

Energy Cloud: real-time cloud-native Energy Management System to monitor and analyze energy consumption in multiple industrial sites

Hugo Sequeira and Paulo Carreira
Dept. of Computer Science and Engineering
Instituto Superior Técnico
Lisbon, Portugal

Email: {hugo.sequeira, paulo.carreira}@tecnico.ulisboa.pt

Thomas Goldschmidt and Philipp Vorst
ABB Corporate Research
Ladenburg, Germany

Email: {thomas.goldschmidt, philipp.vorst}@de.abb.com

Abstract—Industrial organizations use Energy Management Systems (EMS) to monitor, control, and optimize their energy consumption. Industrial EMS are complex and expensive systems due to the unique requirements of performance, reliability, and interoperability. Moreover, industry is facing challenges with current EMS implementations such as cross-site monitoring of energy consumption and CO_2 emissions, integration between energy and production data, and meaningful energy efficiency benchmarking. Additionally, big data has emerged because of recent advances in field instrumentation that led to the generation of large quantities of machine data, with much more detail and higher sampling rates. This created a challenge for real-time analytics. In order to address all these needs and challenges, we propose a cloud-native industrial EMS solution with cloud computing capabilities. Through this innovative approach we expect to generate useful knowledge in a shorter time period, enabling organizations to react quicker to changes of events and detect hidden patterns that compromise efficiency.

Index Terms—Energy efficiency, energy management systems, cloud computing, big data

I. INTRODUCTION

The industrial sector produces the most CO_2 and is one of the largest consumers of electricity worldwide, at a rate that continues to grow annually [1]. However, due to limited resources and high costs, energy production is not growing at the same ratio, resulting in a demand-supply mismatch [2]. In an effort to close this ever-widening gap, energy suppliers and consumers are working together to keep demand under acceptable and secure levels. Energy suppliers run a set of Demand Response (DR) programs that influence consumers to amend their energy consumption, through changes in the price of electricity or by financial incentives [3]. On the other hand, energy consumers can use their available energy more efficiently by their own initiative. In order to accomplish this, industries must find inefficiencies and reduce energy consumption without affecting their business and production processes. In residential and commercial sectors, this essentially involves using energy efficient equipment or dimming out lights and heaters. In contrast, energy efficiency initiatives encounter unique difficulties in the industrial sector, including production and quality constraints, multiple energy tariffs, and

consumption and emission restrictions that make the task of saving energy more complex. EMS are tools that monitor, control, and optimize energy consumption [4]. Nevertheless, literature, research projects, studies and industry expertise, make clear that there is a need for a novel and robust platform capable of providing more and better energy information monitoring, integration, repository, and analytics towards a future energy efficient manufacturing (see Table I). In addition, recent advances in hardware, networking and software of sensor and control equipment, resulted in a massive increase of machine data in terms of volume, variety and velocity (the three V's of big data), that hinders the capability to collect, monitor, and analyze all these data [5].

In the meantime, cloud computing is coming to the forefront and being applied to various fields. Its main application is solving large scale computation problems by optimizing and combining distributed resources [6]. Cloud computing is a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction [7]. It provides several benefits such as saving of IT costs and maintenance, strong integration capabilities, short time-to-benefit, and scalable computation on demand that keep up with customer needs [8]. These benefits have pushed many residential and commercial EMS solutions to the cloud [9]–[11]. Hence, based on our research across publications, research projects, and industry trends (e.g., IMC-AESOP, IMS2020, Industry 4.0, Internet of Things, Cyber-Physical Systems), we believe that the industrial sector will also incorporate cloud technologies in their energy management and production processes in the near future. Therefore, the central motivation for this work is to study the migration of industrial EMS to the cloud, evaluate the possibilities to achieve more energy and cost savings, taking advantage of the latest cloud computing and big data technologies, and finally propose a solution (see Section III). The conclusions of this work will then be validated by developing, deploying and evaluating the

feasibility and performance of a cloud-native industrial EMS proof of concept with big data capabilities.

A. Research methodology

A research methodology road map was outlined that involved a comprehensive research on current literature, surveys, and discussions with industry experts from ABB, in order to identify the needs of industrial companies. ABB is an engineering company, leader in power and automation technologies that improve the performance of utility and industrial customers, while lowering environmental impacts. The research finished with topic consolidations and refinements in order to shape the solution proposal. Finally, to prove the feasibility of this work's claims, a proof of concept was developed and evaluated.

B. Document organization

The following Section II addresses the research background regarding this work. The solution proposal can be found in Section III and its evaluation in the following Section IV. Finally, Section V presents the final conclusions and summarizes the most important aspects.

II. RESEARCH BACKGROUND

Energy management stands for all the measures and activities which are planned or executed in order to minimize the energy consumption of an organization [4]. It affects all organizational and technical processes, in order to reduce, within economical constraints, the consumption of energy and increase energy efficiency. In fact, energy management is beneficial for economic, environmental and social reasons [12].

EMS are energy management tools used in a wide variety of applications, to effectively monitor, optimize and control power generation, distribution and consumption. The main goal of these systems is to increase energy efficiency and thus achieve energy savings, through continuous monitoring and maintenance of sites, improving the operation of equipment and decreasing energy consumption without compromising the customer needs [15]. Since lowering the energy costs increases the profit, the search for energy saving potentials always affects the bottom line. On the other side, one of the barriers to the adoption of EMS is the capital investment necessary to deploy and maintain these systems.

A. Survey on industry trends and current EMS

At present there are many researchers actively researching the future of smart and green manufacturing architectures.

IMC-AESOP is an EU-funded research project that researches a novel cross-layer Service Oriented Architecture (SOA)-based architecture, entitled "Cloud of Services", as the future of industrial infrastructures.

Cyber-physical system (CPS) refers to a new generation of intelligent and real-time control and Internet of Things (IoT) systems with integrated computational, communication, sensor and actuation capabilities.

Industry 4.0 is a term coined by the German government under the premise of "Smart Factories" [17]. It stands for the




Fig. 1. Real-time is a requirement in industrial decision-making because any disturbance affects several domains at the same time such as energy, production, and quality (adapted from [16]).

coming industrial revolution guided by the introduction of the IoT, CPS, big data, and cloud computing into the manufacturing environment. The basic principle is that by connecting machines, work pieces and systems, via the Internet we are creating intelligent networks along the entire value chain that can control each other autonomously.

To understand the current state of EMS solutions we conducted a survey on industrial EMS solutions standards (ABB cpmPlus EM, GE XA/21) and on EMS with cloud capabilities (DemandSMART, ClockWorks, EEMSuite, Powertech). Our findings show that the current cloud-base EMS are more focused on the residential and building sectors to mainly manage utilities or heating systems. In addition, current on premise industrial EMS have stronger monitoring and controlling capabilities, but are weaker on the integration and analytics side (see Table II).

III. SOLUTION PROPOSAL

The functionalities that a cloud-native industrial EMS may offer as a service are endless. However in regard to the current industrial energy management needs, the key features that we propose that any novel industrial EMS system must have are real-time cross-site energy consumption monitoring, integration and benchmarking between energy and machine data, and meaningful energy data analysis and cost allocation.

A. Risks and concerns

The adoption of a cloud approach on a traditional and sensitive domain as the industrial domain, raises some valid concerns:

1) *Security and privacy*: Handing over sensitive data and sharing infrastructures worries managers. The solution relies on the right cloud deployment architecture such as private or hybrid clouds—with sensitive data on premise—and proper secure and encryption techniques.

2) *Performance and reliability*: Real-time monitoring systems are hard to implement in the cloud, due to latency issues. However, energy monitoring is a non-sensitive operation, usually performed per min. or every 15 min. The right cloud setup can provide at least equal or faster times.

TABLE I
SUMMARY OF CURRENT INDUSTRY CHALLENGES AND HOW CLOUD TECHNOLOGIES COULD ADDRESS THEM [2], [12]–[15]

Industry challenges		Cloud computing resolutions	
Challenges	Implications	Cloud benefits	Added value
Industries have complex infrastructures and operations	Industrial EMS are expensive and harder to implement	Cloud solutions reduce investment and maintenance costs and with less time-to-benefit	Organizations can now afford to have an EMS with less investment and maintenance costs
Sites are often geographically distributed and self-managed	Hinders the ability to have a multi-site energy management	Cloud solutions are centralized with easy access to data	Achieve energy monitoring and analytics across sites
Weak integration between energy and production data	Manufacturing is not energy-aware and decision making is not informed nor based on real-time data	Cloud can provide interoperability and integrate data from external systems	React faster to changes with data driven and informed decisions
Weak and decentralized energy efficiency and cost analysis tools	Harder to benchmark energy usage and strategies across sites	Centralized correlation of data from all production levels	Derive knowledge from energy use and identify inefficiencies
Huge amounts of machine and energy data to analyze and process	Guesswork to find consumption inefficiencies and optimal schedules	Cloud can optimize the performance of scheduling algorithms	Find hidden patterns and produce new, faster and richer knowledge

TABLE II
SUMMARY OF GENERAL FEATURES PROVIDED BY CURRENT EMS:
✓ SUPPORTED, ✗ : NOT SUPPORTED, - : UNKNOWN INFORMATION

	EEMSuite	ClockWorks	PowertechEMS	DemandSMART	GE XA/Z1	ABB cpmPlusEM
Target Market						
Building Sector (HVAC)	✓	✓	✓	✓	-	-
Industrial Sector	✓	-	✓	-	✓	✓
Cloud Technologies						
Cloud storage	-	✓	✓	✓	✗	✗
Cloud capabilities (Scalability, etc.)	-	✓	-	-	✗	✓
Cloud deployment (SaaS, Paas, IaaS)	✓	✓	✓	✓	✗	✗
Data Integration						
Energy meter data gathering	✓	✓	✓	✓	✓	✓
Equipment status gathering	✗	✓	✓	✓	✓	✓
Environmental data gathering	✗	✓	✓	-	✓	✓
Production data integration	✗	✗	✗	✗	-	-
Data Mining and Analysis						
Unusual pattern detection	✓	✓	✗	-	-	-
Efficiency advisory	-	✓	✗	✓	-	✓
Cost allocation	-	✗	✓	✓	✓	-
Demand forecast	✓	✗	✗	✓	✓	✓
Energy-use Evaluation						
Equipment efficiency benchmarking	-	✓	-	✓	-	-
Process efficiency benchmarking	✗	✗	-	✗	-	-
Multi-site benchmarking	✓	✓	-	-	-	✓
KPI for energy efficiency	✓	-	-	✓	-	✓
Monitor and Control						
Real-time monitor	✓	✓	✓	✓	✓	✓
Remote control	-	-	✗	✓	✓	-
Load Management capabilities	-	✗	✓	-	✓	✓
Demand Response capabilities	-	✗	✗	✓	-	-
Alarms	✓	✗	✓	✓	✓	✓

B. Bigdata challenge

Recent advances in control and sensor devices have resulted in the generation of large quantities of data (the volume of bigdata). These equipment now produce data with much more detail and in a much higher frequency (the velocity of bigdata). Therefore, time series data automatically originated from

thousands of different sensors and customers (the variety of bigdata), presents a technical challenge in terms of collecting, storing, and processing these data in a real-time basis, and at the same time produce meaningful knowledge. Our solution proposal applies big data techniques, like batch and real-time computation, to provide real-time analytics on large quantities of incoming energy and machine data streams.

C. Conceptualization

Based on Marz and Warren [18] lambda architecture and Goldschmidt et al. [19] time-series cloud architecture for industrial processes, our architecture relies on two parallel layers that process continuous data streams as follows (see Figure 2):

1) *Energy Monitoring (Real-time computing)*: This high-speed pipeline layer processes all incoming time-series data, applies simple data transformations, and outputs the processed data without storing it. The main concern of this layer is to transfer data from one end of the system (sites) to the other (desktop and mobile clients) as fast as possible, so that users can react faster to quick changes of events.

2) *Energy Analytics (Batch processing)*: This slow-speed layer also processes all incoming data. However, it autonomously stores and analyzes these data using a series of pre-defined programs (“jobs”). Its goal is to compute complex algorithms that produce knowledge from all available data sources, also known as Knowledge Discovery in Databases (KDD).

D. Implementation

This work intends to provide a sustainable and conceptual architectural solution common to any cloud-native industrial EMS. Nevertheless, to prove the feasibility and claims of this work’s contributions, the authors developed a proof of concept, entitled *Energy Cloud*. Running entirely in the cloud, this implementation is a scalable and real-time industrial energy management system capable of monitoring (Energy Monitoring) and analyzing (Energy Analytics) time series energy metering data from external or simulated energy meters.

In this prototype, the authors only focus on a specific set of use cases that highlight the use of the cloud in this domain (see




Fig. 2. Conceptual architecture for a cloud-native industrial EMS.

section IV). The following sections depict the technologies and implementations for each architectural component and how they work together as a unified solution.

1) **Energy Cloud: Dashboard** is the front-end system that enables interaction between the user and the system. The monitoring page, based on the innovative Freeboard.io library, uses a set of widgets to visualize real-time data in lines (to display the variation of energy consumption over time), pies (to depict the ratio of consumption per site), and columns (to have a side-by-side comparison of energy usage per site).

By default, the screen comes with a fixed widget that shows an overview of the energy consumption of all sites. The user can then dynamically customize the rest of the screen with other widgets. All these widgets subscribe to internal and external data sources (using e.g. WebSockets or JSON) that update their visualization automatically (see Figure 3).

Analytics API is a back-end system that handles requests from the Dashboard or any other external system. Using Spring, this web API offers a set of operations to pull historical energy consumption measurements or previously calculated KPI from the Timeseries DB. It also offers options to execute on-demand Energy Analytics computation jobs. Background threaded daemons hold on to these requests and reply to them as soon as the operations are completed.

Message-oriented Middleware is the cross-layer component that enables communication of messages between several distributed systems. In this case, there are many open source options to choose from. But since the AMQP industry standard protocol was needed to support a wider variety of developer platforms, the authors chose RabbitMQ because it is a robust,

yet easy to use and deploy queue messaging system. A collection of exchanges and queues were created to queue all input data that comes from the Data Collector to be consumed by both computing layers, and all output data from the Energy Monitoring layer, to be consumed by the Real-Time Messaging Servers.

2) **Energy Monitoring: Storm Cluster** is essentially a distributed set of Apache Storm virtual machines that continuously compute unbounded streams of data in a very fast manner. This open-source, fully scalable, fault tolerant, and distributed system, recently clocked one million 100 byte messages per second per node (2xIntel E5645 2.4Ghz processors and 24Gb memory). A custom topology (a graph of parallel computation nodes) was developed to consume, filter, normalize, sample and compute incoming energy metering data streams in short time windows (e.g. 5 or 10 seconds) and find the most significant energy consumers per site (e.g. top 5 consumers).

Real-Time Messaging Servers are network servers that consume data from the MoM and distribute them to any subscribed client through WebSockets. The authors used Node.js to manage all the WebSockets connections and to push in real-time any incoming data from the MoM queues to every client listening for it.

3) **Energy Analytics: Hadoop Cluster** is a set of batch processing clusters that perform complex computations. There are endless opportunities here and much research has been made in the last years in the fields of data mining, machine learning, pattern recognition and others. In this proof of concept, a simple use case was chosen to demonstrate the inherent possibilities. Based on the work of Anna Koufakou et al. [20], the authors developed a simple outlier detection system (MR-AVF) to identify energy consumption peaks per site. The user is then aware of significant energy usage that might be inefficient or unexpected, and later lead to higher energy costs, or worst, result in penalties from the energy providers from crossing a contracted limit. Therefore, it is of the utmost importance to keep these peaks under control.

The MR-AVF is based on the Map Reduce paradigm for parallel programming. It provides high-speed and scalable outlier detection by analyzing each individual point and categorizing it by the average rate of occurrence. The more infrequent or irregular a value is, the more likely it is to be an outlier. This rather simple approach is easy to parallel and implement. Even though, it is faster and sometimes more efficient than other more complex calculations.

Timeseries DB and Distributed Historical Storage are the persistent data store components that archive data from the computation layer. Following the lambda architecture principals, every detailed raw data is time-base stored and immutable. The authors used KairosDB, a fast distributed scalable time series database, that works on top of Cassandra, an industry standard for distributed NoSQL databases. KairosDB provides us with an easy to use abstraction layer (API) to push and pull time series data from the Cassandra cluster. This setup was proved to be very robust, scalable, and suitable for

industrial processes by Goldschmidt et al. [19]. A cluster of 24 KairosDB nodes could handle the workload of a large city (6 million smart meters).

E. Deployment

This work's proof of concept was deployed on ABB's development cluster that runs Openstack, an open source software for building private and public clouds. All the components described before, except the Storm cluster, were deployed on top of Cloud Foundry, a Platform as a Service (PaaS) system that made the deployment of the developed applications in Spring and Node.js, and services, such as the KairosDB, RabbitMQ and Cassandra, much easier. Furthermore, it provided us with easy to use commands to vertically and horizontally scale all the components independently.

IV. EVALUATION

To prove the application and benefits of the cloud in current EMS, the authors focused on five use cases, extrapolated from recent literature and industry experts' expertise, that represent innovative functionalities that the industry lacks or that could not be easily obtained in a multi-site management without the cloud. Each use case was then evaluated regarding the following questions:

Q1: What is the value added by the use case? Q2: How easy is it to obtain the same results without the use of the cloud? Q3: Which aspects of the cloud does this use case highlight? Q4: How is this use case implemented?

A. UC1: Monitor the most significant energy consumers

The most significant energy consumers need to be autonomously identified, monitored and analyzed in real-time because they are the ones that contribute the most for the energy costs. This use case also supports the judgment as to whether anticipated energy savings like DR, the scheduling of tasks to avoid peaks loads, and as to whether take advantage of on-site power generation during production [12]. This requires standardization of data collection, on-line data processing and visualization techniques [21].

Q1: Increase industrial energy efficiency by monitoring and identifying unaware high consumers of energy. Q2: Moderate, because with current on premise systems only results from individual sites can be obtained. Performance is also limited to system's capabilities. Q3: Real-time analytics. Q4: The Energy Monitoring layer computes the top consumers per site in a real-time basis (see Figure 3).

B. UC2: Calculation of energy performance indicators (KPI)

The system must calculate energy related KPI to enable actors within an organization to react to negative developments. Nowadays, there is a need for effective energy efficiency KPI to track the changes and improvements on both process and plant level [12].

Q1: Increase industrial energy efficiency by identifying energy improvement opportunities, increase energy consumption awareness and assist the development of strategies. Q2: Hard



Fig. 3. Monitoring dashboard has a global overview of the energy consumption and the top consumers per site, and other individual measurements.

for large and complex sites because performance is limited to the local system capabilities. Q3: Data integration and cloud computing. Q4: Both computing layers compute and archive different KPI.

C. UC3: Historical data visualization and correlation

Users must be able to correlate energy, automation, and external data to identify inefficiencies. Especially, energy consumption behaviors should be analyzed using different timescales to identify peak loads that may result in surcharges from the energy provider [21].

Q1: Increase energy consumption awareness. Q2: Hard with multi-site correlation, because without the cloud, this involves having to transfer and import data between systems. Q3: Data integration and easy access to data. Q4: With the Dashboard, the user can correlate and visualize external and historical data stored in the Distributed Persistent Storage (see Figure 4).



Fig. 4. Analytics dashboard to visualize and correlate historical data and previously calculated KPI.

D. UC4: Pattern matching and data usage analysis

The system must derive intelligence from data streams to enable low-latency decisions in response to changing conditions. For example, using pattern matching techniques, the system must be able to unveil patterns in a data stream of events, such as abnormal behaviors, or identify different stages of production to quantify idle time [21].

TABLE III
SUMMARY OF THE ENERGY PERFORMANCE INDICATORS (KPI) USED IN THE ENERGY CLOUD PROJECT [21]

KPI Indicator	Focus	Description
Power	Energy consumption	Instantaneous or average power used by a process
Energy consumption	Energy consumption	Energy input into a process during a defined time period
Production energy consumption	Energy consumption	Energy consumption per manufactured product (items or units)
Energy costs	Energy costs	Monetary cost of energy used including fixed and variable components
Production energy costs	Energy costs	Energy costs per manufactured product
Energy losses	Energy efficiency	Energy use associated with non-value adding process steps or operating states
Energy efficiency	Energy efficiency	Ratio between the total energy used and the one used only in production

Q1: Increase industrial energy efficiency by identifying inefficiencies. Q2: Hard for large and complex sites because performance is limited to the local system's capabilities. Q3: Data integration and cloud computing. Q4: The Energy Analytics layer continuously analyzes and computes incoming data and perform complex computations (e.g. MR-AVF) to autonomously derive more knowledge.

E. UC5: Benchmark energy usage and efficiency

Benchmarks for similar equipment should be facilitated. Benchmarks should be available, stating where other sites or even other companies with the same challenges stand, in order to increase energy efficiency with the same process quality [12].

Q1: Increase energy efficiency by providing benchmark metrics. Q2: Very hard, because on premises, there may not exist the necessary statistic population number to benchmark. Q3: Data integration and cloud computing. Q4: On-demand data mining jobs in the Energy Analytics layer can browse through the multiple datasets and compare energy usage profiles.

V. CONCLUSIONS AND FURTHER WORK

Following the current industry trends regarding smart and future production, this work focuses how industrial EMS can be enhanced with the use of the cloud. We propose a novel architectural solution to address the industrial needs with current EMS implementations and to cope with the big data challenges involved. Our research shows that the application of these technologies can raise new opportunities for energy and cost savings. For further work, we intent to perform extensive performance benchmarking tests using real energy data from ABB's sites.

ACKNOWLEDGMENTS

This work was supported by national funds through *Fundação para a Ciência e Tecnologia* (FCT), under the project PEst-OE/EEI/LA0021/2013.

REFERENCES

- [1] International Energy Agency, "World energy outlook 2008," 2008.
- [2] H. P. Inamdhar and R. P. Hasabe, "It based energy management through demand side in the industrial sector," in *INCACEC 2009*, 2009.
- [3] S. Mohagheghi and N. Raji, "Intelligent demand response scheme for energy management of industrial systems," in *Industry Applications Society Annual Meeting (IAS), 2012 IEEE*, 2012, pp. 1–9.
- [4] T. Fiedler and P.-M. Mircea, "Energy management systems according to the ISO 50001 standard — Challenges and benefits," *2012 ICATE*, Oct. 2012.
- [5] B. Saha and D. Srivastava, "Data quality: The other face of big data," in *Data Engineering (ICDE), 2014 IEEE 30th International Conference on*, March 2014, pp. 1294–1297.
- [6] Y. Jadeja and K. Modi, "Cloud computing - concepts, architecture and challenges," in *2012 International Conference on Computing, Electronics and Electrical Technologies (ICCEET)*. IEEE, Mar. 2012.
- [7] P. Mell and T. Grance, "The NIST Definition of Cloud Computing, Recommendations of the National Institute of Standards and Technology," *National Institute of Standards and Technology*, 2011.
- [8] W. Voorsluys, J. Broberg, and R. Buyya, "Introduction to cloud computing," *Cloud computing: Principles and Paradigms*, 2011.
- [9] N. Motegi, M. A. Piette, S. Kinney, and K. Herter, "Web-based Energy Information Systems for Energy Management and Demand Response in Commercial Buildings," *Lawrence Berkeley National Laboratory*, 2003.
- [10] J. Byun, I. Hong, and S. Park, "Intelligent Cloud Home Energy Management System Using Household Appliance Priority Based Scheduling Based on Prediction of Renewable Energy Capability," vol. 58, no. 4, pp. 1194–1201, 2012.
- [11] I. Hong, J. Byun, and S. Park, "Cloud Computing-Based Building Energy Management System with ZigBee Sensor Network," in *Innovative Mobile and Internet Services in Ubiquitous Computing (IMIS), 2012 Sixth International Conference on*, 2012, pp. 547–551.
- [12] K. Bunse, M. Vodicka, P. Schönsleben, M. Brühlhart, and F. O. Ernst, "Integrating energy efficiency performance in production management - gap analysis between industrial needs and scientific literature," *Journal of Cleaner Production*, vol. 19, no. 6–7, pp. 667–679, 2011.
- [13] P. Thollander and M. Ottosson, "Energy management practices in Swedish energy-intensive industries," *Journal of Cleaner Production*, vol. 18, no. 12, pp. 1125–1133, 2010.
- [14] O. Givehchi, H. Trsek, and J. Jasperneite, "Cloud computing for industrial automation systems — A comprehensive overview," *2013 IEEE 18th Conference on Emerging Technologies & Factory Automation (ETFA)*, pp. 1–4, Sep. 2013.
- [15] J. Arinez and S. Biller, "Integration requirements for manufacturing-based Energy Management Systems," in *Innovative Smart Grid Technologies (ISGT)*, 2010, pp. 1–6.
- [16] X. Ma, R. Cui, Y. Sun, C. Peng, and Z. Wu, "Supervisory and Energy Management System of large public buildings," *2010 IEEE International Conference on Mechatronics and Automation*, pp. 928–933, Aug. 2010.
- [17] W. G. April, "Recommendations for implementing the strategic initiative Industrie 4.0," (*White Paper*), no. April, 2013.
- [18] N. Marz and J. Warren, *Big Data: Principles and Best Practices of Scalable Realtime Data Systems*. Manning Publications Company, 2013.
- [19] T. Goldschmidt, A. Jansen, H. Koziolok, J. Doppelhamer, and H. Pei-Breivold, "Scalability and robustness of time-series databases for cloud-native monitoring of industrial processes," in *Proceedings 7th IEEE Int. Conf. on Cloud Computing (IEEE CLOUD 2014) Industry Track*. IEEE, July 2014.
- [20] A. Koufakou, J. Secretan, J. Reeder, K. Cardona, and M. Georgiopoulos, "Fast parallel outlier detection for categorical datasets using mapreduce," in *IJCNN 2008*, June 2008.
- [21] K. Vikhorev, R. Greenough, and N. Brown, "An advanced energy management framework to promote energy awareness," *Journal of Cleaner Production*, vol. 43, pp. 103–112, Mar. 2013.