise software vendors (e.g., IBM, Oracle, QlikTech, SAP, TIBCO) have begun building in-memory technology into their application systems. These new offerings challenge decision makers in IT departments and the software industry to determine to what extent in-memory-based applications add business value.

In-memory technology can be described in terms of five primary characteristics. First, data is stored entirely in the computer’s main memory instead of on disk-based hard drives [Knabke and Olbrich, 2011], which enables the dramatic decrease in the time required to access it [Garcia-Molina and Salem, 1992]. Because of larger main memory capacities and decreasing prices for the additional capacity, entire databases can be permanently shifted to servers’ main memory. Such systems are referred to as in-memory database management systems (IMDBMS) [Wust et al., 2012].

The second characteristic of in-memory technology is that multiple CPUs can process parallel requests, thereby using the available computing resources fully [Word, 2012]. The use of parallelism through many processing units allows algorithms to be split up across several computing units and many operations to be executed concurrently [Chaudhuri et al., 2011].

Third, a mixed row- and column-oriented storage approach is applied in IMDBMS instead of the row-based-only approach that is implemented in traditional relational database management systems (RDBMS). Row-oriented storage allows single tuples to be read quickly but is not well-suited to reading a set of results from a single column [Plattner and Zeier, 2011]. Column-oriented databases, on the other hand, are well-suited to calculations, such as aggregates, that are executed on a single or only a few columns [Word, 2012]. Plattner [2009] points out that column storage provides significant advantages for modern CPUs over traditional row-oriented RDBMS, as it improves data compression and allows for massive parallel processing and efficient memory access.

Fourth, IMDBMS offer compression techniques like dictionary-encoding and run-length-encoding that significantly reduce data size [Krueger, Grund, Tinnefeld, Eckart, Zeier, and Plattner, 2010]. Plattner [2009] demonstrates a compression from a 35 GB row-store database table size to an 8 GB column-store table size because of the more efficient vertical compression along columns.

Fifth, IMDBMS implement an insert-only approach, so a database does not allow applications to perform updates or deletions on physically stored tuples of data [Grund, Krueger, Tinnefeld, and Zeier, 2009]. Therefore, changes of tuples must be flagged as valid or invalid or by using time stamps either via point representation, which uses a single timestamp and stores only the differences between the new tuple and the old one or via interval representation, which stores the full tuple and any other information needed to determine how long it is valid [Plattner and Zeier, 2011, p. 110].

These five characteristics are key to a significant increase in data management performance. Studies that investigate the opportunities offered by this increased performance refer to the concepts of real-time computing and high-performance computing, both of which have been the subject of research in computer science for decades. The objective of high-performance computing is to minimize the average response time of a task, but this response time is often far from real time (e.g., DNA sequence mining). The objective of real-time computing is to meet the time requirement of a task [Stankovic, 1988]. The term *real* indicates that the reaction of a real-time system to external events must occur *during* those events [Buttazzo, 2011], although a distinction between *hard real-time* systems and *soft real-time* systems can be made [Stankovic, 1992]. While missing a task deadline in soft real-time systems leads "only" to a loss in business value, missing a task deadline in hard real-time systems often results in catastrophic consequences that define the total failure of a system (e.g., an air traffic control system).

Originating in business intelligence software products, such as QlikView in 1997, in-memory technology has been incorporated into the applications of most vendors of enterprise systems [Howarth, 2011]. One example is SAP's in-memory approach, HANA. Oracle first took a critical view of in-memory technology, with Oracle founder and CEO Larry Ellison describing the idea as "wacko" and "ridiculous" [Evans, 2010]. Today, however, even Oracle has entered the market with an in-memory-based solution called Oracle Exalytics In-Memory Machine.

In light of these market developments and the interest in finding additional business insights through Big Data initiatives, in-memory-based application systems have gained increasing relevance for decision makers in the field of enterprise information systems. Still, specific benefits in terms of economic value creation are not yet well understood, so CIOs remain concerned about the extent to which their organizations' information systems may benefit from an investment in in-memory technology. The same is true for software vendors, who seek to understand the principles of value creation through in-memory technology. Information systems research has long investigated the principles of value creation through IT (e.g., Brynjolfsson, 1993; Melville et al., 2004; vom Brocke, 2007, vom Brocke, Recker, and Mendling, 2010) and IT adoption (e.g., Karahanna, Straub, and Chervany, 1999; Rogers, 1995), so the purpose of our study is to contribute to both of these streams of IS research by investigating real-life application scenarios for in-memory-based application systems and to derive general conclusions from these first experiences.

III. PROJECT OVERVIEW

Hilti is a manufacturer of fastening technologies that supplies high-tech tools (e.g., power drills, demolition hammers, electric saws, cutters) and services to construction professionals around the globe. Hilti is an early adopter of new IT in general and is especially interested in the application of in-memory technology for its potential to answer some old questions and add layers of analytics to fashion entirely new insights [Goldberg, 2012]. Hilti was selected as a ramp-up partner to test SAP's in-memory computing appliance, SAP HANA, and we had the opportunity to observe the introduction of this technology. The main objective of this university-industry collaboration was to identify innovative application scenarios for in-memory technology and to evaluate these scenarios against the specific business process requirements of the Hilti Corporation.

Against this background, in the summer of 2012, we set up a team of IS master's students to explore potential application scenarios. By including students into the project team, we hoped to approach the task with a fresh perspective and to stimulate out-of-the-box thinking. The students' work was observed by university researchers and regularly challenged by Hilti's IT managers, including the CIO, to ensure the innovativeness and appropriateness of the students' ideas.

The project task—generating and evaluating innovative application scenarios—is a typical creative task, and therefore, we structured the project according to the four phases of the generic creative process [Lubart, 2001]: preparation, incubation, illumination, and verification. The aim of the preparation phase is to become familiar with the task and domain at hand and to gather relevant information. In the incubation phase, the information obtained in the preparation phase is processed, and some draft ideas for addressing the task are produced. In the illumination phase, multiple and disparate ideas for solving the task are generated. Finally, in the verification phase, the solutions generated during incubation and illumination are tested against the task requirements, and ideas are excluded, refined, and combined. In particular, this process included idea evaluation in terms of novelty, workability, relevance, and specificity [Dean, Hender, Rodgers, and Santanen, 2006].

Next we report on the results of two full iterations of this creative process. While the first iteration concentrated on explorative and divergent thinking, the second iteration focused on convergent thinking and exploitation of previously generated ideas. The iterations and phases of the project, the corresponding events/actions, actors, and information sources, and the number of generated ideas are summarized in Table 1.

Table 1: Overview of the Project

	Phases	Events/Actions	Actors	Information Sources	# of Ideas
Iteration 1	Preparation	Kick-off	CIO of Hilti, researchers, students	Project task, company information	-
		Literature review and Web research	Students	Literature, technical documentation, company information, information on current IT trends and developments	-
	Incubation	Informal discussions	Researchers, students	See Preparation	5
	Illumination	Brainstorming sessions	Students	See Preparation and Incubation	18
	Verification	Interim presentation and challenging of ideas	CIO of Hilti, Hilti experts (heads of Performance Competence Center, HR, Finance, and Reporting), researchers, students	Feedback from Hilti	9
Iteration 2	Preparation	Review of feedback from interim presentation	Researchers, students	Ideas of first iteration, feedback from interim presentation	9
	Incubation	Informal discussions	Students	See Preparation	9
	Illumination	Brainstorming sessions	Students	See Preparation and Incubation	12
	Verification	Final presentation and challenging of ideas	CIO of Hilti, Hilti experts (heads of Performance Competence Center, HR, Finance, and Reporting), researchers, students	Feedback from Hilti	5

The project kicked off with a meeting with Hilti's CIO, who presented the project task to the researchers and students and provided input on business and technical requirements. The kick-off was followed by thorough research on the topic of in-memory technology. As a traditional literature review revealed only a few academic sources on the topic (e.g., Krueger et al., 2010; Loos et al., 2011; Piller and Hagedorn, 2011), nonscientific outlets like magazines, websites, and blog posts were also reviewed to clarify and contextualize in-memory technology. One finding of this phase was that an in-memory database itself enables only limited performance gains unless it is combined with other infrastructural technologies for sensing and acting in order to leverage its full potential. Only swapping the database of an information system, all other components being equal or held constant, we may resolve one bottleneck but immediately run into another one. For example, if sensors or other input devices do not capture and forward data in real time, there is only a limited benefit of implementing a database capable of real-time processing of large volumes of data. Likewise, if end users do not have ready access to information systems, there is little need to update data on a sub-second scale.

Therefore, we examined other emerging information and communication technologies, such as those listed in the Gartner Strategic Technologies report [Gartner, 2011]. We identified four contemporary technologies or trends that could be combined with or are related to in-memory technologies: mobility, the Internet of Things, Big Data, and business analytics. These technologies served as inspiration for the first group brainstorming session with eighteen information systems students. The task was to brainstorm potential application scenarios by combining in-memory technology with one or more of the four IT trends. Then small teams challenged and discussed the generated ideas in light of Hilti's business requirements.

The resulting eighteen application scenarios were discussed among the students and qualitatively evaluated with regard to their novelty, workability, relevance, and specificity. Then the nine most promising scenarios were presented to Hilti IT in an interim presentation. The meeting revealed that the Hilti representatives saw the greatest

Communications of the Association for Information Systems

relevance of in-memory technology in the most customer-facing applications since customer-orientation is among the foremost strategic orientations of Hilti [vom Brocke et al. 2010].

The feedback from the interim presentation formed the basis for a second iteration of the process. The focus of this iteration was on generating new scenarios and refining old ones by combining the ideas created beforehand. This effort resulted in five final scenarios, one of which is fully described in the next section.

IV. APPLICATION SCENARIOS

We chose the four most illustrative application scenarios from Iteration 1—*Track 'n' Aid*, *Listening to the Web*, *Instant Knowledge Finder,* and *Talk to Me*—and one scenario from Iteration 2, the *Sales Area App*. Later, we use these scenarios in discussing the general principles of in-memory value creation that we derived from them.

Selected Scenarios from Iteration 1

Track 'n' Aid

The idea behind *Track 'n' Aid* is to capture the position, status, and usage data of tools in the field continuously in order to enhance customer service and improve product designs. Products equipped with sensors and communication devices (e.g., GSM, UMTS, Wi-Fi, and NFC) can transfer live data to Hilti's information systems [Chong and Kumar, 2003; Tilak, Abu-Ghazaleh, and Heinzelman, 2002] in order, for example, to anticipate maintenance work and prevent costly repair processes and interruptions of work at construction sites.

In-memory technology facilitates this scenario through its ability to process large amounts of data from several sensor networks in a timely manner. Assuming that about 5 million Hilti machines are located on customer sites around the globe, that each tool sends a status update once a day and usage data every fifteen minutes when it is being used, and that 50 percent of the tools are used each day for eight hours, around 100 million database entries will be generated per day. This data must be analyzed in a timely manner to ensure fast reaction to incidents.

Similar scenarios have already been implemented in other industries, such as the airline industry and the elevator industry (cf. Avery et al., 2010; Wang, Du, Zhang, and Li, 2011). The enormous amounts of data create a wide range of opportunities and challenges, as the IT director of Virgin Atlantic Airlines explains: "The challenge is what do you do with that amount of data when you are getting terabytes of data a day of your various airplanes? We are getting to the stage right now where we cannot deal with that much" [Finnegan, 2013]. In-memory databases are one approach to analyzing this data in a timely manner in order to ensure fast reaction to incidents.

Discussions with the Hilti managers identified three levels of implementing *Track 'n' Aid*. First, a simple geographic tracking of tools would provide Hilti visibility into the installed base and help customers locate and manage their tools efficiently. Second, the status and usage data of products and their components (e.g., batteries) would help predict the need for maintenance and repair work and enable Hilti to schedule and trigger maintenance and repair processes automatically. Third, the tools' lifecycle data is a rich source for data mining. Analyses of tools' log files can identify patterns of recurring incidents so that predictive actions can be planned and triggered, the overall life span of the products can be extended, and knowledge about improving product designs can be derived.

Listening to the Web

Listening to the Web relates to crawling the Internet, especially social media, in search of any kind of Hilti-related content and loading this information into a central in-memory database for further text mining. Through methods like named-entity recognition [Nadeau and Sekine, 2007] and sentiment analysis [Liu, 2012], opinions and other subjective information can be extracted from a constant stream of social media posts. The Hilti experts considered learning about customer opinions through this scenario highly important, as Hilti makes more than 200,000 customer contacts every day. Another advantage of social media mining is that the "voice of the customer" [Griffin and Hauser, 1993] can be analyzed without actively surveying customers.

Implementing *Listening to the Web* requires capturing and analyzing huge amounts of unstructured data in a timely manner, and in-memory technology provides the necessary capacity and speed. Even though this application scenario could be implemented with other Big Data processing platforms, such as Hadoop, in-memory technology leads to more timely results through real-time analytics [Bifet and Frank, 2010; Woods, 2012]. Timely results of customer opinions are essential if the company is to react quickly to negative comments and new business opportunities. For example, analyzing messages from Twitter requires that data mining algorithms run in real time in order to keep up with the continuous stream of messages, as an average of 100,000 messages are tweeted per minute [Bifet and Frank, 2010; Tepper, 2012].

In challenging the idea, we discussed additional application areas and found that the scenario can create value for Hilti along the whole supply chain, including product development, marketing, sales, logistics, and product service. For example, in product development, analyzing social media posts can identify opportunities for improving products and services, as the derived knowledge can illuminate specific customer needs that can lead to innovative product features. Such was the case when Hilti learned from customer contacts that the theft of tools from its customers' construction sites was an issue, and Hilti developed an innovative theft-protection feature that is now a standard feature of each valuable Hilti tool.

Instant Knowledge Finder

The scenario *Instant Knowledge Finder* is concerned with all parts of the knowledge management process, but the main advantage of the scenario lies in the identification and capture of existing knowledge [Alavi and Leidner, 2001].

A common obstacle to effective knowledge management is keeping the knowledge base current with newly captured knowledge [Alavi and Leidner, 1999; Zyngier, 2002]. In-memory technology can be used to extract information from various unstructured data sources as it is created. This information can then be stored and indexed in a central repository. Chaudhuri et al. [2011] refer to this approach as "enterprise search." The speed and data storage capacity of in-memory databases would allow Hilti to crawl Hilti's file systems constantly to extract concepts and relationships from textual data in order to build and maintain a structured knowledge base. Combined with the sub-second query response time offered by in-memory technology, this capacity would create a Google-like user experience when employees search for internal information. In addition to retrieving the information itself, the *Instant Knowledge Finder* can analyze who has created, updated, downloaded, or forwarded a document and display these connections in order to connect individuals with similar expertise or interests.

As in the *Listening to the Web* scenario, in-memory technology's ability to process large amounts of unstructured data is the major reason to deploy it in this scenario. The ultra-fast query response time of in-memory database systems is another relevant factor for end users.

The *Instant Knowledge Finder* can be applied in many areas of the organization, but the Hilti experts were especially interested in its application in sales and product development. Through personalized information filters, sales representatives can receive all sales-related information related to a product, such as the latest sales figures and customer complaints. In addition, the *Instant Knowledge Finder* recognizes when more than one sales representative is searching for similar information and can connect the individuals to one another. Product developers can receive the latest product-related information, such as reported failures. For example, if an employee from technical service enters a new repair case into the system, all relevant product developers would receive the information when they search for the corresponding product.

However, the *Instant Knowledge Finder* might also lead to unintended consequences, such as information waste [Wijnhoven, Dietz, and Amrit, 2012], as the ability to discover and retrieve data at zero cost might lead to the generation of unnecessary or useless data. This information waste might cause employee information overflow and/or have a negative effect on the energy needs of information systems, which are directly related to greenhouse gas emissions [Wijnhoven, Dietz, and Amrit, 2012]. Here, additional effort is required in order to ensure a useful, non-redundant, and understandable knowledgebase.

Talk to Me

Talk to Me suggests replacing the traditional mouse-and-keyboard-based input methods for data access with voice control [Rabiner and Juang, 1993]. Speech recognition per se does not require in-memory technology, but combined with the processing speed and ability to handle unstructured data of in-memory technology, speech recognition reduces the query response time significantly, leading to a conversation-like human-computer interaction.

One possible field of application that Hilti identified was the use of *Talk to Me* during executive meetings. Currently, all information that might be relevant to a meeting must be prepared prior to the meeting by requesting reports from various enterprise systems. With *Talk to Me*, the data can be accessed during the meeting and interactively analyzed through verbal commands. As a result, it is expected that fewer decisions will have to be postponed because of missing information during meetings.

Another application is access to data from "the road." Because of Hilti's direct sales model [vom Brocke, Petry, Sinnl, Kristensen, and Sonnenberg, 2010], more than two-thirds of its 20,000 employees work as traveling sales representatives, driving from one customer to another. The sales representatives could use the time in the car to access the latest customer briefings, sales figures, or product information by querying the system through voice control.

Selected Scenario from Iteration 2

The *Sales Area App* is a combination of multiple ideas generated during the first iteration. The application is intended to support salespeople in the field with customer-related information, such as customers in the region, nearby construction sites, location and status of Hilti tools on these sites, and social media posts from people nearby. This may all be accessible in one application.

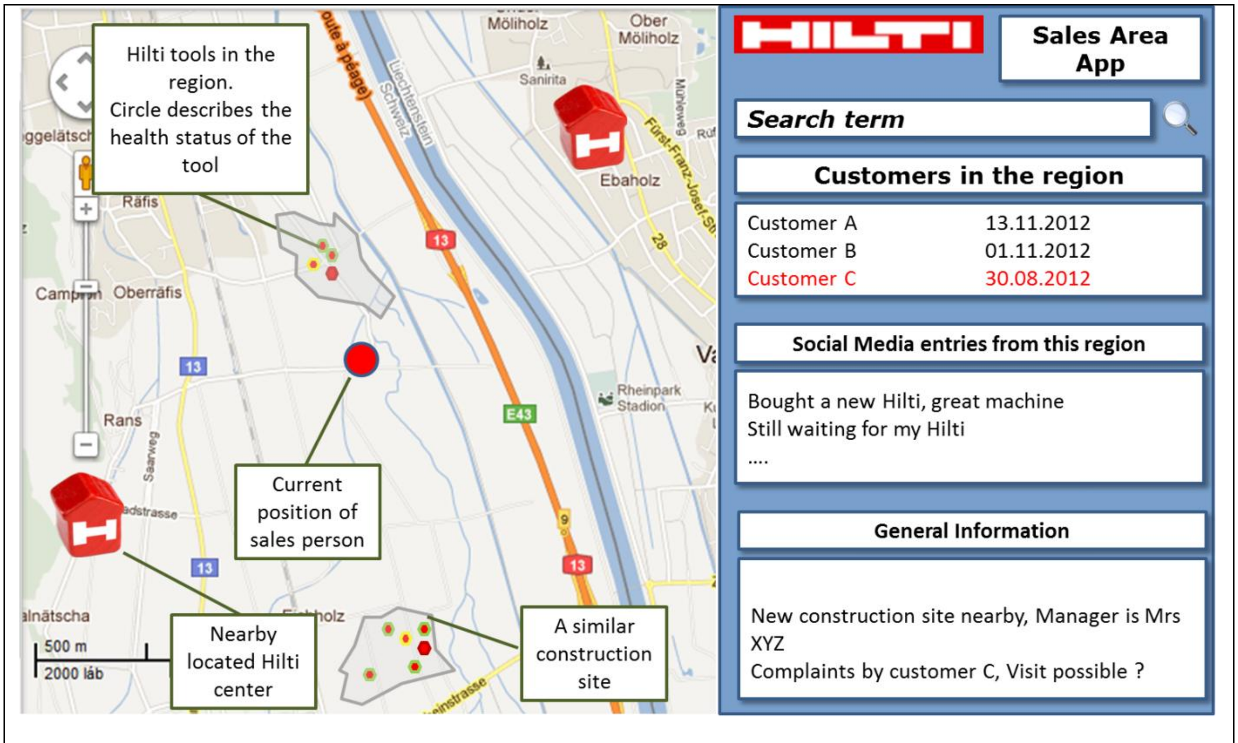

Figure 1. Mock-up of the Sales Area App Cockpit

The *Sales Area App* has two parts (cf. Figure 1). One part of the screen shows a map with sales-related points of interest, and another displays the salesperson's current position. The sales representatives also use the map to find nearby service centers, shops (called "Hilti centers" in this example), and customers (construction sites). In addition to the position of the customer or the construction site, the salesperson can see where Hilti products are being used (cf. *Track 'n' Aid*). Each product is displayed by a colored circle that presents the operational status of the tool, that is, whether it has problems or not. When the salesperson selects one of the products, detailed lifecycle data about the product is displayed, such as the owner, the date of purchase, and the product's repair history.

Other parts of the screen display detailed information about a geographical region, listing customers in the region and information about the last visits to these customers. Customers who have not been visited for a defined time span are flagged, and by clicking on the customer name, the salesperson is linked to the customer details based on the information stored in the customer relationship management system. The application also searches social media networks for relevant posts from the region (cf. *Listening to the Web*).

The *Sales Area App* can be operated through voice control (cf. *Talk to Me*). For example, the salesperson may ask, "Who are the nearest customers?" to which the application may reply that the nearest customer is customer BigBuild, who was last visited a month ago and who has one product that needs maintenance (cf. *Track 'n' Aid*). The *Sales Area App* checks the availability of the required parts, places an order at the Hilti center, and plans the salesperson's driving route via the center to the customer. After visiting the Hilti center and the customer, the sales representative may see on the map that a new construction site is being established a short distance away, so the representative places a query about the project manager in charge, and the system returns his or her latest projects, projects with Hilti, and colleagues at Hilti who were in charge of these projects and may have additional background information on the person (cf. *Instant Knowledge Finder*).

Since the *Sales Area App* pushes information to salespeople according to their location, new business opportunities can be generated when, for example, maintenance services can be offered proactively. Managers at Hilti perceived that this scenario could generate customer excitement and competitive advantage.

The *Sales Area App* requires that huge amounts of data be processed from several heterogeneous sources, including sensor networks and social media platforms. In addition, some of the processed data is structured and some is unstructured, and both must be analyzed instantly in order to provide needed feedback to the sales representatives. Therefore, the Hilti experts consider in-memory technology a major prerequisite to implementing the scenario. The scenario also illustrates that value creation through in-memory technology might be driven by the right combination of simple application scenarios and technologies.

An overview of all applied application scenarios and the necessity of using an in-memory application system is illustrated in Table 2. We discuss the more general propositions we drew from our study in the next section.

Table 2: Overview of Selected Scenarios and the Need for In-memory Technology

Scenario	Need for processing large data volumes	Need for processing complex data	Need for low latency time
Track 'n' Aid	X		X
Listening to the Web	X	X	X
Instant Knowledge Finder	X	X	X
Talk to Me		X	X
Sales Area App	X	X	X

V. FINDINGS

Studies about the organizational impact of information technology provide various views on the phenomenon (e.g., Bakos, 1987; Bharadwaj, 2000; Brynjolfsson and Hitt, 2003; Melville et al., 2004; Soh and Markus, 1995; Tallon, Kraemer, and Gurbaxani, 2000). Business value through IT systems can refer to the impact of IT on organizational performance, including productivity enhancements, profitability improvements, cost reduction, competitive advantage, inventory reduction, and other measures of performance [Melville et al., 2004]. Bakos [1987] observes three areas for impact research: the impact of information technology on organizational performance, the impact of information technology on organizational structure and processes (including socio-technical elements), and the impact of organizational process changes on organizational performance.

Based on our study at Hilti, we derived four general principles of value creation through in-memory technology, which we discuss against the background of the model described in Figure 2. Figure 2 was derived from Bakos [1987] by mapping "IT" against the five in-memory technology characteristics and "Organizational Performance" against the list of performance characteristics as described by Melville et al. [2004]. Bakos's "Structures and Processes" are represented as a non-exhaustive set of value-adding effects that were observed in our Hilti-specific application scenarios.

Principle (1): The technical characteristics of in-memory technology provide first-order and second-order effects.

Our study suggests distinguishing different levels of effect of in-memory technology. As a first-order effect (Barua, Kriebel, and Mukhopadhyay, 1995), we observe an increase in the information-processing capability provided by in-memory technology, such as increased speed of single calculations or transactions. However, this increase in information processing capability may have limited direct value-creating effects. Critical statements in the workshops included, "Where is the benefit? I used to wait three seconds for a calculation. Now I wait one millisecond. Who cares?" or "We used to provide sales reports in the morning at 7 AM, now we would be able to provide them at 3 AM or much earlier. But who would care. No sales representative would wake up three hours earlier just to see these reports." In order to investigate more deeply into organizational performance effects, we have to focus on what possibilities arise from this increase in information processing capacity. We refer to these effects as the second-order effects of in-memory technology.

Principle (2): The second-order effects of in-memory technology manifest in advanced business analytics and the convergence of online transaction processing (OLTP) and online analytical processing (OLAP).

Advanced business analytics includes techniques like data mining and process mining [Chen, Chiang, and Storey, 2012]. These techniques were extant prior to the emergence of in-memory technology. Still, new scenarios such as the mining of streaming, unstructured, or noisy data (e.g., from sensor networks, social networks, event logs) poses new requirements to analytical models and technologies. Through increased information processing capacity, in-memory technology facilitates the creation of statistical models that consider an unprecedented number of inputs and allows analyses to be performed in a timely manner. An example can be seen in the *Instant Knowledge Finder* scenario, where text mining algorithms are applied when crawling large amounts of complex enterprise content to process queries for knowledge sources.

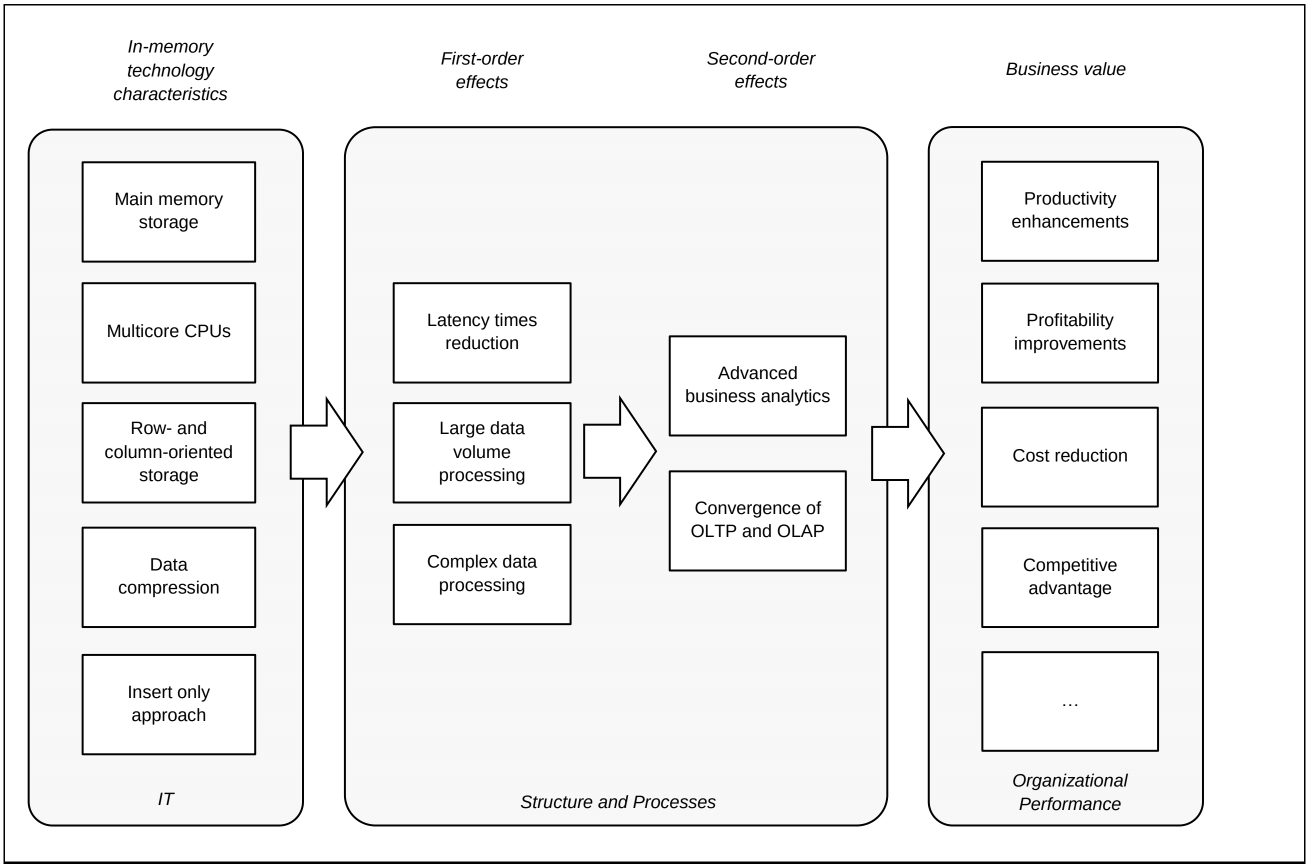

Figure 2. Generating Business Value Through In-memory Technology

Plattner [2009] states that, historically, the introduction of data warehouses for running analytical applications was a compromise that was needed for performance reasons only. Likewise, Färber et al. [2011] state that "it is no longer reasonable to continue the classical distinction between transactional and analytical access patterns" (p. 45). In-memory technology allows OLAP and OLTP systems to be converged by eliminating the information latency of conventional ETL (extract, transfer, load) processes and enables analyses of "historical" data at the time a transaction is performed [Janiesch, Matzner, and Müller, 2012]. This phenomenon is also called embedded analytics [Nijkamp and Oberhofer, 2009], operational business intelligence [Marjanovic, 2007], or real-time business intelligence [Chaudhuri et al., 2011]. A concrete OLTP/OLAP convergence example from our case is the *Track 'n' Aid* scenario, in which preventive maintenance and repair processes can be automatically scheduled and triggered by analysis of historical tool-usage data.

Principle (3): The second-order effects of in-memory technology are driven by reduced latency times and the ability to process large volumes of complex data.

We identify three first-order effects as drivers of second-order effects: the reduction of latency times, the ability to process large data volumes, and the ability to process complex data. The reduction of latency times is illustrated in Figure 3 (cf. zur Mühlen and Shapiro, 2010), which shows the time required to react to an event. The figure distinguishes the time required to store the data (data latency), to analyze the data (analysis latency), to make a decision based on that data (decision latency), and to take appropriate action (implementation latency). In addition, the figure expresses a negative correlation between the time required to react and business value. Against this background, in-memory technology can be understood as a way to shrink latency times and reduce the value lost. Two mechanisms that facilitate this effect can be distinguished: real-time computing and high-performance computing. While real-time computing allows the data of an economic event to be captured when it occurs (data latency), high-performance computing increases the speed of the data analysis process (analysis latency). We found examples of both mechanisms in our study, but we also found that real-time scenarios required that in-memory technology be combined with other data-capturing technologies, such as sensor networks, Web crawlers, or voice recognition. We also found scenarios in which we observed the possibility to reduce several types of latency times at once. In automated processes, such a combination enables dynamic process control based on real-time information about the process context. The *Sales Area App* scenario provides a number of examples of such context-aware process control, including real-time capture and storage of product health data and automatic scheduling of a sales representative's visit if issues arise. An analysis of system log files helps measure the various latency times in a productive system environment.

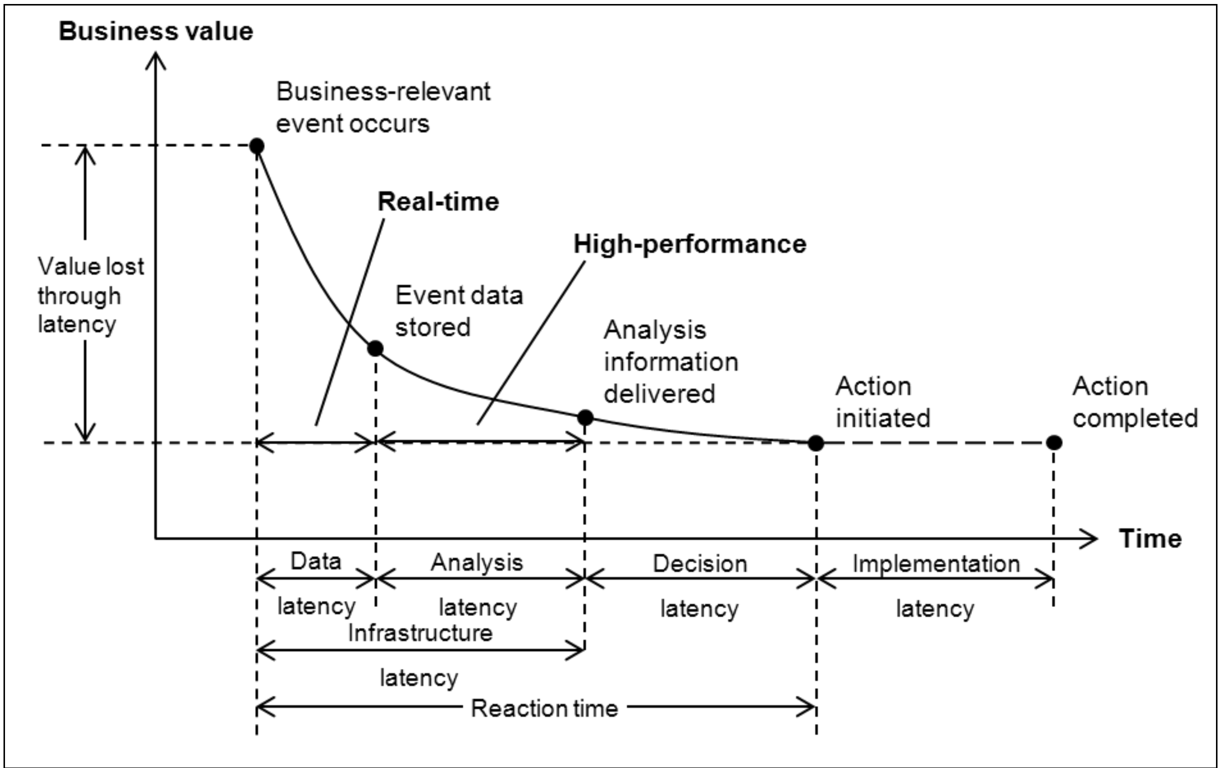

Figure 3. Latency Types (cf. Zur Mühlen and Shapiro, 2010)

In-memory technology also allows large volumes of data and complex data formats to be processed. Because of the technical characteristics of in-memory technology, larger and more complex data sets can be processed in a given time, so data can be taken into account that has not been considered before. Data can be qualified as complex if it has one or more of five characteristics: multi-format, multi-structure, multi-source, multi-modal, and/or multi-version [Darmont, Boussaid, Ralaivao, and Aouiche, 2005]. The scenarios *Track 'n' Aid* and *Listening to the Web* illustrate this opportunity. Because of the increased capacity, large data sets, such as those on the health status of products, can be measured and applicable services offered. Similarly, more complex information, such as user comments, can be incorporated into planning and decision-making processes. These larger and more complex data sets have at least two types of effects: First, business process management can consider an entirely new set of factors (e.g., news feeds, tweets, sensor data) to monitor and control processes. Second, process mining [van der Aalst, 2012] can use richer log data to identify patterns in process execution that may shed new light on how to manage processes. Therefore, we expect increasing numbers of refined planning algorithms and business models to be designed that make use of data of various types.

Principle (4): The value creation through in-memory technology is restricted by the capabilities of the overall socio-technical structures and processes.

We understand that any increase of an application system's information processing capacity must be supported by or is constrained by the socio-technical structures and processes the application system serves [Trist, 1981]. For instance, in the *Track 'n' Aid* scenario, analytics may be able to anticipate the spare parts needed to maintain tools, but in order to create value through these analytics, the supply chain must be sufficiently fast and agile to initiate and complete appropriate actions. Discussions with the industry experts from Hilti also demonstrated that, in many cases, a mature IT landscape, including high-quality data, harmonization of diverse internal data sources, and information-sharing along the supply chain is required in order to realize the full potential of in-memory technology.

We also found application scenarios in which the realization of benefits will be challenging because roles and responsibilities are not clearly defined, processes are characterized by improvisation and ad-hoc management, or human judgment and decision-making are required. Before in-memory technology can leverage its full potential in such situations, standards must be established in order to reduce variations, or governance structures must be adapted, for example, to allow for more autonomous decision making. In addition, many scenarios involve multiple institutions that must be considered; sharing the available or required data with all partners along the supply chain requires a strong relationship of trust, which can be seen as a game theoretical problem [Wijnhoven, Schuur, and Timmer, 2010]. If the business strategy does not embrace such strong partnerships, many supply chain scenarios may not be applicable. In any case, we may conclude that a careful IT-business alignment is needed to choose the application scenario that is most valuable for a specific organizational context.

VI. DISCUSSION

Our study was motivated by a practical consideration, to investigate how in-memory technology can create business value. In addition to IT value creation research, the diffusion of in-memory technology also relates to wider fields of IS research, which we discuss in this section using two examples.

In-memory databases can be viewed as an instance of organizational memory information systems (OMIS), that is, systems that “provide a means by which knowledge from the past is brought to bear on present activities, thus resulting in increased levels of effectiveness for the organization” [Stein and Zwass, 1995]. In-memory databases offer an unprecedented technological basis for the mnemonic functions of OMIS (information acquisition, retention, maintenance, search, and retrieval). Conceptualizing in-memory databases as OMIS leads to the proposition that they bear the potential to foster experiential organizational learning [Argyris and Schön, 1978]. By capturing an increased amount of and increased richness of internal and external stimuli and reducing the latency between when an experience occurs and memorization of that experience, they can speed up single-loop learning [Argyris and Schön, 1978]. While this theoretically grounded argument sounds appealing, it remains a hypothesis. Numerous researchers have underscored the importance of reflection—the process of stepping back from an experience to contemplate its meaning [Daudelin, 1997]—for organizational learning. Therefore, in-memory technology may also inhibit organizational learning by eliminating buffers, such as reflection and sense-making, that are important for cognitive processes. We propose that future empirical studies investigate both of these standpoints.

That said, we need to examine not only the new action possibilities emerging from this new technology, but also its limitations and constraints that result from the structures and processes of the overall socio-technical system, as illustrated in Principle (4). With in-memory technology the technical subsystem, or parts of it, provides increased capabilities, but the social subsystem in particular must adapt to these capabilities. For example, the increased performance of the technology may lead to unintended effects on people when the ubiquitous availability of massive amounts of real-time data may leads to phenomena like technostress, resulting in perceived work overload, demoralization, frustration, information fatigue, loss of motivation, and dissatisfaction at work [Ayyagari, Grover, and Purvis, 2011; Ragu-Nathan, Tarafdar, and Ragu-Nathan, 2008; Riedl, 2013]. Ragu-Nathan et al. [2008] identify factors like constant connectivity, simultaneity of streams of information, the need to learn new applications, rapid change of functionalities, constant customization and modification to systems, and multitasking as creators of technostress. It is easy to recognize these technostress creators in the elements of the model in Figure 2 and in the application scenarios outlined in Section 4. Hence, we propose that future studies use the concept of technostress as a theoretical lens through which to examine the potential for negative impacts of in-memory technology on end users and organizational performance.

VII. CONCLUSION

This “innovation and novel concepts” article contributes to IT value creation research by providing exemplary application scenarios and general principles for leveraging the potential of in-memory technology in business contexts. The study goes beyond the identification of application scenarios by discussing general principles of value creation through in-memory technology. We distinguish between first- and second-order effects of value creation, discuss drivers of second-order effects, and identify four principles for realizing the value proposition of in-memory technology. These findings help to clarify the potential of in-memory technology for creating business value and describe a set of required actions to leverage its full potential. In short, our study uncovers some essential principles of business value creation through in-memory technology that can inform in-memory adoption in practice. In addition, we set out first directions for future research.

Our findings should be viewed in light of the study’s limitations. The study was a collaborative university-industry project with the goal of generating application scenarios of in-memory technology and evaluating these scenarios against the specific business process requirements of the Hilti Corporation. Hence, the identified application scenarios are based on Hilti-specific requirements and manufacturing-related circumstances. In addition, the list of identified scenarios is neither exhaustive nor technically specified. As our findings are based on a single project, we suggest that future research collect additional data from companies in other industries in order to revise and refine our findings and yield more generalizable results.

While technological innovations, such as in-memory technology, clearly offer new possibilities for the design, use, and impact of information systems, substantial challenges are also apparent, particularly as they relate to the organizational and social aspects of the technology in use. These challenges make in-memory technology an important topic for information systems research in many ways, and with the focus on socio-technical systems design, information systems is the discipline to take on these challenges. Such contributions will be rewarding for both theory and practice.

REFERENCES

Editor's Note: The following reference list contains hyperlinks to World Wide Web pages. Readers who have the ability to access the Web directly from their word processor or are reading the article on the Web, can gain direct access to these linked references. Readers are warned, however, that:

1. These links existed as of the date of publication but are not guaranteed to be working thereafter.
2. The contents of Web pages may change over time. Where version information is provided in the References, different versions may not contain the information or the conclusions referenced.
3. The author(s) of the Web pages, not AIS, is (are) responsible for the accuracy of their content.
4. The author(s) of this article, not AIS, is (are) responsible for the accuracy of the URL and version information.

Alavi, M., and D.E. Leidner (1999) "Knowledge Management Systems: Issues, Challenges, and Benefits", *Communications of the Association for Information Systems Articles,* (1)7, pp. 1–37.

Alavi, M., and D.E. Leidner (2001) "Review: Knowledge Management and Knowledge Management Systems: Conceptual Foundations and Research Issues", *MIS Quarterly,* (25)1, pp. 107–136.

Argyris, C., and D. Schön (1978) *Organizational Learning: A Theory of Action Approach,* Reading, MA: Addison Wesley.

Avery, R.L., W. Cronie, E. Fromm, P. Rhodes, C. Khera, P.M. Kirkish, D.R. Leonhardi, J.P. Maloney, M. Nomi, and M. Bueser (2010) "Integrated Maintenance and Materials Services for Fleet Aircraft Using Aircraft Data to Improve Quality of Materials", U.S. Patent 7,689,329, pp. 1–22.

Ayyagari, R., V. Grover, and R. Purvis (2011), "Technostress: Technological Antecedents and Implications", *MIS Quarterly,* (35)4, pp. 831–858.

Bakos, J.Y. (1987) "Dependent Variables for the Study of Firm and Industry-level Impacts of Information Technology", *Proceedings of the Eighth International Conference on Information Systems,* pp. 10–23.

Barua, A., C.H. Kriebel, and T. Mukhopadhyay (1995) "Information Technologies and Business Value: An Analytic and Empirical Investigation", *Information Systems Research,* (3)1, pp. 3–23.

Bharadwaj, A.S. (2000) "A Resource-based Perspective on Information Technology Capability and Firm Performance: An Empirical Investigation", *MIS Quarterly,* (24)1, pp. 169–196.

Bifet, A., and E. Frank (2010) "Sentiment Knowledge Discovery in Twitter Streaming Data", *Proceedings of the 13th International Conference on Discovery Science,* pp. 1–15.

Brynjolfsson, E. (1993) "The Productivity Paradox of Information Technology: Review and Assessment", *Communications of the ACM,* (36)12, pp. 66–77.

Brynjolfsson, E., and L.M. Hitt (2003) "Computing Productivity: Firm-Level Evidence", *Review of Economics and Statistics,* (85)4, pp. 793–808.

Buttazzo, G. (2011) *Hard* Real-time *Computing Systems: Predictable Scheduling Algorithms and Applications, third edition,* New York: Springer.

Chau, P.Y.K., and K.Y. Tam (2000) "Organizational Adoption of Open Systems: A 'Technology-push, Need-pull' Perspective", *Information & Management,* (37)5, pp. 229–239.

Chaudhuri, S., U. Dayal, and V. Narasayya (2011) "An Overview of Business Intelligence Technology", *Communications of the ACM,* (54)8, pp. 88–98.

Chen, H., R.H.L. Chiang, and V.C. Storey (2012) "Business Intelligence and Analytics: From Big Data to Big Impact", *MIS Quarterly,* (36)4, pp. 1165–1188.

Chidamber, S.R., and H.B. Kon (1994) "A Research Retrospective of Innovation Inception and Success: The Technology-push, Demand-pull Question", *International Journal of Technology Management,* (9)1, pp. 94–112.

Chong, C., and S. Kumar (2003) “Sensor Networks: Evolution, Opportunities, and Challenges”, *Proceedings of the IEEE,* (91).

Conn, S.S. (2005) “OLTP and OLAP Data Integration: A Review of Feasible Implementation Methods and Architectures for Real Time Data Analysis”, *Proceedings of the IEEE Southeast Conference,* pp. 515–520.

Darmont, J., O. Boussaid, J.-C. Ralaivao, and K. Aouiche (2005) “An Architecture Framework for Complex Data Warehouses”, *7th International Conference on Enterprise Information Systems,* pp. 370–373.

Daudelin, M.W. (1997) “Learning from Experience Through Reflection”, *Organizational Dynamics,* (24)3, pp. 36–48.

Dean, D.L., J.M. Hender, T.L. Rodgers, and E.L. Santanen (2006) “Identifying Quality, Novel, and Creative Ideas: Constructs and Scales for Idea Evaluation”, *Journal of the Association for Information Systems,* (7)10, pp. 646–699.

Dean, J., and S. Ghemawat (2008) “MapReduce: Simplified Data Processing on Large Clusters”, *Communications of the Association for Information Systems*, (51)1, pp. 107–113.

Dutta, S., and B. Bilbao-Osorio (2012) “The Global Information Technology Report 2012: Living in a Hyperconnected World”, Geneva, Switzerland: World Economic Forum and INSEAD.

Eich, M. (1987) “MARS: The Design of a Main Memory Database Machine”, *Proceedings of 5th International Workshop on Database Machines,* pp. 332–339.

Evans, B. (2010) “Global CIO: SAP’s Hasso Plattner on Databases And Oracle’s Larry Ellison”, *InformationWeek,* http://www.informationweek.com/news/global-cio/interviews/224700913 (current March 7, 2013).

Färber, F., S.K. Cha, J. Primsch, C. Bornhövd, S. Sigg, and W. Lehner (2011) “SAP HANA Database—Data Management for Modern Business Applications”, *ACM SIGMOD Record,* (40)4, pp. 45–51.

Finnegan, M. (2013) “Boeing 787s to Create Half a Terabyte of Data per Flight, Says Virgin Atlantic”, *Computerworld UK,* http://www.computerworlduk.com/news/infrastructure/3433595/boeing-787s-create-half-terabyte-of-data-per-flight-says-virgin-atlantic/ (current March 7, 2013).

Garcia-Molina, H., and K. Salem (1992), “Main Memory Database Systems: An Overview”, *IEEE Transactions on Knowledge and Data Engineering,* (4)6, pp. 509–516.

Gartner. (2011) “Gartner Identifies the Top 10 Strategic Technologies for 2012”, http://www.gartner.com/it/page.jsp?id=1826214 (current March 7, 2013).

Gersick, C. (1991) “Revolutionary Change Theories: A Multilevel Exploration of the Punctuated Equilibrium Paradigm”, *Academy of Management Review,* (16)1, pp. 10–36.

Goldberg, M. (2012) “Early SAP HANA Adopters Describe Analytics Possibilities and Performance”, *Data Informed,* http://data-informed.com/early-sap-hana-adopters-describe-analytics-possibilities-and-performance/ (current March 7, 2013).

Griffin, A., and J. Hauser (1993) “The Voice of the Customer”, *Marketing Science,* (12)1, pp. 1–27.

Grund, M., J. Krueger, C. Tinnefeld, and A. Zeier (2009) “Vertical Partioning in Insert-only Scenarios for Enterprise Applications”, *16th International Conference on Industrial Engineering and Engineering Management,* pp. 760–765.

Howarth, B. (2011) “In-memory Computing”, CIO, http://www.cio.com.au/article/373945/in-memory_computing/ (current February 28, 2013).

Janiesch, C., M. Matzner, and O. Müller (2012) “Beyond Process Monitoring: A Proof-of-concept of Event-driven Business Activity Management”, *Business Process Management Journal,* (18)4, pp. 625–643.

Karahanna, E., D.W. Straub, and N.L. Chervany (1999) “Information Technology Adoption Across Time: A Cross-sectional Comparison of Pre-adoption and Post-adoption Belief”, *MIS Quarterly,* (23)2, pp. 183–213.

Knabke, T., and S. Olbrich (2011) “Towards Agile BI: Applying In-memory Technology to Data Warehouse Architectures”, *Proceedings des Workshops Innovative Unternehmensanwendungen mit In-memory Data Management,* pp. 101–114.

Krueger, J., M. Grund, C. Tinnefeld, B. Eckart, A. Zeier, and H. Plattner (2010) “Hauptspeicherdatenbanken für Unternehmensanwendungen”, *Datenbank-Spektrum,* (10)3, pp. 143–158.

Langrish, J. (1972) *Wealth from Knowledge: A Study of Innovation in Industry,* New York: Wiley, Halstead Press Division.

Liu, B. (2012) "Sentiment Analysis and Opinion Mining", *Synthesis Lectures on Human Language Technologies,* (5)1, pp. 1–167.

Loos, P., J. Lechtenbörger, G. Vossen, A. Zeier, J. Krüger, J. Müller, W. Lehner, D. Kossmann, B. Fabian, O. Günther, and R. Winter (2011) "In-memory Databases in Business Information Systems", *Business & Information Systems Engineering,* (3)6, pp. 389–395.

Lubart, T.I. (2001) "Models of the Creative Process: Past, Present and Future", *Creativity Research Journal,* (13)3, pp. 295–308.

Lyytinen, K., and M. Newman (2008) "Explaining Information Systems Change: A Punctuated Socio-technical Change Model", *European Journal of Information Systems,* (17)6, pp. 589–613.

Manyika, J., M. Chui, B. Brown, J. Bughin, R. Dobbs, C. Roxburgh, and A.H. Byers (2011) "Big Data: The Next Frontier for Innovation, Competition, and Productivity", *McKinsey Global Institute,* (June).

Marjanovic, O. (2007) "The Next Stage of Operational Business Intelligence: Creating New Challenges for Business Process Management", *Proceedings of the 40th Hawaii International Conference on System Sciences.*

McAfee, A., and E. Brynjolfsson (2012) "Big Data: The Management Revolution", *Harvard Business Review,* (1), pp. 59–68.

Melville, N., K. Kraemer, and V. Gurbaxani (2004) "Review: Information Technology and Organizational Performance: An Integrative Model of IT Business Value", *MIS Quarterly,* (28)2, pp. 283–322.

Nadeau, D., and S. Sekine (2007) "A Survey of Named Entity Recognition and Classification", *Lingvisticae Investigationes,* (30)1, pp. 3–26.

Nerney, C. (2013) "SAP Updates University Alliances Program to Include Big Data Analytics", *Data Informed,* http://data-informed.com/saps-updates-university-alliances-program-to-include-big-data-analytics/ (current March 7, 2013)

Nijkamp, E., and M. Oberhofer (2009) "Embedded Analytics in Front Office Applications", *Proceedings of BTW,* pp. 449–459.

Piller, G., and J. Hagedorn (2011) "Business Benefits and Application Capabilities Enabled by In-memory Data Management", *Proceedings des Workshops Innovative Unternehmensanwendungen mit In-memory Data Management* (IMDM2011), pp. 45–56.

Plattner, H. (2009) "A Common Database Approach for OLTP and OLAP Using an In-memory Column Database", *Proceedings of the 35th SIGMOD International Conference on Management of Data—SIGMOD '09,* pp. 1–2.

Plattner, H., and A. Zeier (2011) *In-memory Data Management: An Inflection Point for Enterprise Applications,* Heidelberg: Dordrecht, and London, New York: Springer Verlag.

Rabiner, L., and B.-H. Juang (1993) *Fundamentals of Speech Recognition,* Englewood Cliffs, NJ: PTR Prentice-Hall, Inc.

Ragu-Nathan, T.S., M. Tarafdar, M., and B.S. Ragu-Nathan (2008) "The Consequences of Technostress for End Users in Organizations: Conceptual Development and Empirical Validation", Information Systems Research, (19)4, pp. 417–433.

Riedl, R. (2013) "On the Biology of Technostress: Literature Review and Research Agenda", *The DATA BASE for Advances in Information Systems,* (44)1, pp. 18–55.

Rogers, E.M. (1995) "Diffusion of Innovations" in Salwen, M.B., and D.W. Stacks (eds.) *An Integrated Approach to Communication Theory and Research, Volume 65,* New York: Free Press.

SAP (2012a) "The "One Hundred Thousand" Club", http://www.experiencesaphana.com/blogs/experts/2012/02/06/the-one-hundred-thousand-club (current March 7, 2013).

SAP (2012b) "SAP HANA Use Cases", http://www.saphana.com/community/implement/use-cases (current March 7, 2013).

Soh, C., and M.L. Markus (1995) "How IT Creates Business Value: A Process Theory Synthesis", *International Conference on Information Systems,* pp. 29–41.

Stankovic, J. (1988) "Misconceptions About Real-time Computing", *IEEE Computer,* (21)10, pp. 10–19.

Stankovic, J. (1992) "Distributed Real-time Computing: The Next Generation", *Journal of the Society of Instrument and Control.*

Stein, E., and V. Zwass (1995) "Actualizing Organizational Memory with Information Systems", *Information Systems Research,* (6)2, pp. 85–117.

Tallon, P., K.L. Kraemer, and V. Gurbaxani (2000) "Executives' Perceptions of the Business Value of Information Technology: A Process-oriented Approach", *Journal of Management Information,* pp. 145–173.

Tepper, A. (2012) "How Much Data Is Created Every Minute?", mashable.com, http://mashable.com/2012/06/22/data-created-every-minute/ (current April 8, 2013)

Tilak, S., N.B. Abu-Ghazaleh, and W. Heinzelman (2002) "A Taxonomy of Wireless Micro-sensor Network Models", *ACM SIGMOBILE Mobile Computing and Communications Review,* (6)2, pp. 28–36.

Trist, E. (1981) "The Evolution of Socio-technical Systems", Occasional Paper, No. 2.

van der Aalst, W. (2012) "Process Mining: Overview and Opportunities", *ACM Transactions on Management Information Systems,* (3)2, pp. 1–17.

vom Brocke, J. (2007) "Service Portfolio Measurement: Evaluating Financial Performance of Service-oriented Business Processes", *International Journal of Web Services Research* (IJWSR), (4)2, pp. 1–32.

vom Brocke, J., J. Recker, and J. Mendling (2010) "Value-oriented Process Modeling: Integrating Financial Perspectives into Business Process Re-design", *Business Process Management Journal,* (16)2, pp. 333–356.

vom Brocke, J., M. Petry, T. Sinnl, B.Ø. Kristensen, and C. Sonnenberg (2010) "Global Processes and Data. The Culture Journey at Hilti Corporation" in vom Brocke, J., and M. Rosemann (eds.) *Handbook on Business Process Management, Volume 2,* New York: Springer-Verlag, pp. 537–556.

Voss, C.A. (1984) "Technology Push and Need Pull: A New Perspective", *R&D Management,* (14)3, pp. 147–151.

Wang, H., X. Du, F. Zhang, and W. Li (2011) "Application of RFID Technology in Elevator Maintenance Supervision", *International Conference on Computer Science and Service System* (CSSS), pp. 2112–2115.

White, T. (2009) *Hadoop: The Definitive Guide, first edition,* Sebastopol, CA: O'Reilly.

Wijnhoven, F., P. Dietz, and C. Amrit (2012) "Information Waste, the Environment and Human Action: Concepts and Research", *ICT Critical Infrastructures and Society,* Berlin, Heidelberg, Germany: Springer, pp. 134–142.

Wijnhoven, F., P. Schuur, and J. Timmer (2010) "The Inventor Game: Game-theoretical Analysis of Knowledge-sharing Between Inventors and Employers", *Knowledge Management Research & Practice,* (8)1, pp. 61–75.

Woods, D. (2012) "How Hadoop and SAP HANA Can Accelerate Big Data Startups", *Forbes,* http://www.forbes.com/sites/danwoods/2012/05/16/how-hadoop-and-sap-hana-can-accelerate-big-data-startups/ (current March 7, 2013).

Word, J. (2012) *SAP HANA Essentials, second edition,* Epistemy Press.

Wust, J., J.-H. Boese, F. Renkes, S. Blessing, J. Krueger, and H. Plattner (2012) "Efficient Logging for Enterprise Workloads on Column-oriented In-memory Databases", *Proceedings of the 21st ACM international conference on Information and Knowledge Management—CIKM '12,* New York: ACM Press.

Zmud, R.W. (1984) "An Examination of "Push-Pull" Theory Applied to Process Innovation in Knowledge Work", *Management Science,* (30)6, pp. 727–738.

zur Mühlen, M., and R. Shapiro (2010) "Business Process Analytics" in vom Brocke, J., and M. Rosemann (eds.), *Handbook on Business Process Management, Volume 2,* New York: Springer-Verlag.

Zyngier, S. (2002) "Knowledge Management Obstacles in Australia", *10th European Conference on Information Systems,* pp. 919–928.

ABOUT THE AUTHORS

Jan vom Brocke is a professor of Information Systems, the Hilti Chair of Business Process Management, and Director of the Institute of Information Systems at the University of Liechtenstein. Jan's research specializes in the areas of IT-enabled business innovation, Green IS, and Neuro IS. He has published more than 200 research papers, including in MIS Quarterly, the Journal of Management Information Systems, and Information and Management. He is author and co-editor of 19 books, including the International Handbook on Business Process Management. He serves on the editorial review board of the Journal of the Association for Information Systems, is an editor of Business & Information Systems Engineering, and is a co-editor-in-chief of the Journal of Information Technology Theory and Application.

Stefan Debortoli is a research assistant and Ph.D. candidate at the Hilti Chair of Business Process Management (BPM) at the University of Liechtenstein. His doctoral studies focus on the creation of business value through real-time and high-performance computing, including phenomena and innovations like Big Data and in-memory technology. Before joining the team, Stefan has gained over five years of working experience in the field of software engineering and IT project management across various industries.

Oliver Müller is Assistant Professor at the Institute of Information Systems at the University of Liechtenstein. Before joining the team, he worked as a researcher at the European Research Center for Information Systems (ERCIS) at the University of Muenster, Germany. In 2011, Oliver finished his doctoral thesis on customer decision support systems. Prior to his academic career, he gained industry experience as a consultant for supply chain management and as a visiting researcher at SAP Research. Oliver's research interests are business process management and decision support systems. His work has been published in peer-reviewed academic journals and presented at major international conferences. Moreover, he has been awarded with prizes for his research.

Nadine Reuter is research assistant and Ph.D. candidate at the Hilti Chair of Business Process Management at the University of Liechtenstein. Her doctoral studies focused on Green IS, in particular on the question on how information system can contribute to sustainability transformation. Before joining the team, Nadine has worked in the IT department of an international organization which was interested in in-memory technology.

Communications of the Association for Information Systems

ISSN: 1529-3181