RC25348 (WAT1301-071) January 27, 2013
Computer Science

IBM Research Report

Short Papers from the 6th Advanced Summer School on Service Oriented Computing (SummerSOC '12)

Rania Khalaf and Kostas Magoutis, Editors
IBM Research Division
Thomas J. Watson Research Center
P.O. Box 218
Yorktown Heights, NY 10598

Research Division
Almaden - Austin - Beijing - Cambridge - Haifa - India - T. J. Watson - Tokyo - Zurich

The 6th Advanced Summer School on Service Oriented Computing

July 2-7, 2012,

Hersonissos, Crete, Greece

The 6th Advanced Summer School on Service Oriented Computing (SummerSOC'12) continued a successful series of summer schools that started in 2007, regularly attracting world-class experts in Service Oriented Computing to present state-of-the-art research during a week-long program organized in four thematic tracks: Foundations of service-oriented computing; Computing in the Clouds; People in Services; and Emerging Topics. The advanced summer school is regularly attended by top researchers from academia and industry and by graduate students from programs with international acclaim, such as the Erasmus Mundus International Master in Service Engineering.

As part of education curriculum of the 6th Advanced Summer School, graduate students taking the school for credit were asked to form small teams with the purpose of undertaking a project of their choice inspired by a regular-track presentation. Each team researched and discussed their topic during the week of the 6th Advanced Summer School and presented their view of the state of the art, challenges, and promising future research directions in their particular area of focus in Services Science. The teams presented their findings in a special session during the last day of the summer school. The best teams were invited to further explore their work and compose short report papers based on their findings. These papers were then edited and integrated into the volume included in this Technical Report.

The two papers focus on key areas of Services Science, namely

- Flexible business processes, including an overview and analysis of their behavior. A key focus was ways to quickly build and improve process models and to enrich the functionality of flexible processes
- Service processes used by Data Center Infrastructure Management service providers, with a focus of how wireless-sensor technologies can improve these service processes in a variety of cases.

Rania Khalaf, Kostas Magoutis

Editors

Contents

Flexible Processes and Process Mining: A Brief Survey……………………………………………..……… 1
A. Murguzur, R. Makhani, B. Tao, and D. Zambon

Wireless Technologies in Datacenter Management…………..…………………………………………… 5
G. Frangiadaki, P. Kazmirchuk, N. Phaphoom, O. Smyrnaki, R. Zabolotnyi

Flexible Processes and Process Mining: A Brief Survey

Aitor Murguzur
IK4-IKERLAN
20500, Arrasate-Mondragon
Spain
amurguzur@ikerlan.es

Rahim Makhani
Erasmus Mundus -
International Master in Service
Engineering (IMSE)
rahim.makhani@gmail.com

Binjiang Tao
Erasmus Mundus -
International Master in Service
Engineering (IMSE)
taobinjiang@gmail.com

Davide Zambon
Rover Laboratories
25019 Sirmione (BS)
Italy
dav.zambon@gmail.com

ABSTRACT

Flexible and unstructured process handling has become very common in business primarily due to continuous changes in the business environment. Flexible processes consist of a series of activities that cannot be predefined. In this paper, we provide an overview of flexible processes and analysis used to understand their behavior. We also suggest four open areas of research regarding how to quickly build and improve process models and to enrich the functionality of flexible processes. The work in this paper is an extended version of a project presentation, based on topics from [3], presented at the 6th Advanced Summer School on Service Oriented Computing (`http://www.summersoc.eu/summersoc2012/`), held on July 2-7, 2012 in Hersonissos, Greece.

Keywords

Process management, Case handling, ACM, process mining

1. INTRODUCTION

Nowadays organizations have to deal with constant change and uncertainty in order to adapt and gain competitive advantages, i.e., they need to handle the increasing challenges of customer satisfaction, emerging markets, and new user requirements. Traditional Business Process Management (BPM) systems have focused on repetitive and predictable processes. Still, those processes need to operate in dynamic business environments where flexibility becomes a prerequisite. For these evolving situations, case-handling concept supports flexible and knowledge intensive processes, which combines document and process management [10].

At the same time, the amount of data is increasing at a rapid pace. Thus, the corresponding treatment and management methodologies for handling it are becoming critical in order to improve decision making and upgrade business processes. Process mining techniques can help to extract valuable information and patterns from the data collected [6]. Bearing flexible processes and process mining in mind, the main objective of the study presented here is to offer a comprehensive picture of these technologies and shed light on future research directions.

The remainder of the paper is organized as follows. First, case-handling is introduced in Section 2. Then, Section 3 details Adaptive Case Management (ACM). In Section 4, some of the process mining tools are illustrated and Section 5 deals with process mining at runtime. The last part of the paper, Section 6, draws conclusions and gives future recommendations.

2. CASE HANDLING

Knowledge-intensive business processes have proven to be very difficult (when not impossible) to be described by traditional BPM. The classical approaches to workflow are too restrictive and have problems dealing with change [10].

Examples of knowledge-intensive processes are: evaluating of a job application, the outcome of a tax assessment or the ruling of an insurance claim. The lack of flexibility of BPM appears when it is considered that in knowledge-intensive processes:

- not all the activities can be modeled as ACID,
- distribution and authorization may occur in different time/space,
- context tunneling may occlude some possibilities,
- *"what can be done"* is sometimes more important than *"what should be done"*.

In the paper [10], Aalst et al. suggest a new paradigm to solve such problems for handling knowledge-intensive business processes. The new approach increases the priority of the data (eliminating context tunneling), separates the authorization and the distribution of processes and allows workers to change the data before and after their activities. A short summary of the main differences between BPM and the Case Handling Paradigm is shown in Table 1.

Table 1: Short comparison between classical BPM and Case Handling [10]

	Workflow management	Case handling
Focus	Work-item	Whole case
Primary driver	Control flow	Case data
Separation of case data and process control	Yes	No
Separation of authorization and distribution	No	Yes
Types of roles associated with tasks	Execute	Execute, Skip, Redo

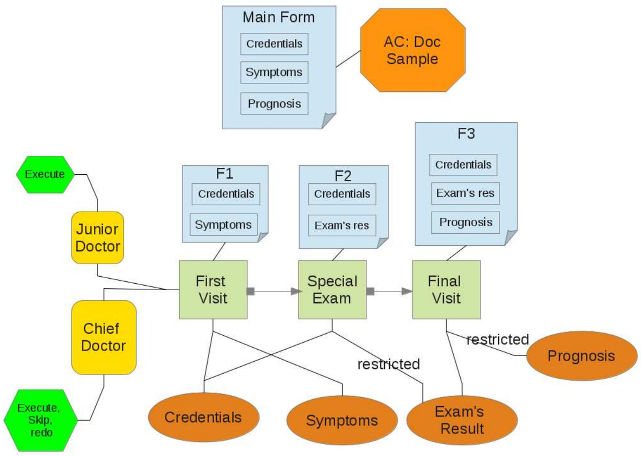

Figure 1: Abstract example of application of the case handling model.

The UML representation and a formal description of the model is also explained in [10]. In that paper, the case handling problem is described with a proper automaton which uses the following components:

- *case definition C*: complex case
- *form F*: data descriptor offered to the user,
- *activity A*: action [may be active or not, may produce new data,...],
- *activity role R*: class of worker trained to execute the role "R",
- *data object D*: any data representation.
- Relationships among activities and data objects may be:
  - *mandatory*: the data value MUST be entered BEFORE the activity execution
  - *restricted*: the data value MUST be entered DURING the activity execution

In order to give a better idea of the application of the described paradigm, a simple example is illustrated in Figure 1 (many details are hidden and the shapes/colors are added only to enhance the comprehensibility). The example represents a very simple version of one of the classical case management problems : the "Blind Surgeon Metaphor" in which a physician needs to diagnose a patient and if some complex exams are requested all the decisions must be taken accordingly with the patient state and history. In Figure 1, the following are represented:

- a case definition : "AC Doc Sample" contains all the activities (with its form),
- 3 activities: "First Visit", "Special Exam" and "Final Visit",
- 4 data objects: Credentials, Symptoms, Exam's Result and Prognosis,
- all the activity-data relationships are of type "mandatory" except for the ones marked with the label "restricted",
- every activity is associated with the corresponding form which contains the I/O interface of the procedure,
- the activities are ordered and to start the first activity, the "credentials" and the "symptoms" data objects must be provided,
- there are 2 roles defined: "Junior Doctor" and "Chief Doctor", these roles are valid for all the activities.

The example is a great simplification but it is an easy exercise trying to imagine how it can be extended (and still be formalized by the case handling model above), for example:

- adding new roles for the laboratory's technicians,
- adding different types of examination (blood, MRI, X-rays,...).....

Of course structural changes still raise problems even in this formalization, but the case handling approach extends the limits of BPM and can be considered an interesting balance between flexibility and support. Currently, taking the dynamic and flexible nature of cases, the initial concept of a "case" has transformed into a new way of approaching support of work.

3. ADAPTIVE CASE MANAGEMENT

The nature of work is changing: predictable processes are turning into knowledge-intensive processes where businesses require more agility to manage processes, tasks, people and documents. In such situations, process modeling in advance is an arduous and almost an impossible task, so knowledge workers take control of their own processes based on their previous experience. A knowledge worker is an employee who performs complex decision-making based on knowledge and judgment. ACM *"sits at the center of gravity for process, content, and customer relationship management"* [5] to help knowledge workers manage the unpredictable, ad-hoc processes they engage in everyday.

ACM, also known as Dynamic or Advanced Case Management, represents a radical new approach to supporting knowledge workers in today's leading edge organizations. ACM addresses the same concept of data-driven approaches but not in the same way. With ACM, *"a case is not a folder containing activities, processes, and documents; it is a con-*

solidated, connected, and evolving entity progressing towards completion" [8]. ACM adapts to the context of the case, guiding the outcome based on rules, data, etc. Rather than sequentially routing the case to the next task, department, or user role by action or condition, ACM advances through both external (e.g., receipt of a phone call or an email) and internal events (e.g., case task assignation). Hence, ACM addresses the issue of managing knowledge workers processes without overly constraining them, thereby increasing productivity. An ACM platform enables business actors to not only execute a specific task, but to actually add knowledge to the case.

Although some approaches have been proposed as process-aware tools to enhance knowledge worker's productivity (e.g., Cognocare[1]), ACM is still an emerging discipline. In addition, these solutions need to focus on process analysis and mining to extract information from stored data in order to create and manage best practices for knowledge workers, enabling the knowledge workers themselves to make their work more efficient and solving the unpredictable side of cases.

4. PROCESS MINING

Process mining [1] is a technique that helps in the analysis of business processes primarily using event logs. By tracing the event logs from an application, organizational workflows can be modeled based on the sequence of events. Various key insights can be gained from the event logs such as the relation of the event to a particular process instance, time on which the event was performed, and the performer of the event. In addition, organizational structure can also be modeled.

There have been various tools and techniques that are developed for process mining and ad-hoc workflows. A few of them are described below.

Business Process Insight (BPI) is a system that is used for the analysis of semi-structured processes. *"It includes a process intelligence toolset containing techniques related to process mining, business intelligence, business process management, complex event monitoring, and business activity monitoring"* [6]. Thus it is a framework that is used not only for managing semi-structured processes but also for discovering and analyzing process behavior.

Another system developed for ad-hoc workflow support is called Caramba. *"Caramba is one of a few process-aware collaboration systems that help people to collaborate across geographical distances and organizational boundaries"* [2]. It converts the event logs into a generic structure so that the other tools can extract this information in a manner that is suitable for process mining.

ProM is a widely popular framework that has been developed in order to implement process mining tools using a standard structure. The basic idea is that ProM framework defined a standard format to store process logs eliminating the problem of different formats of process logs from different information systems. *"One important feature of ProM framework is that it allows for interaction between a large number of so-called plug-ins"* [11], where a plug-in is primarily the implementation of an algorithm.

[1]http://www.cognocare.com/

5. RUNTIME PROCESS MINING

Based on the fast paced business environment, it is important to perform process mining at runtime and guide the users accordingly. One of the options is to train the system with historical information about previous workflow instances. By training the system, various patterns of actions can be developed based on the steps that are already performed by the users. This will help in providing suggestions to the user for selecting the next activity of the workflow during runtime. These hidden patterns can be identified by using predictive data mining techniques. The techniques could be especially helpful during case based analysis. For instance, by having an appropriate amount of historical data for a process model, classification algorithms based on decision trees can be applied in order to predict the next activity. When the same set of activities is followed again, the user can be suggested the next recommended activity based on the decision tree developed using the classification technique.

Runtime behavior analysis can be particularly helpful to eradicate bottlenecks and to improve the overall efficiency of a given case or a process instance. For instance, if a user is taking too much time to perform an activity, as compared to the time taken by the users based on historical information, the case could be escalated at run-time, or could even be transferred to another user.

There are a number of runtime techniques that have been developed. BPI has a predictive analytics component, where a decision tree can be trained in order to predict the next step of the process [6]. ProM has a recommendation service that recommends the next work item from the historical information [7]. In [9], an explicit process model is constructed to predict the completion time of running process instances by using an abstraction mechanism to represent the processes with time annotations. In [12], a collaborative task management system called KISSmir is developed that allow a user to adapt tasks, and also manage process related knowledge during the execution of tasks. The case health system in [4] provides ongoing analytics for cases in a case management platform to give an indication of how a case is doing based on: current metrics such as whether tasks are completing in time and predictions such as whether an undesired task is likely to execute.

6. OPEN AREAS

Having presented an overview of flexible processes and analysis used to understand their behavior, we conlcude by presenting some suggestions for open areas in this research domain. We highlight four such suggestions. We call the first the Q-Model, where "Q" stands for "Quickly Build". When a process-mining project is ongoing, it is always better to see the value of project in the early stages. Therefore, a quick and dirty way of building a prototype is necessary. As a consequence, an evaluation standard should be formalized to anticipate the value of the project.

The second is a process-mining technique on privacy-hidden data. It may be difficult to provide data to these systems,

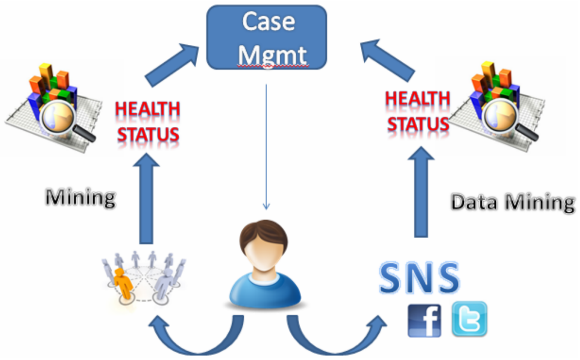

Figure 2: Health Care Involving Social Network.

especially in the case of sensitive or confidential information in the data itself. Instead of simply erasing some fields, privacy-preserving techniques can be used, like k-anonymity, on the set of log data. As a result, process-mining techniques need to adapt to the change. For example, handling cases where a data field is anonymized and where that data field may be crucial for determining the task flow. Thus, a probabilistic process mining technique can be further studied.

The third is a wikipedia-pattern process model. Everyone is able to contribute to improving the business processes. We suggest using what we call a "Hint Form" and associating such a form with each task. When a staff is conducting that task, and he/she finds some valuable data is missing, he/she will fill feedback in the Hint Form. The knowledgeable worker goes over the forms regularly. In this way, he can take the initiative to improve the business process, since the submitted form will raise notification.

Finally, there is a potential value in combining case management with social network. Taking the medical case for example, prevention is better than treatment. A potential corresponding architecture is shown in Figure 2. By collecting all sorts of information, like status, photos and blogs, on the patient's social network site, data mining techniques can be applied to detect the mood and emotion changes of the patient. Also, family members can get involved to provide more direct and accurate mental or physical information about the patient. A case is then designed to monitor the health status, and corresponding activities should be carried out by a knowledgeable worker to treat the patient in an early stage.

7. REFERENCES

[1] W. M. P. v. d. Aalst, A. Adriansyah, A. K. A. d. Medeiros, F. Arcieri, T. Baier, T. Blickle, R. P. J. C. Bose, P. v. d. Brand, R. Brandtjen, J. C. A. M. Buijs, A. Burattin, J. Carmona, M. Castellanos, J. Claes, J. Cook, N. Costantini, F. Curbera, E. Damiani, M. de Leoni, P. Delias, B. F. v. Dongen, M. Dumas, S. Dustdar, D. Fahland, D. R. Ferreira, W. Gaaloul, F. v. Geffen, S. Goel, C. W. Günther, A. Guzzo, P. Harmon, A. H. M. t. Hofstede, J. Hoogland, J. Espen Ingvaldsen, K. Kato, R. Kuhn, A. Kumar, M. La Rosa, F. Maggi, D. Malerba, R. S. Mans, A. Manuel, M. McCreesh, P. Mello, J. Mendling, M. Montali, H. Motahari Nezhad, M. zur Muehlen, J. Munoz-Gama, L. Pontieri, J. Ribeiro, A. Rozinat, H. Seguel Pérez, R. Seguel Pérez, M. Sepúlveda, J. Sinur, P. Soffer, M. S. Song, A. Sperduti, G. Stilo, C. Stoel, K. Swenson, M. Talamo, W. Tan, C. Turner, J. Vanthienen, G. Varvaressos, H. M. W. Verbeek, M. Verdonk, R. Vigo, J. Wang, B. Weber, M. Weidlich, A. J. M. M. Weijters, L. Wen, M. Westergaard, and M. T. Wynn. Process mining manifesto. In *BPM 2011 Workshops, Part I*, volume 99, pages 169–194. Springer-Verlag, 2012.

[2] S. Dustdar, T. Hoffmann, and W. van der Aalst. Mining of ad-hoc business processes with teamlog. *Data and Knowledge Engineering*, 2005.

[3] R. Khalaf. When people know better: Flexible ad-hoc flows and the analytics that help manage them. `http://www.summersoc.eu/summersoc2012/wp-content/uploads/2012/07/WhenPeopleKnowBetter-Compatibility-Mode.pdf`, July 2012. Advanced Services Summer School, Hersonissos, Greece.

[4] G. T. Lakshmanan, N. Mukhi, R. Khalaf, A. Martens, and S. Rozsnyai. Assessing the health of case-oriented semi-structured business processes. In *International Conference on Service Computing*, 2012.

[5] M. J. Pucher. Mastering the unpredictable with adaptive case management. `http://www.articlesbase.com/information-technology-articles/mastering-the-unpredictable-with-adaptive-case-manageme html`.

[6] S. Rozsnyai, G. T. Lakshmanan, V.Muthusamy, R. Khalaf, and M. J. Duftler. Business process insight: Platform for the discovery and analysis of end-to-end business processes. In *SRII*, 2012.

[7] H. Schonenberg, B. Weber, B. van Dongen, and W. van der Aalst. Supporting flexible processes through recommendations based on history. In *International Conference on Business Process Management*, 2008.

[8] K. D. Swenson. *Mastering the Unpredictable: How Adaptive Case Management Will Revolutionize the Way That Knowledge Workers Get Things Done.* Meghan-Kiffer Press, 2010.

[9] W. van der Aalst, M. Schonenberg, and M. Song. Time prediction based on process mining. *Information Systems*, 2011.

[10] W. M. van der Aalst, M. Weske, and D. Grünbauer. Case handling: A new paradigm for business process support. *Data and Knowledge Engineering*, 2005.

[11] B. van Dongen, A. de Medeiros, H. Verbeek, A. Weijters, and W. van der Aalst. The prom framework: A new era in process mining tool support. *Lecture Notes in Computer Science*, 2005.

[12] H. F. Witschel, B. Hu, U. V. Riss, B. Thönssen, R. Brun, A. Martin, and K. Hinkelmann. A collaborative approach to maturing process-related knowledge. In *International Conference on Business Process Management*, 2010.

Wireless technologies in datacenter management

Georgia Fragkiadaki
University of Crete
gogofrag@gmail.com

Petro Kazmirchuk
University of Crete
petro.kazmirchuk@gmail.com

Nattakarn Phaphoom
Free University of Bolzano
phaphoom@inf.unibz.it

Ourania Smyrnaki
University of Crete
smyrnak@csd.uoc.gr

Rostyslav Zabolotnyi
Vienna University of Technology
rstzab@infosys.tuwien.ac.at

ABSTRACT
Outsourcing of IT processes to specialized service providers is becoming the prevailing way for modern enterprises to handle an exponential growth of complexity in their data center infrastructures. A large number of IT service processes belong to the domain of datacenter infrastructure management. To gain a competitive advantage in this market, a service provider must continuously innovate, introducing solutions that are nonintrusive and simple to install and maintain. Devices based on wireless communication protocols, such as RFID and ZigBee, that are integrated with a variety of sensors connected in a wireless communication network often support solutions that satisfy these requirements. In this article we present a survey of recent advances in academia and industry aiming to improve datacenter infrastructure with the help of wireless sensor networks. Also we define the scope where such solutions can be applied. The work in this paper is an extended version of a project presentation, based on topics from [1], presented at the 6th Advanced Summer School on Service Oriented Computing[1], which was held on July 2-7, 2012, in Hersonissos, Greece.

Categories and Subject Descriptors
C.2.1 [**Computer-Communication Networks**]: Network Architecture and Design – *wireless communication*; C.4 [**Performance of Systems**]: *Measurement techniques, Reliability, availability, and serviceability.*

General Terms
Management, measurement, performance.

Keywords
Datacenter infrastructure management, IT process outsourcing, wireless sensor networks, green datacenter, health monitoring.

1. INTRODUCTION
Maintaining and improving data center (DC) infrastructure is an important task for IT departments in all modern enterprises. Recently, the size, complexity and energy consumption of datacenters has been growing significantly as the data processing and storage needs of enterprises have been growing exponentially. Outsourcing datacenter infrastructure management (hereinafter referred to as DCIM) processes to external service providers with specialized expertise is becoming the prevalent way to handle this complexity. A service provider is thus often invited to manage a private DC of an enterprise on-premise. Several large IT service providers (such as IBM, TCS, and Accenture) have added DCIM service processes to their portfolio of offerings; other more specialized service providers (such as Nlyte, Rackwise, and Asset VUE) focus solely on DCIM.

Although recently designed modular datacenters such as Microsoft's IT Pre-Assembled Components (ITPAC) come with all the management infrastructure built-in, a service provider has to work with a customer's existing IT infrastructure. Tools and processes used by IT service providers are expected to be as non-intrusive as possible and easy to use for system administrators of the service consumer. Therefore, the disruption to existing processes is minimized. In this article we provide a survey of how wireless technologies are currently used in ways that can be of direct use and benefit to DCIM service processes:

1. Monitoring air temperature, pressure and humidity, which is useful for cooling optimization to reduce power consumption and carbon emission.
2. Location discovery, which simplifies inventory tracking tasks and speeds up failure management, when physical intervention is required.
3. Reserving an alternative communication channel, that can be temporarily used when wired channels break.

We believe that these are among the most important service processes in the portfolio of any IT service provider engaged in DCIM.

On the contrary, if a service provider is providing IT services for an enterprise off the latter's premises (i.e. in the server provider's datacenter with modern hardware and well-established workflows), then in most cases all the needed infrastructure is expected to be already available. Therefore this use case is out of scope of our paper.

The remainder of this paper is organized as follows. Section 2 discusses the current issues of cooling optimization in the data center and proposed solutions using wireless sensor networks (WSNs). Section 3 summarizes how RFID tags and WSNs can be used to solve the location discovery problem. Section 4 outlines the advantages of wireless sensors for equipment health monitoring. Finally, in Section 5 we provide our conclusions.

2. COOLING OPTIMIZATION
Data center temperature and environmental monitoring has become critical due to the rapid growth in the IT industry. An inefficient energy management not only increases operational costs but it also has an adverse impact on the environment. Empirical findings have shown that the heat density in data processing and computer systems has increased significantly in this decade [4]. This is partly a result of the better physical

[1] http://www.summersoc.eu/summersoc2012/

configuration of data centers that allows servers to be vertically racked to conserve floor space. The traditional cooling method, i.e. using regular air conditioning that provides cool air from the floor and pumps the heat out through ceiling, is no longer sufficient in this setting where hundreds of racks are densely packed in limited space. Such an inefficient cooling approach leads to a data center that is over-cooling in some areas and over-heating in other areas.

One approach to overcome the cooling challenge is to implement an automated and integrated temperature control system that adjusts the data center's cooling behavior according to environmental variables detected by sensors installed in different locations in the data center. Wireless sensors networks (WSN) can be used to serve this purpose. WSN poses several advantages over wired methods, such as better manageability and non-disruption to existing infrastructure. General requirements of a wireless environmental monitoring system include reliability which is measured by a percentage of data loss, high performance, usability which measures how easy it is to install and maintain the system, and low power consumption.

In 2011 Rodriguez et al. [14] proposed a prototype system that integrates WSN for distributed data collection, data transmission and analysis for environmental monitoring. The prototype collects temperature and humidity data and creates a real-time thermal map of the rack environment to facilitate dynamic temperature control. The implemented prototype could be conceptualized into three modules including network architecture design, network hardware, and software. Firstly, the prototype employs a multihop cluster network to capture environmental data in the dense rack area. Secondly, the hardware is composed of sensor units that create packets of sensed data, a WSN based on XBee 2.5 RF modules and power supplies. Thirdly, the software module is composed of sensor interface software that converts streams of sensor data to a specific data format and embeds a data package with a sensor id. An XBee node can be wirelessly configured to behave as a coordinator, router or sensor node. The data management system includes a part that reads the sensed data and saves to the database, a part that manages different types of sensed data, and lastly, a part that controls graphic user interfaces and visualizes data over time.

The proposed system was deployed to the Argonne data center to monitor its environment in ten locations. The results demonstrated that the system is able to capture fluctuations in temperature and humidity in different locations. It is easy to be deployed throughout the opaque rack area and flexible to relocate, reflecting its usability without the need of wiring for power and data transmission. In terms of reliability, within a test period of 24 hours no data packet was lost during the transmission.

Seeking for a better alternative to the existing cooling designs such as “cold-aisle-hot-aisle” and Computational Fluid Dynamics Simulations, in 2008 Microsoft developed and deployed a WSN monitoring system that could harvest data from different sources – the Data Center Genome System (DCGS) [11]. The used IEEE 802.15.4 standard has its pros and cons, as it achieves certain cost reduction, but at the same time it has higher bit error rates than other wireless technologies. The nodes in this system are called Genomotes, and are divided into master and slaves. Slaves contain humidity and temperature sensors, whose data is periodically sent over a wired interface to master nodes. There the data is stored in 1MB flash memory and later transmitted over the wireless channel. The collected data provides both heat distribution visibility and adaptive air conditioning in real time.

Two years after Microsoft's Genome Project, the Server Environment Monitoring and Controlling System (SENVM) [6] for a small data center using a wireless sensor network was proposed and implemented by the Faculty of Engineering of the Department of Computer Engineering in Chulalongkorn University on Bangkok, Thailand. The authors of this approach characterize it as the first proposed monitoring and controlling system using WSN, since, in contrast to the Microsoft project, it goes one step further to cover the controlling part as well.

3. LOCATION DISCOVERY

Another domain where wireless technologies may come to the rescue is the location discovery problem; in other words, how to map a given server with its IP and MAC addresses to a physical location in a room. Supporting a manual registry will become a tedious, error-prone and time-consuming task in a highly dynamic and fast-growing DC, where servers and racks are often added, moved and replaced. This problem becomes even more complicated, if DCIM is outsourced to a service company. This location information is required for both inventory tracking and resolving critical problems, when physical intervention is required.

Various solutions are offered in different price ranges. E.g. HP Location discovery services included with HP ProLiant Gen8 servers rely on the intelligence that is built in racks to report the server's position inside of a rack and a rack's ID [3]. Still, this is not an exact position in a room. Moreover, in this paper we are focusing on less intrusive solutions than purchasing new expensive servers.

In the low-cost segment, there is the LightsOn solution from Vizualiiz [15]. They offer software for mobile platforms that, together with RFID tags, simplifies the inventory process. However, each change must be accounted manually.

In the WSN domain, the problem of determining the node's location based on triangulation principles (measuring the received signal strength or the time to arrival to special nodes with known location) is well researched, but it is usable only in the case of open spaces. In a room filled with metal, the accuracy of these methods is expected to be significantly reduced. RFID zoning is an alternative solution that may be applicable. For example, the VIZBEE solution [16] is comparatively cheap and easy to install and use. A DC administrator will be able to determine not the exact location though, but only a zone (within e.g. 2-3m of radius) where a server is located.

To our knowledge the most promising solution was offered by the researchers from Southern Illinois University and Intel [17]. Instead of measuring distances to a few predefined nodes with known locations in a room, the authors proposed to exploit the knowledge about server racks' geometry, which is mostly standardized, such as width of a rack. Therefore, they install a wireless node on each server and measure distances between servers at first in the same rack, and then - distances between racks. This incremental process allows discovering and keeping track of physical locations completely automatically after all the wireless nodes have been placed. Moreover, these nodes provide an alternative communication channel in case a top-of-rack switch fails.

4. SERVER HEALTH MONITORING

One of the most important requirements of a data center is reliability. To be able to offer guaranteed SLAs, all equipment that takes part in the data center functionality has to be continuously monitored and validated. Existing solutions are either committed to a specific hardware vendor or even a specific model or are expensive and have complex installation procedures that may disable the data center for a long period of time. Most of the existing systems have their own server monitoring protocols and standards and cannot be merged together or sometimes even conflict in case of parallel installation. Here we try to investigate what are the options that might solve the described challenges and if there can be a simple and yet powerful solution that will provide functionality comparable with industrial systems, but cheaper and simpler in installation in an existing DC.

The problem of equipment status monitoring is well-known and covered in the literature [10] [13] and in real-world implementations [5] [7], but some issues are still not addressed. Nowadays there are different ways to monitor QoS provided by a datacenter, so we will consider the pros and cons of some of the most popular ones. Most of the data center hardware vendors understand the importance of equipment health monitoring and provide their own solutions to solve this issue [10] [9]. Usually these systems are robust and easy to install (if they are not already preinstalled with the equipment), but still there are some issues. The biggest challenge with vendor-provided health monitoring systems is that they are usually locked to the provided equipment and do not integrate well with equipment and status monitoring software provided by other vendors. Alternatively, third-party monitoring solutions can be selected [12]. To achieve big market share, such companies usually try to support equipment provided by multiple vendors. To achieve this, monitoring software cannot use vendor-specific means of retrieval of diagnostic information, therefore the quality and set of supported features provided by such systems usually tend to be lower than vendor-specific software can provide. Additionally, third-party systems can require administrators to install additional hardware that will perform server control and monitoring, that requires complex installation sequences, obfuscates existing network and can influence overall network reliability and stability. Another option to provide data center health monitoring is to use smart load balancers or border routers [2], that provide server status detection functionality. Of course, such solutions make sense only in data centers that require load balancing or border control. To achieve at least some level of server status monitoring, often system administrators tend to develop their own simple solutions that work exactly in their environment. Such an option is quite popular in small and private data centers, but it is hardly scalable and maintainable, so it is acceptable only if equipment in the data center is very diverse and a company cannot afford a professional health monitoring solution.

One of the possible solutions to this problem is to use small health monitoring sensors based on single-board microcontrollers that gain popularity nowadays. These sensors are cheap and easily programmed [8]; they can support different means of communication and can easily be tailored to work with different software and hardware. To control the state of a single server, they can be attached to a console, USB or management network port of the server. Additionally, to exclude any influence and load on the existing network communication, they could communicate with the help of the wireless network. This solution has some benefits over classic solutions. At first, it adds no load on the existing network. Additionally, it does not require new network infrastructure or additional cables. Instead, it creates a new scalable network that can be used for maintenance or fault management. As the controllers are programmable, they can support special features of the hardware they have to work with, but still remain vendor-independent, as they can be easily programmed to support only generic functionality. Due to these reasons, we think that this solution can be easily introduced in data centers of any size without disruption of their operations.

5. CONCLUSIONS

In this article we have described how wireless technologies and especially WSNs can significantly simplify inventory and fault management in a DC, provide an alternative communication channel and reduce the total cost of ownership, particularly due to the reduced power consumption and increased PUE. We believe that all these tasks can be combined in a single sensor node, installed on the per-server (preferably) or per-rack basis. As of mid-2012 we are unaware of a single IT service company that offers such an integrated solution. The main issue that arises in this case (compared to wired solutions) and must not be overlooked is the security of the chosen wireless communication channel. Therefore, proper cryptographic protection and authentication patterns should be applied.

6. REFERENCES

[1] K. Magoutis, “Recent advances in large-scale IT service management technologies”, Advanced Summer School in Service-Oriented Computing (SummerSOC 2012), http://www.summersoc.eu/summersoc2012/wp-content/uploads/2012/07/SummerSOC-magoutis-talk-4jul2012-1.pdf, July 2012, Hersonissos, Greece.

[2] Barracuda Networks. Barracuda Load Balancer, 2012. https://www.barracudanetworks.com/ns/products/balancer_overview.php.

[3] Mike Bedford. Universal EPROM Programmer. *Electronics Today International,* 45–51, 37–39, 1983

[4] K. Brill. Heat Density Trends in Data Processing, Computer Systems, and Telecommunications Equipment: Perspectives, Implications, and the Current Reality in Many Data Centers, 2006. http://www.fibertown.com/pdf/uptimeinstituteheatdensitytrends.pdf

[5] B.J. Casper, J.H. Dunn, S.B. Hirsch, B.D. Markee and D.C. Simone. Method and system for monitoring and dynamically reporting a status of a remote server. US Patent 6505248, 7 January 2003.

[6] S. Choochaisri, V. Niennattrakul, S. Jenjaturong, C. Intanagonwiwat, and C. Ratanamahatana. SENVM: Server Environment Monitoring and Controlling System for a Small Data Center Using Wireless Sensor Network. In *Proceedings of International Computer Science and Engineering Conference (ICSEC2010),* Chiang Mai, Thailand, November 2010 http://arxiv.org/ftp/arxiv/papers/1105/1105.6160.pdf

[7] IBM Corporation Software Group. IBM Tivoli Monitoring for Energy Management. *IBM Corporation,* NY, 2009.

[8] HP Inc. Innovative technologies in HP ProLiant Gen8 servers, 2012. h20000.www2.hp.com/bc/docs/support/SupportManual/c03227849/c03227849.pdf.

[9] HP Inc. The HP Network Management Center, 2012. http://h17007.www1.hp.com/us/en/products/network-management/index.aspx#comprehensive.

[10] L. Liu, H. Wang, X. Liu, X. Jin, W.B. He, Q.B. Wang and Y. Chen. GreenCloud: a new architecture for green data center. In *Proceedings of the 2009 ACM/IEEE Conference on International Conference on Autonomic Computing & Co-Located Workshops (ICAC'09)*, Barcelona, Spain, June 2009.

[11] Microsoft Corp. Data Center Genome Project, 2011 http://research.microsoft.com/en-us/projects/dcgenome/

[12] Molex. Brad Diagnostics and Tools, 2012. http://www.molex.com/molex/products/family?key=enetmeter_dn&channel=products&chanName=family&pageTitle=Introduction.

[13] J. Moore, J. Chase, K. Farkas, and P. Ranganathan. Data center workload monitoring, analysis, and emulation. In *Proceedings of Eighth Workshop on Computer Architecture Evaluation using Commercial Workloads (CAECW-8)*, San Francisco, CA, February 2005.

[14] M.G Rodriguez, L.E. Ortiz Uriarte, Y. Jia, K. Yoshii, R. Ross and P.H. Beckman. Wireless sensor network for data-center environmental monitoring. In *Proceedings of the Fifth International Conference on Sensing Technology (ICST 2011)*, Palmerston North, New Zealand, November 2011, pp. 533-537.

[15] Vizualiiz. LightsOn Solution, 2012. http://www.vizualiiz.com/lightson

[16] Vizbee Inc. Performances and cost effectiveness of Zoning Real-time location in comparison with Triangulation methods for indoor applications, 2012 www.vizbee-rfid.com/performances-and-cost-effectiveness-zoning-real-time-location-comparison-triangulation-methods

[17] N. Udar, K. Kant, and R. Viswanathan. Asset localization in data centers using WUSB radios. In *Proceedings of the Seventh International Networking Conference on Ad Hoc and Sensor Networks, Wireless Networks, Next Generation Internet (NETWORKING 2008)*, Singapore, May 2008, pp. 756-767.