ENERGY 2014

The Fourth International Conference on Smart Grids, Green Communications and IT Energy-aware Technologies

ISBN: 978-1-61208-332-2

April 20 - 24, 2014

Chamonix, France

ENERGY 2014 Editors

Steffen Fries, Siemens AG, Germany

Petre Dini, Concordia University, Canada / China Space Agency Center, China

ENERGY 2014

Foreword

The Fourth International Conference on Smart Grids, Green Communications and IT Energy-aware Technologies (ENERGY 2014), held between April 20 - 24, 2014 in Chamonix, France, continued the event considering Green approaches for Smart Grids and IT-aware technologies. It addressed fundamentals, technologies, hardware and software needed support, and applications and challenges.

There is a perceived need for a fundamental transformation in IP communications, energy-aware technologies and the way all energy sources are integrated. This is accelerated by the complexity of smart devices, the need for special interfaces for an easy and remote access, and the new achievements in energy production. Smart Grid technologies promote ways to enhance efficiency and reliability of the electric grid, while addressing increasing demand and incorporating more renewable and distributed electricity generation. The adoption of data centers, penetration of new energy resources, large dissemination of smart sensing and control devices, including smart home, and new vehicular energy approaches demand a new position for distributed communications, energy storage, and integration of various sources of energy.

We take here the opportunity to warmly thank all the members of the ENERGY 2014 Technical Program Committee, as well as the numerous reviewers. The creation of such a high quality conference program would not have been possible without their involvement. We also kindly thank all the authors who dedicated much of their time and efforts to contribute to ENERGY 2014. We truly believe that, thanks to all these efforts, the final conference program consisted of top quality contributions.

Also, this event could not have been a reality without the support of many individuals, organizations, and sponsors. We are grateful to the members of the ENERGY 2014 organizing committee for their help in handling the logistics and for their work to make this professional meeting a success.

We hope that ENERGY 2014 was a successful international forum for the exchange of ideas and results between academia and industry and for the promotion of progress in the fields of smart grids, green communications and IT energy-aware technologies.

We are convinced that the participants found the event useful and communications very open. We also hope the attendees enjoyed the charm of Chamonix, France.

ENERGY 2014 Chairs:

ENERGY Advisory Chairs

Stefan Mozar, CCM Consulting / CQ University - Sydney International Centre, Australia
Mardavij Roozbehani, Massachusetts Institute of Technology, USA
Mark Apperley, University of Waikato, New Zealand
Steffen Fries, Siemens, Germany

ENERGY Industry Liaison Chairs
Dave Cavalcanti, Philips Research North America, USA
Marco Di Girolamo, Hewlett-Packard Company - Cernusco sul Naviglio, Italy
Dragan Obradovic, Siemens AG, Germany

ENERGY Special Area Chairs on Nano-Grids
Peter Müller, IBM-Zurich, Switzeland

ENERGY Special Area Chairs on Smart Grids
Ritwik Majumder, ABB AB / Corporate Research Center - Vasteras, Sweden

ENERGY Special Area Chairs on IT-energy- aware, Planning
Brian P. Gaucher, IBM Research Division - Yorktown Heights, USA

ENERGY Special Area Chairs on Grid, Green Communication
Gargi Bag, ABB Corporate Research, Sweden
Ken Christensen, University of South Florida, USA

ENERGY Special Area Chairs on Vehicular
Grzegorz Swirszcz, IBM Watson Laboratory, USA

ENERGY 2014

Committee

ENERGY Advisory Chairs

Stefan Mozar, CCM Consulting / CQ University - Sydney International Centre, Australia
Mardavij Roozbehani, Massachusetts Institute of Technology, USA
Mark Apperley, University of Waikato, New Zealand
Steffen Fries, Siemens, Germany

ENERGY Industry Liaison Chairs

Dave Cavalcanti, Philips Research North America, USA
Marco Di Girolamo, Hewlett-Packard Company - Cernusco sul Naviglio, Italy
Dragan Obradovic, Siemens AG, Germany

ENERGY Special Area Chairs on Nano-Grids

Peter Müller, IBM-Zurich, Switzeland

ENERGY Special Area Chairs on Smart Grids

Ritwik Majumder, ABB AB / Corporate Research Center - Vasteras, Sweden

ENERGY Special Area Chairs on IT-energy- aware, Planning

Brian P. Gaucher, IBM Research Division - Yorktown Heights, USA

ENERGY Special Area Chairs on Grid, Green Communication

Gargi Bag, ABB Corporate Research, Sweden
Ken Christensen, University of South Florida, USA

ENERGY Special Area Chairs on Vehicular

Grzegorz Swirszcz, IBM Watson Laboratory, USA

ENERGY 2014 Technical Program Committee

Amir Abtahi, Florida Atlantic University - Boca Raton, US
Nizar Al-Holou, University of Detroit Mercy, USA
Lachlan Andrew, Swinburne University of Technology – Melbourne, Australia
Luca Ardito, Politecnico di Torino, Italy
Mehdi Bahrami, University of California, Merced, USA
Chakib Bekara, Fraunhofer Fokus Institute – Berlin, Germany

Fabio Luigi Bellifemine, Telecomitalia, Italy
Siegfried Benkner, University of Vienna, Austria
Frede Blaabjerg, Aalborg University, Denmark
Robert Brewer, Aarhus University, Denmark
Antonio Caló, NorTech Oulu -Thule Institute / University of Oulu, Finland
Davide Careglio, Universitat Politècnica de Catalunya - Barcelona, Spain
Mari Carmen Domingo, Barcelona Tech University, Spain
Dave Cavalcanti, Philips Research North America - Briarcliff Manor, USA
Antonin Chazalet, IT&Labs, France
Mohamed Cheriet, ETS/GreenStar - Montreal, Canada
Howard Choe, Raytheon - McKinney, TX, USA
Sangho Choe, The Catholic University of Korea, Korea
William Cheng-Chung Chu, Tunghai University, Taiwan
Delia Ciullo, EURECOM Sophia Antipolis, France
Peter Corcoran, College of Engineering & Informatics, NUI Galway, Ireland
Margot Deruyck, Ghent University/IBBT, Belgium
Marco Di Girolamo, Hewlett-Packard Company, Italy
Yong Ding, Karlsruhe Institute of Technology (KIT), Germany
Ron Doyle, IBM | Software Group Strategy and Technology - RTP, USA
Venizelos Efthymiou, Electricity Authority of Cyprus / Sustainable Energy Research Centre at University of Cyprus, Cyprus
Tullio Facchinetti, University of Pavia, Italy
Farid Farahmand, Sonoma State University, USA
Eugene A. Feinberg, Stony Brook University - New York, USA
Alexandre Peixoto Ferreira, IBM Austin Research Laboratory, USA
Steffen Fries, Siemens Corporate Technology - Munich, Germany
Brian P. Gaucher, IBM Research Division - Yorktown Heights, USA
Xiaohu Ge, Huazhong University of Science and Technology, China
Erol Gelenbe, Imperial College London, UK
Hamid Gharavi, National Institute of Standards and Technology, USA
Georgios B. Giannakis, University of Minnesota - Minneapolis, USA
Manimaran Govindarasu, Iowa State University, USA
Mesut Günes, Freie Universität Berlin, Germany
Rune Gustavsson, KTH, Sweden
Helmut Hlavacs, University of Vienna, Austria
Chun-Hsi Huang, University of Connecticut - Storrs, USA
Zhiyi Huang, University of Otago, New Zealand
Muhammad Ali Imran, University of Surrey, UK
Canturk Isci, IBM TJ Watson Research Center, - New York, USA
Philip Johnson, University of Hawaii - Honolulu, USA
Aravind Kailas, University of North Carolina - Charlotte, USA
Young Sun Kim, Korea Electrotechnology Research Institute (KERI), Korea
Thierry E. Klein, Bell Labs / Alcatel-Lucent, USA
Janine Kniess, Santa Catarina State University, Brazil
Sileshi Kore, Dilla University, Ethiopia
Dejan Kostic, Institute IMDEA Networks, Spain
Paul J. Kuehn, Institute of Communication Networks and Computer Engineering - University of Stuttgart, Germany

Dimosthenis Kyriazis, University of Piraeus, Greece
DongJin Lee, University of Auckland, New Zealand
Eugene Litvinov, ISO New England, USA
Shanshan Liu, Electric Power Research Institute, USA
Marco Lützenberger, Technische Universität Berlin, Germany
Priya Mahadevan, PARC, USA
Thair Shakir Mahmoud, Edith Cowan University - Western Australia, Australia
Mikko Majanen, VTT Technical Research Centre of Finland, Finland
Michael Massoth, University of Applied Sciences - Darmstadt, Germany
Satoshi Matsuoka, Tokyo Institute of Technology, Japan
Rafael Mayo-García, CIEMAT, Spain
Jean-Marc Menaud, Ecole des Mines de Nantes, France
George Michailidis, University of Michigan, USA
Marilena Minou, Athens University of Economics & Business, Greece
Nicolas Montavont, Telecom Bretagne, France
Christina Morfopoulou, Imperial College London, UK
Daniel Mossé, University of Pittsburgh, USA
Gero Mühl, Universitaet Rostock, Germany
Masayuki Murata, Osaka University, Japan
Bruce Nordman, Lawrence Berkeley National Laboratory, USA
Dragan Obradovic, Siemens AG - München, Deutschland
Jacob Østergaard, Technical University of Denmark, Denmark
Cathryn Peoples, University of Ulster, UK
Mirek Piechowski, Meinhardt (Vic) Pty Ltd. - Melbourne, Australia
Philip W. T. Pong, The University of Hong Kong, Hong Kong
Evangelos Pournaras, Delft University of Technology, Netherlands
Judy Qiu, Indiana University, USA
Enrique S. Quintana-Orti, Universidad Jaume I, Spain
Shankar Raman, Indian Institute of Technology - Madras, India
Jan Richling, TU Berlin, Germany
Jan Ringelstein, Fraunhofer Institut für Windenergie und Energiesystemtechnik IWES, Germany
Darren Robinson, University of Nottingham, UK
Ivan Rodero, Rutgers University - Piscataway, USA
Sebnem Rusitschka, Siemens AG - München, Germany
Eliot Salant, IBM Haifa Research Labs / Haifa University, Israel
Dave Saraansh, University of Bristol / Toshiba Research Europe Limited, UK
Dirk Uwe Sauer, ISEA / RWTH - Aachen University, Germany
Andrey V. Savkin, The University of New South Wales - Sydney, Australia
Harald Schrom, Technische Universitaet Braunschweig , Germany
Sandra Sendra, Universidad Politécnica de Valencia, Spain
Björn Skubic, Ericsson Research, Sweden
Gerard Smit, University Of Twente - Enschede, The Netherlands
Fernando Solano, Warsaw University of Technology, Poland
Pavel Somavat, CCS Haryana Agricultural University, India
Grzegorz Swirszcz, IBM Watson Laboratory, USA
Nick A. Tahamtan, Vienna University of Technology, Austria
Zhibin Tan, East Tennessee State University, USA
Lefteri H. Tsoukalas, Purdue University - West Lafayette, USA

Copyright Information

Table of Contents

Enabling Distributed Meter Data Management Using Mediation System

Nipun Sharma
Research & Development, Center of Excellence
Mediation and Device
Ericsson India Global Services Private Limited
Gurgaon, Haryana, India
nipun.sharma@ericsson.com

Sandeep Akhouri
Research & Development, Center of Excellence
Mediation and Device
Ericsson India Global Services Private Limited
Gurgaon, Haryana, India
sandeep.akhouri@ericsson.com

***Abstract*—This paper proposes the use of mediation system as a light weight embedded software module in Distributed Meter Data Management (DMDM) platform in Utility domain. In DMDM platform, a head end system would collect the information from all the connected devices in a house/neighbourhood, perform meaningful processing and forward the processed information to centralized repository (or centralized MDM) system. Mediation as a module would be embedded in each head end system to perform distributed data management and send the limited yet useful information to centralized repository reducing the processing and storage requirement on the whole. Processed data by mediation system at head end system would also be available to the end customer(s) for real time updates rather than waiting for processing of data at centralized MDM in batch process.**

Keywords-Meter Data Management; Smart Grid; Data Management; Mediation system

I. INTRODUCTION

Utilities sector is undergoing a massive transformation driven by government regulation and commercial factors. Global leaders across the world have committed to reduce emissions significantly by 2020. EU targets to reduce the CO_2 footprint by 20% below 1990 level. United States targets to reduce CO_2 footprint by 17% below 2005 level. China targets to reduce the CO_2 level of GDP by 40 to 45 % from 2005 level. Smart grid deployments improve the efficiency within a utility's network, open up new revenue streams and provide enhanced customer services. Over the past few years, the rollout of large-scale smart meter deployments and the Automated Meter Infrastructure (AMI) that supports connectivity and bi-directional information flow has become the most promising and well proven domain to modernize.

Smart grid along with smart metering would play a leading role in achieving the defined targets. As per a leading research organization [2], global installation of smart meters will be more than 95 million in CY2013 primarily dominated by developed nations across North America and Europe. As per projected growth, there would be more than 450 million smart meters deployed across the globe in CY2018 primarily lead by Asia. Investment for smart grids would exceed $278 billion dominated by China, which accounted 22% of world's electricity production in CY2011 [3]. On an average, Compounded Annual Growth Rate (CAGR) for utility space is expected around 12-17% annually [2]. Also, going from one meter reading a month to smart meter readings every 15 minutes works out to 96 million reads per day for every million meters. The result is a 3,000-fold increase in data that can be overwhelming if not properly managed [6]. Telecom vendors, being established player in all communicating world, could play a pivotal role to further empower utility vendors with their existing expertize for effective and efficient deployments.

The transition to the smart grid is fundamentally driven by market forces. The smart grid's ability to make better use of existing assets, enhance reliability and minimize costly environmental impacts are all market forces that have substantial economic value for all the involved stakeholders like standardization bodies, government institutes, utility companies as well as end customer.

From end customer perspective, Smart Meter enables customers with greater control over their energy use and costs by allowing them to monitor their energy use online and determining which activities are contributing to their bills. Customers can online view projected monthly bills based on their current usage allowing action to be taken before prior to bill generation. Near real-time processing of this information can provide valuable insight e.g., smart meters that collect information; say every minute can help identify appliances that consume the maximum electricity.

Typically, Meter Data Management is centrally located where all the raw information from different smart meters are collected and saved for processing the information in the required format. This kind of centralized MDM solution leads to big requirement on storage as well as non-real-time update to end consumers.

In this paper, we discuss the use of mediation systems as a data management platform for Distributed Meter Data Management. Mediation systems allow events generated in asynchronous as well as synchronous fashion to be collected

 ISBN: 978-1-61208-332-2

and processed. The events can be enriched and consolidated as per the requirement of ecosystem. Mediation system provides extensive capabilities for data management including DB support, archiving, complex data calculation, data analysis, auditing and version data management. Near real-time analysis and visualization of events can produce alerts and notifications that can be used as a feed back to end customer for real-time action.

To summarize, we will start by introducing Mediation system and elaborate on the importance of mediation system in telecom and IT domain. We then discuss the challenges posed for Meter data management due to ever growing data from smart meters and how traditional features of mediation including data management and respective processing capabilities could lead to efficient and effective deployment of scalable Distributed Meter Data Management in utility space.

II. MEDIATION SYSTEM

Traditionally, a mediation system collect data from the network, enrich that data via processing and then delivers the enriched data to the desired system as highlighted in Figure 1. Mediation system could be generalized as software component used for data transformation depending upon the business requirement.

Operators manage a large computing infrastructure comprising of a diverse set of hardware, software and applications. Mediation system plays a critical role in accurate and timely charging & billing of a variety of services including voice, data and multi-media services provided by an operator. These systems operate in a highly complex, multi-vendor environment and provide reliable operations, high performance and secure access in a data-intensive environment.

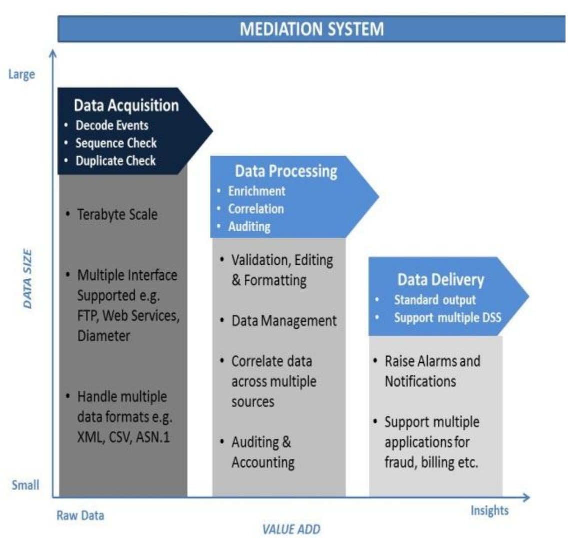

Figure 1. Overview of Mediation System.

Mediation systems collect data from multiple networks - 2G, 2.5G, 3G, IP, Fixed line. As part of basic processing, the events generated in different formats by different NE are decoded from XML, CSV, ASN.1, etc. to an internal data format for processing. Validation of event records is performed to ensure that they are received in the correct order [7].

Events can be enriched with information from external databases e.g., subscriber profile information can be added to the usage data. Event records originating from the same NE can be aggregated by the mediation system e.g., a long running call might generate partial CDR that can be aggregated to produce a single billable CDR. Event Records that are generated across NE can be correlated using a unique key (say, MSISDN, IMSI) to stitch together the multiple event records into one billable CDR [7].

Since billing can be based upon duration, volume, content value, time, Quality of Service (QoS) or a combination of these parameters, custom processing logic can be written to define these business rules. The logic also defines interfaces with external systems, error handling and configures notifications and alarms that can be raised.

Finally, the data records are encoded in a standardized format for downstream OSS/ BSS applications. Thus, these applications are completely abstracted from the multiple data formats, transmission protocols and receive standardized and complete records.

Near real-time and real-time mediation systems handle events in a continuous streaming format. Real-time systems can process information between 1 to 10 milliseconds whereas near real-time systems take up to a second for processing. As soon as events occur, usage data can immediately flow through the mediation system towards billing, network planning, fraud management and / or other OSS/ BSS systems.

Mediation system is software component which could easily be deployed, co-located with other telecommunication as well as IT based solutions on Commercial-off-the-shelf (COTS) hardware.

III. SMART GRID & DISTRIBUTED METER DATA MANAGEMENT

On high level, Smart Grid could be defined as convergence of IT, automation and communication technology in Energy Distribution Network. Smart grid is a modernized electrical grid that uses information and communications technology to gather and act on information, such as information about the behaviors of suppliers and consumers, in an automated fashion to improve the efficiency, reliability, economics and sustainability of the production and distribution of electricity.

Smart grids are intelligent energy networks which further add intelligence through new bi-directional communications and computer controls. There are many smart grid technology areas each consisting of sets of individual technologies that span the entire grid, from generation through transmission and distribution to various types of consumers like electricity, gas, etc. [9].

Meter data management (MDM) refers to a business process in the Smart Grid infrastructure primarily responsible for data management and VEE component as highlighted by Ericsson Smart Grid offering [4]. VEE component stands for Validation, Editing and Estimation modules.

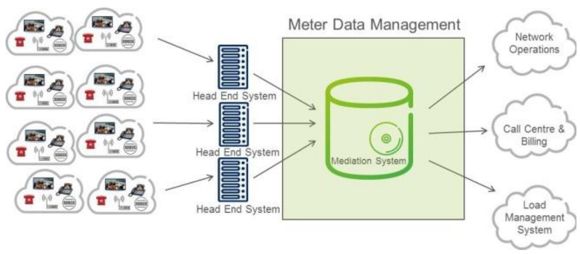

Figure 2. Deployment view for Meter Data Management.

As highlighted above, MDM is positioned between Head End systems collecting information from customer's premises and operator's backend system like Network operations, billing, load management system to name a few. In centralized MDM solution, all the raw information is collected at one centralized point and processed before forwarding it to respective backend system. This leads to extensive data flow among various nodes in MDM solution, support to handle BIG data as well as delayed update to the end customers including consumed unit, consumption pattern, etc.

MDM systems provide a database repository that automates and streamlines the complex process of collecting meter data from multiple collection points supported and deliver the data in an appropriate format to a utility billing system and other smart grid related applications.

As per IEC standard [1], the meter data management system is used to provide a common repository, and point of management and access of meter data that is collected from disparate metering systems. In addition to data aggregation, quite often the MDM will also make an effort to scrutinize the data collected from the various metering systems, and provide a validating, editing, and estimating (VEE) capability.

"VEE" means validation, estimating and editing of Meter Reads to identify and account for missed and inaccurate reads used to derive billing data.

Data Management & Repository is key component in MDM module. Data repository consist of saving raw data as well as processed data in required formats for long term storage as part of legal requirement. Stored data would be further analyzed in case of conflicts/complaints. This data should also be used for estimation purpose.

In real scenarios, more functions have been embedded in MDM module as highlighted in Figure 3. Various functions which are part of MDM solution offerings are Billing, Customer information, Revenue assurance and Auditing & Reporting, etc.

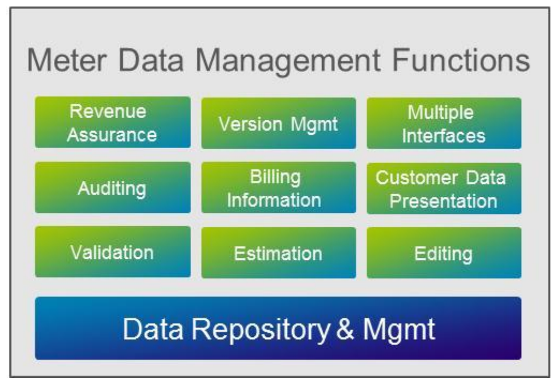

Figure 3. Functional breakdown of Meter Data Management (MDM) processes.

Billing customers is a critical process for a utility space and lot of resources, financial investment and time is invested since it is directly connected to revenue. Going forward, more efficient billing techniques would be implemented like Time of Use Billing, Critical peak pricing to reduce peak load and even out the load.

Auditing and Reporting process is an absolutely necessary module to manage and be in control of all the events performed in the system. Auditing should record all the modification, involving processed data like when, how and why an action was performed. End to End auditing for all the historical data and respective reporting framework are important aspects of MDM module.

One of the important advantages of Smart grid over traditional grid deployment is real-time meaningful update for the end consumers. MDM system interfaces with a number of different smart grid related applications supporting multiple interfaces.

Revenue assurance is the module that ensures a Utility Service Provider (USP) to accurately capture revenue for all services rendered. Possibility to track events through the business logic or to report events from within the processing flows are few of the mandatory requirements of the domain.

 ISBN: 978-1-61208-332-2

In distributed MDM solution, primary responsibility of MDM module i.e. data repository and VEE component are migrated from centralized MDM platform to Head End system leading to win-win situation for all the stakeholders as highlighted in DMDM Patent [5].

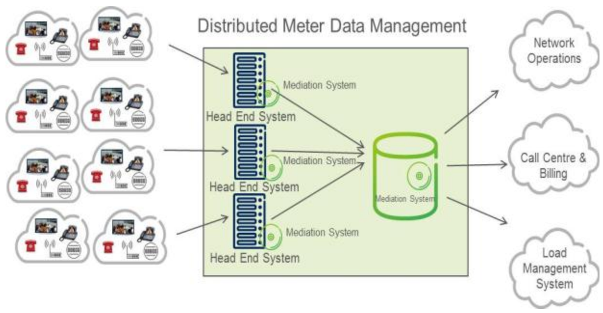

Figure 4. Deployment view for Distributed Meter Data Management.

IV. MEDIATION SYSTEM IN DISTRIBUTED METER DATA MANAGEMENT

As highlighted, Meter Data Management (MDM) refers to a business process in the Smart Grid infrastructure primarily responsible for data management and VEE component. Mediation system provides the inbuilt capability, and flexibility to realize the business requirements of Distributed Meter Data Management in scalable manner. In the below section, we will cover the two primary aspect of DMDM i.e. Data Acquisition and Delivery and Data Processing and how mediation system is a suitable fit for implementing DMDM module.

A. Data Acquisition and Delivery

Smart Grid Infrastructure is projected to expand rapidly, leading to complex network supporting new protocols, new technologies and new formats in real time.

As IEC recommends [1] , the system interfaces of a compliant utility inter-application infrastructure has been defined using Unified Modelling Language (UML). The Extensible Markup Language (XML) is a data format for structured document interchange particularly on the Internet. One of its primary uses is information exchange between different and potentially incompatible computer systems. XML is thus well-suited to the domain of system interfaces for distribution management. Meter Data Management module is supposed to interact with almost all the logical entity modules including, but not limited to, Metering system, Planning and Scheduling, Customer Information and billing, Work management, Network Operations and Load Management System.

Additionally, central to the customer benefit is the provision and presentation of data. Data can be presented in a number of different formats and on a number of different media (smartphone, web, customer bill, in-home display unit, and so on). Utilities with a commitment towards improving customers' visibility into their personal consumption behavior must ensure that Mediation system interfaces with a number of different applications supporting customer information.

Mediation systems provide additional applications and open API standards to integrate into customer management systems, Traffic management and network monitoring. Mediation also provides the support for flow based real time protocols which are extensively used in network monitoring domain and to build over the top applications as per customer specific needs like netFlow and related flow based proprietary protocols [7].

Mediation systems could be used as light weight Service-oriented architecture (SOA) application to communicate between multiple modules on different protocol like Web services, SOAP, FTP, and SFTP. Support for XML format is an integral part of the communication framework.

In MDM deployment, the requirement on communication interfaces would be same, if not more, to support different technologies and products available at customer premises in future.

As IEC suggests [1], MultiSpeak is also a distribution focused specification / standard that defines standardized interfaces among software applications commonly used by electric utilities. Also, it defines details of data that need to be exchanged between software applications in order to support different processes commonly applied at utilities. MultiSpeak recommends file based transfers and SOAP messages over HTTP, TCP/IP sockets.

There is no specification on transport level communication however flow of the information should be via standard data format, named XML, which is ubiquitously used in IT domain.

Web Services is widely used technique to send web request to support inter-operable machine-to-machine interaction over a network. It has an interface described in a machine-process able format (specifically Web Services Description Language WSDL). The basic Web Services platform is XML + HTTP and both these technologies are well supported by Mediation System.

Additionally, Mediation system also supports wide variety of data formats like ASN.1/BER, Comma-Separated-Value (CSV), XML, MIXED format to name few [4].

To minimize the amount of data transmission between different logical entities in utility domain mediation system also support transferring of compressed data using various standard techniques like zip, gzip, pkzip, etc. This will overall reduce the load on the network however without compromising the quality of the processing

B. Data Processing

Data processing is key responsibility of Meter Data Management allowing operator specific business logic to be implemented in faster and reliable fashion. Data processing includes, but not limited to, Validation, Estimation, Editing, Auditing, Billing data management, Version data Management & Revenue Assurance functionality.

1) Validation, Estimation & Editing

Validation process consists of various checks to ensure the integrity of the received data. The mediation system checks meter readings for anomalies. Validation process could further be broken down on syntactic check, semantic check and business logic related checks. Mediation system can check the syntax of each data transfer that is received from an AMI. This can be expected to include checking that the data structure is consistent with the required standards in utility space and respective checksum calculations.

Mediation system can also verify the semantics of each data transferred that is received from an AMI. For efficiency (e.g. loading prioritization), the data in any one data transfer should be from the same meter and the data should also correlate with information in the respective header section.

Also, Mediation system should perform checks to ensure that all meters for which data is provided are currently active.

Once the data is received and validated by the Head end system, the raw information is stored in the Database module for further processing. Traditional use case of Mediation in IT domain is to perform Extract-Transform-Load (ETL) functionality which is used for data integration, data migration, and data analysis and data management. ETL is an essential component used to load the data into the data warehouses from different sources in multiple formats across different domains.

Each operator would need to define additional business logic as per specific utility operators need. Time gap check, Interval Data Management and missing Interval checks are three important issues encountered and are expected to be well handled by MDM system.

Mediations system also support Record sequence number check, Time Gap, Zero byte file check, Scheduled Collection and Post collection actions which are best fit feature for supporting VEE module.

Mediation system is capable of data estimation techniques when actual Meter Reads are not available in order to create estimated Meter Reads and Billing Quantities. The method of estimation shall be automated.

Mediation system accommodates the ability to apply customized estimation techniques for non-conforming commercial and industrial loads, or for groups of smart meters that require special rules, or that should not have estimation applied to them.

Estimation techniques could further be classified based on two different estimation techniques on

- Historical Data
- Linear Interpolation

Historical estimation routine calculates average daily load shapes using data from historical reference days. It derives interval-by-interval averages from valid intervals from the 3 nearest days of the same or like day of the week, accounting for holidays.

In the event the correction occurs at the very beginning of the data being validated, Linear Interpolation will attempt to find the last interval from the previous reading group in order to use it as the starting “anchor point” for the estimation. If there is no earlier reading group, Linear Interpolation will use the value from the first good interval for all corrected intervals typically less than two hours [8]. Similarly, in the event the correction occurs at the end of the data set, the routine will use the value from the last good interval for all corrected intervals.

Mediation system provides rich set of functions based on common language like Java where user could build the required business logic based on the specific required.

Additionally, Database interfaces are also supported to store the required information and perform the logical operations using procedures based on database. Relation Databases is supported by mediation systems primarily supporting JDBC based database like Oracle, PostGreSQL, etc.

Mediation system provides the editing capabilities of Meter data by both the MDM operator and authorized external users, though the mechanisms may vary. Viewing and editing of Meter data by the external entity will be restricted to respective owners and related meter ID numbers for which the external entity has been identified as primary authority.

2) Auditing, Version Management & Revenue Assurance

Auditing is defined as step-by-step record by which accounting data can be traced to their source. Every change made to the raw as well as processed information is tracked, along with who made the change, why they made it, and references to problems fixed, or enhancements introduced, by the change.

Auditing is an important, legal and operational requirement in the telecom and IT world. Mediation system provides detailed auditing capabilities with options to export the information to the external system for further analysis.

Versioned Data management provides the capability to provide the snapshot of the meter load at particular time interval in the past to resolve billing and legal disputes. Mediation system provides log records indicating which user or VEE process made changes to the data. For example, if a reading changes five times, the MDM creates five versions of that reading, each of which also has a reference time period, indicating when it was the current version.

Revenue assurance is a critical business process for utility industry since an efficient revenue assurance implementation could minimize the revenue leakage and maximize the revenue. Mediation system is integral part of revenue assurance function across OSS and BSS domain in Telecom. It provides inbuilt capabilities to store the selective information from the meter reads in the file or database, which could further be used for revenue assurance reports.

V. CONCLUSION

Driven by market forces including rising energy costs and technological advancements, Smart Grid along with Smart meters are poised to be an important revenue stream for Information and Communication Technology provider. This paper highlights the inherent advantages, in terms of features, that traditional mediation systems would provide for handling metering data from Smart meters in Distributed Meter Data Management solution.

Meter Data Management application is primarily responsible for data repository and management including Validation, Estimation and Editing. Mediation systems are already equipped to support centralized as well as distributed Meter Data Management.

In the future, distributed meter data management, enabled using mediation system, would play a pivotal role for Smart grid deployments across the globe.

ACKNOWLEDGMENT

We would like to thank Rakesh Pandey, Munish Agarwal and Manas Kumar Jena for their valuable feedback, support and guidance.

REFERENCES

[1] International Electrotechnical Commission (IEC) 61968-9 Ed.1 Application integration at electric utilities – System interfaces for distribution management.
https://webstore.iec.ch/Webstore/webstore.nsf/ArtNum_PK/48719?OpenDocument (Janurary 2014)

[2] ABI research report “THE SMART GRID AS A SERVICE”. https://www.abiresearch.com/research/product/1013765-the-smart-grid-as-a-service-a-game-changer/ (October 2013)

[3] ABI research report “SMART GRIDS”. https://www.abiresearch.com/research/product/1009947-smart-grids/ (October 2013)

[4] Smart Grid & Mediation product offering from Ericsson AB
http://www.ericsson.com/ourportfolio/products/multi-mediation (Novemeber 2013)
http://www.ericsson.com/ourportfolio/energy-and-utilities-industry/smart-grid-communications (Novemeber 2013)

[5] Bruce Christopher Angelis, James Alexander Papp and Timothy James Driscoll on “Distributed Meter Data Management,” Patent Pub. No.: US 2012/0326883 A1 (Octeber 2013)

[6] IBM Software White paper “Managing big data for smart grids and smart meters”. www.smartgridnews.com/artman/uploads/1/IBM_anaytics_paper.pdf (Janurary 2014)

[7] M S Rohmad,R Kaspin,Mazani Manaf and Farok Azmat "Network Mediation for Future IP Billing System: An Overview" http://ieeexplore.ieee.org/xpl/abstractCitations.jsp?arnumber=4339313 (Oct 2013)

[8] Don Matheson,Chaoyang Jmg and Frank Monforte, "Meter Data Management for the Electricity Market " 8th International Conference on Probabilistic Methods Applied to Power Systems, Iowa State University, September 12-16,2004

[9] Fadi Aloul, A. R. Al-Ali, Rami Al-Dalky, Mamoun Al-Mardini and Wassim El-Hajj “Smart Grid Security: Threats, Vulnerabilities and Solutions”, International Journal of Smart Grid and Clean Energy, vol. 1, no. 1, September 2012

An Open Data-Based Discrete Simulation Environment for Testing Smart Grid Messaging

Eric MSP Veith[1], Bernd Steinbach[2] and Jürgen Otten[3]

[1,3]Institute of Computer Science
Wilhelm Büchner Hochschule
Pfungstadt, Germany
e-mail: eric.veith@wb-fernstudium.de

[1,2]Institute of Computer Science
Freiberg University of Mining and Technology
Freiberg, Germany
e-mail: veith@mailserver.tu-freiberg.de

***Abstract*—Including more renewable energy sources in the energy mix will increase the necessity for a finer grained, automatic control of changes in the energy level. Any such software needs extensive testing before it can be released for general availability. Simulation environments will be a part in these testing stacks, but need realistic input data in order to yield expressive and therefore useful results. However, comprehensive input data is often not available or fragmented. We therefore propose a simulation environment that can use open data sources, switch them dynamically, and attribute the testing results with a possible data quality impact.**

Keywords*—*smart grid; simulation; messaging; open data; renewable energy sources

I. INTRODUCTION

The shift from traditional, fossil energy sources to renewable ones such as photovoltaic or wind poses challenges on the existing energy grid. The inclusion of these renewable energy sources becomes more and more difficult with higher number of deployments. A multi-agent based approach has been suggested, for example, in [1], in order to counter the increased complexity in grid management.

Simulation runs have been widely accepted as a means to test software, especially network-based, distributed software architectures. While unit testing can assert the correct functional behavior of a particular module, or unit, of the software being tested, the interaction of several instances of the complete software architecture cannot be asserted by unit testing alone. Several simulation environments have been proposed, both generic network simulators, such as OMNeT++ [2], and smart grid-specific solutions, such as MASGriP [3].

Descriptive and therefore successful testing always depends on the right choice of test data. In the case of renewable energy sources, a large portion of these simulation data is formed by weather measurements. Even source data regarding a single topic such as the aforementioned weather measurements can be heterogeneous, for example, when a portion of the data has a high resolution while another portion of it only offers a low resolution. But a testing environment usually does not explicitly care about the actual source of data; the task of asserting data quality therefore falls to the scientist. He, in turn, typically handles data quality differences by partitioning the source data and, therefore, doing different test runs.

Data quality is especially important when using open data. Using open data can be useful or even necessary for different reasons, like, e.g., budget limits. In this paper, we propose a simulation environment for smart grids that utilizes open data and annotates simulation runs based on these sources and can continue even if the data source is fragmented or incomplete.

II. MOTIVATION

In [4], we have proposed a lightweight smart grid messaging protocol for a distributed agent-based environment that, by design, treats all items within the energy grid equal, but forms microgrids automatically. Using this protocol, consumers and producers interact with each other and calculate demand and supply in a distributed manner. The nodes in the grid, which are represented by agents, commit themselves to deliver or consume a certain amount of energy based on the preceding message exchange. The goal of this system of distributed agents, which exhibit these messages, is to use renewable energy sources more effectively. In order to do so, short-term forecasts are used in a grid-wide, distributed planning phase that leads to an automated, more dynamic grid management.

Since the location of some types of renewable energy sources, such as wind farms, are dictated by the source itself—a wind farm must be built on sites where the wind currents are strong and steady—, we expect the general behavior of smart grid agent messaging to show patterns based on locations.

As such, a simulation environment for this kind of smart grid distributed agent approach must focus on map positions. Hence, it becomes tightly coupled to spatial information, which, in turn, allow us to treat simulated items more effectively. Consider, for example, wind speed measurements. These are valid for a certain area and thus can be applied to any number of items within this area. Without this spatial relationship, each simulated item would have to look up measurements individually, thereby causing more queries and calculations being made. In fact, each participant in our simulation environment needs to be locatable and thus automatically becomes an item on the map when entering the environment.

Having spatial data, i.e., a map, as the common basis also allow us to use realistic weather data. Although [5] suggests that many typical weather conditions can be synthetically modeled, the use of real measurements allow us to test agent behavior for location-specific weather conditions and phenomena.

 ISBN: 978-1-61208-332-2

The most precise information source for weather data is typically a national weather service. Often, however, data needs to be bought. Supplying data to a long-running simulation for a whole country can therefore imply a financial impact that is not desirable. The same can hold true for spatial data such as the position of wind parks or other power plants.

We therefore propose the use of open data within our simulation environment. But since open data can be less exact, care must be taken. Although this could simply mean setting up a separate simulation run, one would loose the internal state of the simulation environment. A combination of both the more exact but expensive national weather service's and the open data source can therefore be desirable. This combination, however, must be carefully augmented in order to try to assert the impact of using different data sources throughout one single simulation run.

Transparently switching between different data sources allows us to take advantage of the more exact data whenever possible and still enables us to have a long-running simulation, thereby observing our agent software's behavior and the message exchange caused by it over a longer time, which potentially yields more diverse data. The assessment additionally shows the impact this switching to a less exact data source had on a particular simulation run.

The remainder of this paper is structured as follows. We briefly outline the general architecture of the simulator in the following Section III, where we outline how the spatial indexing helps us to express the locality of certain events. Afterwards, in Section IV, we show how our proposed environment can be used for larger-scale simulations spanning multiple computing nodes. The extensive use of open data is described in Section V, followed by a discussion of the implications of using open data sources in Section VI. Finally, we conclude and show pointers to future work in the final Section VII.

III. Architecture of the Spatially-Indexed Simulator

Each simulation is controlled by a `Controller` class instance that forms the core of this time-discrete simulation environment. It is responsible for tracking the current time within the simulation and issuing events. `Event` objects are issued during a *tick* and reach all relevant participants within the simulation, which process them and finally return them to the controller. As soon as all have returned, the controller advances to the next time in the simulation at which events are scheduled.

The `Controller` also contains a list of all items participating in the simulation. These items are wrapped in `MapItem` classes. Since every simulation participant is ultimately a `MapItem`, this provides an unified interface between simulation controller and simulated item. By acting as adapters, these wrapper classes feed the native agent interfaces with the input data generated within the simulation environment.

When agents communicate with each other, the same wrapper technique is applied. The simulation provides virtual data connections, which facilitate the transition to the virtual simulation environment. The original agent software can therefore remain unchanged: testing is done as black-box testing. Consequently, agent code is not part of this paper.

Mainly the targets of events such as a simple time change or a simulated sensor reading are `MapItem` objects. This group is constituted of nodes within the power grid that are subject to the simulation and thus our agents that are being tested, namely, power lines, substations, consumers or power generators. However, all data sources that deliver input like, for example, weather data, are also a part of this group.

The simulation replaces actual hardware sensors of an agent by artificial stimuli. These stimuli are modeled as discrete events within the simulation. They also have coordinates attached to them: A wind speed measurement, for example, is valid for a particular region; the same applies to a load profile. The sources for artificial sensor data provided by the simulation are thus instances of the `MapItem` class, called `ValueGeneratingMapItem`. Their coordinates are constituted by the area they provide valid data for.

Consequently, each `Event` instance created by such a map item also has an *area of effect* for which it is valid. Any event is thus delivered to those map items that reside inside this area. Thus, events can be seen as tuples (A, V), containing the area of effect and a value. Typically, an event's area of effect equals the polygon forming the position of the map item. For example, for wind speed changes, this is the region for which the measured data are valid.

Finding a particular item on the map is done by using a R*-Tree structure [6]. For each item, its coordinates are reprojected to WGS84 [7], if necessary. WGS84 is the shorthand symbol for the Spherical Mercator projection, used, e.g., by the Global Positioning System (GPS). It uses latitude/longitude values instead of metric units and thus forms a common basis for all coordinates stored in the database.

After reprojecting, a bounding box for the new coordinates is computed. The bounding box is the rectangle used by the R*-Tree structure, into which it is then inserted.

Thus, the proposed simulation environment is able to integrate the artificial stimuli in the same manner as it integrates the actual agents being tested, which allows for greater flexibility and simplicity in the overall design of the underlying software architecture.

Prior to running, the simulation environment is configured via an object of the class `Description`. This class' attributes contain the information required in order to set up the simulation itself, which are the simulation's start and possibly stop time—the latter only if the simulation does not run endlessly—, a polygon designating the area within which the simulation takes place (the "bounding box"), and the definition of a factory class responsible for populating the environment. The latter one is accompanied by an enum set that specifies, for the three categories consumers, power generators and power grid, whether items from the respective category should be

 ISBN: 978-1-61208-332-2

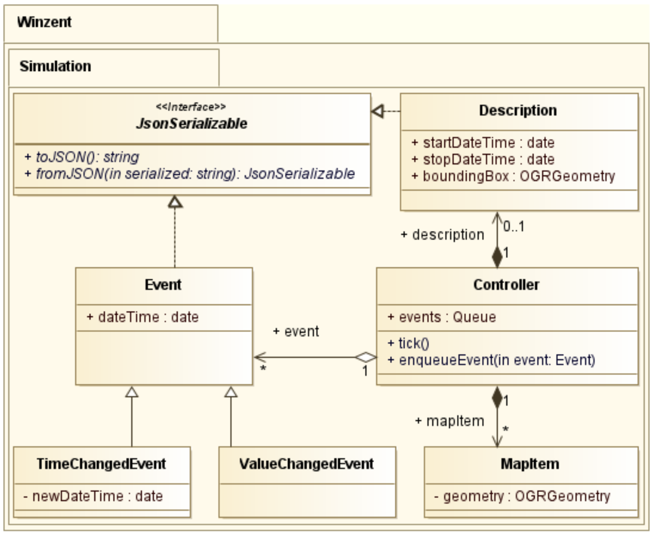

Figure 1. The `Simulation` namespace

created or not.

The description can be serialized to and from JSON [8], allowing an user to configure a simulation run by just creating a description in the text editor and feeding it to the simulator upon start. Thus, "simulation recipes" can be created which allow for repeated simulation runs in a defined manner. This is important for asserting the correct working of a new agent implementation and measuring possible improvements of such a new software version.

This basic layout of the relevant classes is depicted in the UML diagram of Figure 1.

The simulation controller offers another interface based on Qt's signal-slot principle [9], which can be used for other programs to link to the controller and observe the events within the simulation. All classes within the `Winzent::Simulation` namespace form one library; any frontend can link to it and display the simulation and its results in the way it chooses. Thus, simulation logic and view are separated which also enables the simulation to run headless. A headless simulation service can be used for running a clustered simulation as we outline in the next Section.

IV. Distributed Simulation

The simulation environment is multi-threaded; it keeps the actual event processing in separated worker threads to speed up the duration of one tick. When running standalone, the number of worker threads spawned equals the number of CPU cores. Each worker thread has an incoming event queue filled by the controller, as well as an outgoing event queue, wherein processed events are stored.

On local execution, these worker threads also store the map items themselves, i.e., event processing takes place on the local machine within these worker threads. For remote execution, the class `TcpSimulationThread` serializes and deserializes events and transmits them to remote slave workers instead of feeding a calculation.

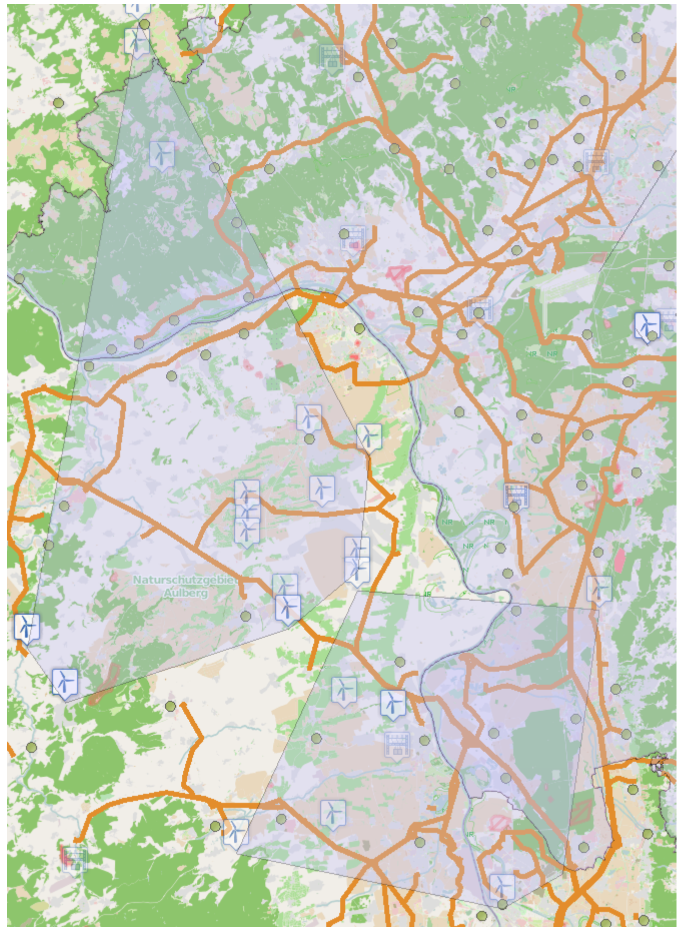

Figure 2. The simulator showing a part of Hesse, Germany, with areas of effect highlighted

These remote slave workers use the same classes as the master machine, i.e., the ones described in Section III. Such a slave machine uses one instance of the `TcpSimulationThread` to receive new events, which are enqueued into the slave `Controller`'s own event queue. The events then get processed on this slave machine via worker threads; processed events are serialized and transmitted back to the master controller.

This way, the simulation test bed can easily span multiple machines although the same classes are used. Since there is per definition one master controller that issues a new tick only once all other events for a simulation time has returned, all machines are in sync at the very beginning of a new tick.

Configuration is done via the `Description` class which also supplies all other parameters as per Section III. The simulation description additionally allows to specify a fixed mapping of a node or a set of nodes to a particular machine. This allows to break out specified nodes and run them on an embedded board, thereby additionally simulating hardware and memory constraints of a real smart grid node.

 ISBN: 978-1-61208-332-2

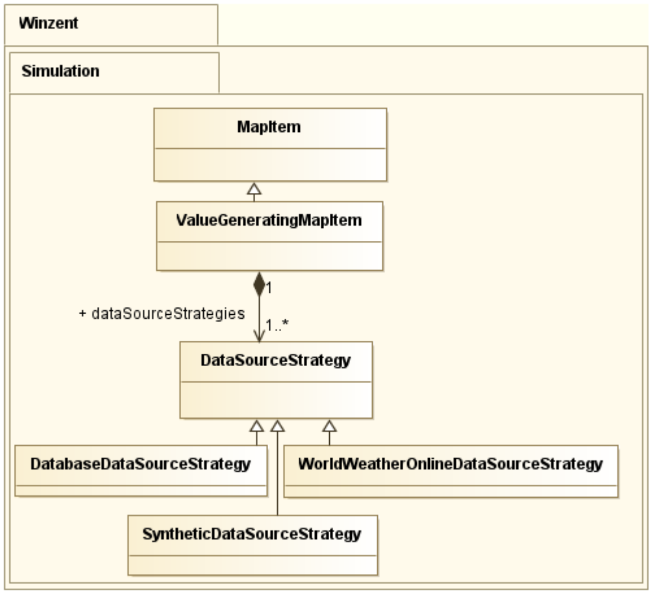

Figure 3. UML class diagram of the data source strategy pattern

V. Use of Open Data

The Winzent simulator does not rely on a single source of data for one particular input. Given, for example, weather data, it can query a national weather service's database or an online data provider such as OpenWeatherMap [10] or WorldWeatherOnline [11]. It is also possible to switch to completely synthetic input value generation.

However, not all sources have the same accuracy. Continuing with the weather data example, a dataset bought from the national weather service is likely to have a higher number of measurements available than an online service, which might serve data at three-hour intervals. As such, the user is likely to prefer the national weather service's dataset to other sources.

The simulation environment allows to express this by simple ordering through priorities. Only if a data source cannot offer any more data, the simulation environment will switch to one with a lower priority. These data sources are an implementation of the strategy pattern [12] used in object-oriented design.

Every time a `ValueGeneratingMapItem` receives a `TimeChangedEvent`, it queries the attached data source with the highest priority for data. If it cannot offer any, the `DataSourceStrategy` instance throws a `DataSourceDepletedException`. This prompts the `ValueGeneratingMapItem` instance to query the next one with a lower priority. This is done until the last data source has thrown an exception, which is then escalated to the `Controller`, which finally stops the simulation: A simulation which is not completely covered by valid data has no use in going on. Figure 3 shows the class architecture accomplishing this behavior.

While this simple mechanism allows the simulation to proceed even if one data source is depleted, the potential change in data quality obviously has a huge influence on the final results of the simulation run. An assessment of the data quality has to be done; at the least, the user has to be notified of the change.

Therefore, each data source has a derivation from the optimal data source attached to it. The optimal data source is, per definition, the first data source in the queue of possible data sources. Each other data source therefore is compared against the optimal source as long as this first source has data available. We discuss the metrics offered by this simulation environment in Section VI.

Another public source for data is the OpenStreetMap project (OSM) [13], which offers geospatial data free according to their licence [14]. OSM works like a wiki for geospatial data, i.e., everybody can add GPS traces to the database, tag ways and nodes on this map and supply additional data for existing items on the map.

This database can be imported into one's own PostgreSQL/PostGIS server [15]. Together with a Mapnik/mod_tile stack, this database can be used to render map tiles, such as the one which forms the background in Figure 2.

We also extract the locations of elements participating in the electricity grid from OSM where they are not covered by more exact databases.

The OSM database does not only contain spatial data, but also additional information such as the voltages a specific part of the power transmission system carries. This information is made available via hstore-based attributes called tags, which can be queried and extracted as a standard PostgreSQL feature. These tags, in OSM terms *Map Features*, are documented in the project's wiki [16].

VI. Discussion

The proposed simulation environment offers two benefits.

First, it treats both tested agents as well as data sources as items on the map. The attachment of spatial data to map items and events alike allows us to model realistic scenarios using real data sets. Together with the discrete event design, a simulation run can easily scale to run on a cluster.

Second, the simulator can switch data source strategies on the fly. However, with a change in the data source, there is also a possible change in the result of the simulation run. The simulator therefore must record that a switch has occurred, at which point this happened, and what data source strategy was subject to the change. This allows the simulation run to continue even if the preferred data source is unavailable, but also gives the user an indication that the quality of the result data may also have changed.

However, simply stating the fact that such a change occurred does not yet help in assessing the impact of the change. Therefore, the simulation environment provides two additional metrics to aid the user. We discuss these using the example of the German national weather service, *Deutscher Wetterdienst* (DWD) and World Weather Online (WWO).

For the simulation, the number of issued events is of foremost importance. An event triggers actions in the agents which participate in the simulation, ergo without events, there

 ISBN: 978-1-61208-332-2

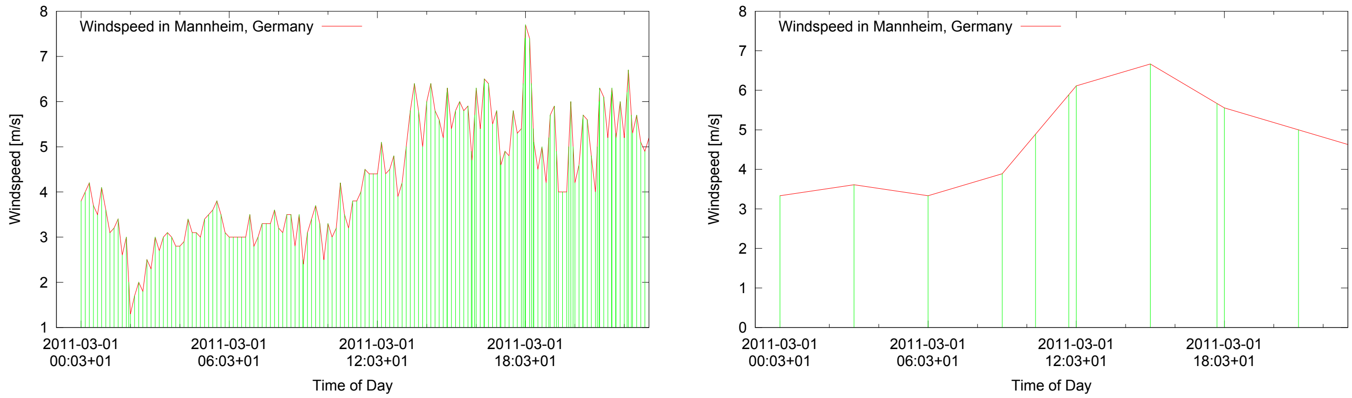

(a) DWD (b) WorldWeatherOnline.com

Figure 4. Wind speed at March 3, 2011 in Mannheim; impulses symbolize triggered events

will be no action. Thus, the number of events issued is the first available measurement for the quality of a data source. Figure 4 shows how different sources trigger a different number of effects: The vertical impulses show an event that would have been fired. One can observe that the open data source leads to a significant lower number of events in the system than the bought, high-resolution data from the DWD.

Using the second data source therefore reduces the expressiveness of the simulation run: The two data sources differ in their entropy. We obtain the per-source entropy using $H(X) = \sum_i P(x_i)I(x_i) = -\sum_i P(x_i)log_{10}P(x_i)$.

For our example day, the DWD source obtains $H(DWD) \approx 1.57\ Sh$ while the open data source has an entropy of $H(WWO) \approx 1.04\ Sh$.

In order to arrive at these numbers, records the number of events that are triggered by the primary data source as long as it is available, and also measures how many events would have been issued by all other, lower-priorized data sources along with the values both sources yield.

Obviously, not only the number of events, but also the values conveyed by particular events are important. Whenever a `ValueChangedEvent` is triggered, it simulates a sensor reading coming from a data source within the simulation test bed. When a change in the data source strategy has occurred, this possibly also changes the values for those types events. The simulator therefore also records all possible event values in the same manner as it records the number of events triggered and calculates a derivation to the primary data source strategy.

This derivation is the source for an error calculation that allows us to express how much the current data source differs from the ideal one. The possible derivations introduced by the error of a data source that is not the reference source allow us to judge where a simulation run may have failed even though the data itself suggests it succeeded.

Figure 5 compares two data sources for errors.

These values are finally summed up to produce an overall result of the data source strategy change. Such a comparison has been done in Table I. The values hold true for one particular day, which was also the basis for the previous Figures 4 and 5.

However, the data source strategies are not the only place where open data are being used by the simulation environment. The map itself and the position data of most items are read from a local copy of the OSM database. There is, however, no fallback as it is employed with the data source strategies.

The trustworthiness of this database is important regarding two requirements: First, because it provides positioning data for the power grid, power substations, producers and consumers, and even map tile images. Second, because the additional attributes attached to these nodes, like voltages carried by a part of the power transmission network or installed types of windmills within a wind farm, form another data set that is used throughout the simulation.

The issue of trusting a community-driven data set has been thoroughly discussed for Wikipedia, for example, in [17], and OSM has also been used to increase efficiency of emergency medical services (EMS) [18]. For our own usage, we've compared the OSM data regarding wind farms to those from TheWindPower.net, an on-line wind turbines and wind farms database [17].

It shows that the overall number of wind farms registered is lower in OSM than in The Windpower's data. However, OpenStreetMap always offers a position since this particular attribute must not be null, whereas The Windpower actully can contain data without a location attached. For the simulation, both cases are not useful: If the wind farm data set has a location, but no other useful information, the simulation

TABLE I. COMPARISON OF DATA SOURCES DWD AND WORLDWEATHERONLINE.COM

	DWD	WorldWeatherOnline.com
Events triggered	166	23
Events derivation	0%	13.9%
Derivation (Avg)	$0\,\frac{m}{s}$	$0.81\,\frac{m}{s}$
Date	2011-03-01	

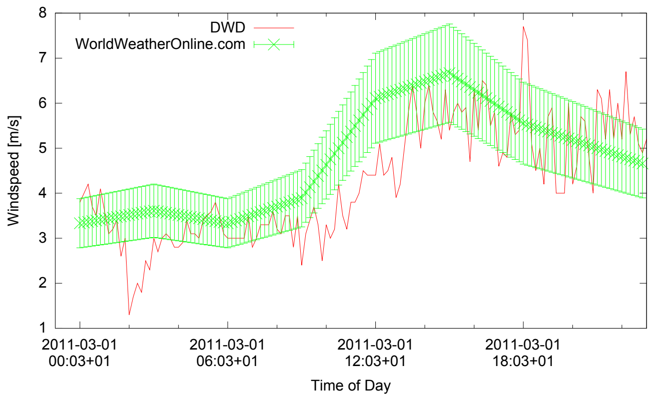

Figure 5. A reference data source (DWD) and an open data source (WWO) with error shown

environment cannot determine the electricity output of the installation; if this data is available, but the position is not, the simulation environment is unable to connect it to the transmission system at the right point.

Table II offers an overview over the most important figures as discussed.

As such, OSM is useful to render map tiles which allow the user to orient himself, but regarding data for simulation purposes, a specialized database should be preferred.

VII. Conclusion and Future Work

In this paper, we have introduced a discrete-event based simulation environment for smart grid messaging with a design adhering to the Don't Repeat Yourself (DRY) principle as much as possible. It is fully spatially indexed, which turns each participant into an item on a map and allows for quickly finding those items which are affected by a change of the environment.

The discrete events and the also discrete, locatable `MapItems` enable the user to run any simulation distributed, while the simulation description provides him with a tool to create reliably repeatable simulation runs.

The simulation environment makes heavy use of open data. It reads location data from OSM and attaches additional information to map items from the same data base. It is able to transparently switch between data sources, mitigating "holes" in one data set. While this allows the simulation to go on, it is also recorded; the simulation environment also provides hints on the possible impact of this change in data sources.

In the future, we will refactor the network code to be based on a Message Passing Interface (MPI), which will allow users to deploy the simulation testbed on any cluster.

TABLE II. COMPARISON OF DATA VOLUME OF OSM AND THE WINDPOWER

Feature	OpenStreetMap	The Windpower
Wind farms registered	508	13649
Wind turbines	1537	20215
Usable datasets	344	10101

VIII. Acknowledgments

This paper has been created as part of a cooperative doctorate program between the TU Bergakademie Freiberg and Wilhelm Büchner Hochschule, Pfungstadt.

The source code of the simulation environment described in this paper will be made available to the public on Bitbucket (*http://www.bitbucket.org/eveith/winzent-simulation*).

References

[1] M. Z. Lu and C. L. P. Chen, "The Design of Multi-agent based Distributed Energy System," *2009 IEEE International Conference on Systems, Man and Cybernetics (Smc 2009), Vols 1-9*, pp. 2001–2006, 2009.

[2] A. Varga, M. Visual, R. Omnet, and S. S. Method, "The OMNeT++ discrete event simulation system," in *Proceedings of the European Simulation Multiconference ESM'2001*, vol. 42, 2001, pp. 319–324.

[3] P. Oliveira, T. Pinto, H. Morais, and Z. Vale, "MASGriP – A Multi-Agent Smart Grid Simulation Platform," 2012, pp. 1–8.

[4] E. M. Veith, B. Steinbach, and J. Windeln, "A Lightweight Messaging Protocol for Smart Grids," in *EMERGING 2013, The Fifth International Conference on Emerging Network Intelligence*. IARIA XPS Press, 2013, p. 6.

[5] D. S. Wilks and R. L. Wilby, "The weather generation game: a review of stochastic weather models," *Progress in Physical Geography*, vol. 23, no. 3, pp. 329–357, 1999.

[6] N. Beckmann, H.-P. Kriegel, R. Schneider, and B. Seeger, "The r*-tree: an efficient and robust access method for points and rectangles," in *International Conference on Management of Data*. ACM, 1990, pp. 322–331.

[7] B. Decker, "World geodetic system 1984," DTIC Document, Tech. Rep., 1986.

[8] D. Crockford, "RFC 4627 - The application/json Media Type for JavaScript Object Notation (JSON)," IETF RFC, IETF, Tech. Rep., July 2006.

[9] J. Blanchette and M. Summerfield, *C++ GUI Programming with Qt 4*. Prentice Hall, 2008.

[10] Extreme Electronics Ltd., "Free weather data api for developers," [retrieved: 2013-10-14]. [Online]. Available: http://www.openweathermap.org/API

[11] WorldWeatherOnline, "World weather online," [retrieved: October, 2013]. [Online]. Available: http://www.worldweatheronline.com

[12] E. Gamma, R. Helm, R. Johnson, and J. Vlissides, *Design patterns: elements of reusable object-oriented software*, 1995, vol. 206. [Online]. Available: http://www.cs.up.ac.za/cs/aboake/sws780/references/patternstoarchitecture/Gamma-DesignPatternsIntro.pdf

[13] OpenStreetMap Contributors, "Openstreetmap," [retrieved: October, 2013]. [Online]. Available: http://www.openstreetmap.org

[14] ——, "Openstreetmap copyright and license," [retrieved: October, 2013]. [Online]. Available: http://www.openstreetmap.org/copyright

[15] ——, "Planet.osm," [retrieved: October, 2013]. [Online]. Available: http://wiki.openstreetmap.org/wiki/Planet.osm

[16] ——, "Map features," [retrieved: October, 2013]. [Online]. Available: http://wiki.openstreetmap.org/wiki/Map_Features

[17] S. Javanmardi, Y. Ganjisaffar, C. Lopes, and P. Baldi, "User contribution and trust in wikipedia," in *Collaborative Computing: Networking, Applications and Worksharing, 2009. CollaborateCom 2009. 5th International Conference on*, 2009, pp. 1–6.

[18] M. Azizan, C. S. Lim, W. Hatta, and L. C. Gan, "Application of openstreetmap data in ambulance location problem," in *Computational Intelligence, Communication Systems and Networks (CICSyN), 2012 Fourth International Conference on*, 2012, pp. 321–325.

Multi-Protocol Transport Layer QoS: A Performance Analysis for The Smart Grid

James Wilcox, Dritan Kaleshi
Center for Communications Research
University of Bristol, UK
James.wilcox@bristol.ac.uk
Dritan.kaleshi@bristol.ac.uk

Mahesh Sooriyabandara
Telecommunication Research Laboratory
Toshiba Research Europe Limited, UK
Mahesh@toshiba-trel.com

***Abstract*— Application specific interaction models will be required to support efficient communications between distributed applications with disparate network requirements in the consumer side Smart Grid (SG) data network. While much work has been done to quantify SG communications requirements in general, there is little public information on how to support the individual applications. This paper will show that specific transport protocols are able to provide increased efficiency of network resource usage under specific network conditions and that by providing a real-time adaptive selection of these transport protocols it would be possible to achieve a distributed embedded system with heterogeneous actors that can react to both application-specified Quality of Service (QoS) requirements and varying network conditions.**

Keywords – Quality of Service, Adaptive Transport Layer, Network Emulation, Middleware, Smart Grid

I. INTRODUCTION

Traditionally networked applications are designed with a pre-selected transport layer protocol. Optimisations for a specific application are done at the application layer and all messages are transported using the same protocol, either TCP (Transmission Control Protocol) [1], UDP (User Datagram Protocol), or with an overlay transport protocol such as RTP (Real Time Protocol) [2]; Fig. 1 represents this paradigm. This is typically fixed at application development; however there is no fundamental requirement for this to be the general rule. Whilst networked applications need to exchange information, there is no reason why application layer code should be concerned with how that information is transported. There are a multitude of existing, mature transport layer protocols available each designed to tackle specific network problems [3]. Utilising these many protocols, a single application could leverage the advantages of each protocol individually at the appropriate time given an environment with dynamic network conditions and application requirements. Acknowledging these points raises the challenge of defining a generic framework that allows for run-time selection of transport protocols to dynamically match specific application requirements and, specifically, message patterns used by the application. If the low level network interactions enforced by a specific transport protocol and higher level architectural messaging pattern are completely decoupled from the application then dynamically modifying the combination can be used as standard. If certain transport protocols and messaging pattern combinations are able to provide higher performance in terms of bandwidth, latency and reliability in certain network environments than others can do, then by supporting adaptive selection of these combinations it becomes possible to have a distributed real-time embedded (DRE) system with heterogeneous actors that can react to both dynamic application QoS requirements and network conditions. This model is shown in Fig. 2. The middleware system required for managing the selection of the large numbers of transport protocols referenced in Fig. 2 is referenced for completeness but is outside the scope of this paper.

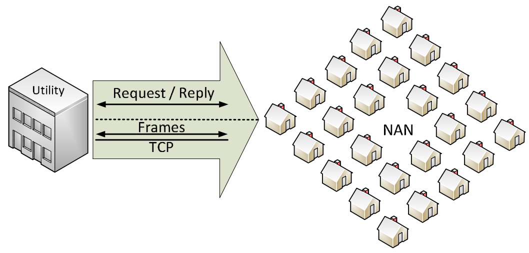

Figure 1. Traditional: Applications are supported by a single transport (in this case, TCP Request / Reply). The utility represents systems providing the SG infrastructure.

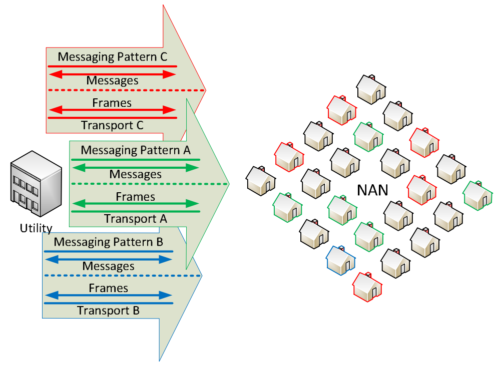

Figure 2. Proposed: Applications using multiple run time optimised transports (managed using middleware) instead of only TCP Request / Reply.

This paper shows experimental results which demonstrate that specific combinations of transport protocol and messaging patterns will provide performance gains over a pre-defined communication transport architecture and that certain combinations provide useful gains over other potentially viable options. The potential of generating a large scale mapping of transport combinations and application requirements will be explored.

The main contributions of this paper are:

- An analysis of performance of a distributed system for different transport protocol and messaging pattern combinations by running application scenarios using a fixed, realistic, network topology.
- Results that show that the performance of the Distributed Real-time Embedded (DRE) applications varies significantly for different messaging pattern and

transport protocol combinations, suggesting that this then can be used to optimise the performance of the individual transactions that make up the application network traffic.

The paper is organized into the following sections: section 2 provides the related material and further motivation for this work. Section 3 presents the experimental emulation test bed setup and the viable communication interaction models. Section 4 presents the experimental parameters and the results. Section 5 presents the conclusions from the experiments and the direction of future work.

II. Related work and motivation

Environments targeted by this work have the following characteristics:

1. They are distributed and built from a large number of heterogeneous embedded devices, running a number of different applications.
2. They are typically loosely-coupled.
3. The majority of the actors are communication network-constrained rather than computing resource-constrained.
4. Each device is expected to run many different applications with varying network requirements.

One example of such a system is the SG, and in particular, the subset of applications that intend to use consumer / demand side equipment and systems to achieve grid specific goals such as load shedding or load shifting or in more general terms, Demand Side Management (DSM).

Section A introduces the SG and its edge applications. Several related works exist which support with the premise of providing DRE software applications, such as those that will operate in a SG, with flexible communication choices in order to either achieve better network resource allocations or meet specific communication requirements. These overviews are presented in sections B and C.

A. The Smart Grid and Demand Side Management (DSM)

The SG can be seen as a large scale distributed system, with a large component being embedded sensor-actuator networks to support distribution power network monitoring and control and DSM interactions. DSM focuses around the control of demand side loads in the electricity distribution network in order to manipulate network conditions [4]. DSM breaks down into a number of related but still significantly different enough sub-applications to warrant different communications approaches. DSM can be broken down into two major sub categories, Demand Control (DC) [5] and Demand Response (DR) [6]. DC is defined as DSM programs that have centralized direct control over consumer loads [5]; DR is defined as DSM programs that use indirect methods (typically pricing) to affect changes [7]. Each approach requires a different communications paradigm in order to utilise network resources efficiently and operate optimally.

The above presents an ideal system for this work. It presents the rare opportunity to take a completely different approach to facilitating machine-to-machine communications in a DRE environment. The SG will eventually call for millions of networked geographically-distributed embedded devices to be deployed into the demand side of the power distribution grid. These devices are either designed to utilise existing networks such as domestic broadband or cellular networks and coexist with the existing traffic, or to utilise purpose built resource constrained networks such as various forms of wireless mesh or power line communications [8]. Both approaches result in strict network resource constraints for the applications. These constraints further increase the impact run-time transport level adaptive QoS will have in such environments. The main argument against dynamically matching transport protocols and messaging patterns to application requirements and real time network conditions has been one of complexity. With sensor networks, the SG and the Internet of Things (IoT) in general becoming more prevalent the environment is changing and these arguments are no longer valid. Our proposition is that the performance gain introduced by dynamically matching transport protocols and messaging patterns outweighs the required increase in complexity of the architecture and embedded hardware.

DSM can be shown to be a good example of how tailored communication paradigms could be beneficial in the SG and similar environments.

B. Transport Mechanisms

The concept of adaptive transport layer services for resource-constraint environments is well explored; however, the approach taken usually considers the transport protocol to be used already pre-selected at development stage, and consider adaptations above the transport layer. They do not propose to provide application optimisations at the lower transport level. However applications, regardless of the type of resource-constrained environment, can benefit from tailored communication service at the transport layer. Mutlu et al. [9] presents a middleware solution for performing transport level QoS focused on Bluetooth application profiles and uses CORBA (Common Object Request Broker Architecture) [10] to facilitate the middleware. While the scope is clearly limited, and transport protocol choices are not part of the QoS mechanism, the motivation is similar. Furthermore, it can be shown that different communication protocols have inherently different QoS characteristics and that using targeted protocols with specific applications can improve performance with a number of chosen metrics. Weishan et al. [11] recognise this and provide experimental results related to protocol switching overhead and also implement the system using a middleware solution. They conclude that protocol switching overhead is minimal with their chosen transport protocols and that protocol switching is beneficial to DRE environments.

C. QoS Architectures for DRE systems

The works highlighted here are attempting to improve or maintain DRE application performance in sub-optimal or resource-constrained networks by utilising real-time adaptive QoS management mechanisms. [12-14] focus on a single messaging pattern and attempt to provide adaptive QoS within these confines. It demonstrates that additional QoS optimisation opportunities are available if the scope of the system includes controlling lower level attributes such as transport protocols and messaging pattern combinations in conjunction with the adaptive QoS mechanisms. For example Wenjie et al. [13] propose a QoS adaptive framework for Publish-Subscribe Service called QoS Adaptive Publish-Subscribe (QAPS). They define several QoS policies and focus on fault tolerance and dependability of services. Schantz et al. [15] present a distributed, real-time embedded system capable of adaptive QoS. They describe in detail several methods of implementing end-to-end adaptive QoS mechanisms and explain how the work gives DRE applications more precise control over how their end-to-end resource allocations are managed. These proposed adaptive QoS mechanisms all address the same problem as this paper, but these implementations are limited to the application layer instead of considering a multi-protocol transport layer to access additional optimisation opportunities. Zieba et al. [14] develop the concept of quality-constrained routing in publish / subscribe messaging architectures. They develop a system which integrates application quality requirements into the message routing architecture in order to better support dealing with varying network conditions such as dynamic network topologies and link characteristics. The idea of integrating the dynamic application requirements into the communication paradigm provides a critical distinction from the others and further reinforces the need for verified optimised communication paradigms in order to meet these dynamic requirements.

III. Experimental Scenarios and Emulation Test Bed

A. *Experimental Scenarios*

The topology used is shown in Fig. 3. It is a simple fan out type topology where one node is distributing data to a group of 300 nodes representing consumer smart meters.

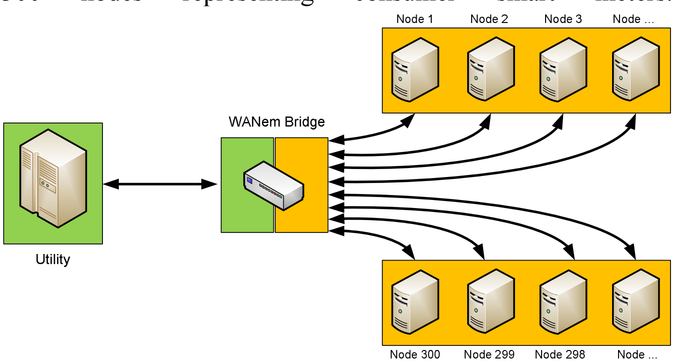

Figure 3 – Test bed network topology. 300 nodes are connected to a utility system through a software Ethernet bridge. The Utility publishes the update.

The topology represents the logical grouping that could be used in a Real Time Pricing DSM operation [16]. Traffic shaping is provided by WANem [17], which is a software wide-area network emulator. It provides the ability to manipulate many common network characteristics including available bandwidth, latency and packet loss.

B. *Selected Communication Paradigms*

Four viable transport protocol and messaging patterns were chosen to experiment with; these are shown in Table I and II. All are tested with both ideal and resource constrained, lossy network conditions in a set of eight experiments.

TABLE I - DOWNLINK (UTILITY→CONSUMERS) TRANSPORT & MESSAGING PATTERN CHOICES.

Scenario	Transport Protocol	Messaging Pattern
1	TCP	Router / Dealer
2	TCP	Publish / Subscribe
3	PGM [18]	Publish / Subscribe
4	UDP	Request / Response

TABLE II - UPLINK (CONSUMERS→UTILITY) TRANSPORT & MESSAGING PATTERN CHOICES

Scenario	Transport Protocol	Messaging Pattern
1	TCP	Request / Response
2	TCP	Request / Response
3	TCP	Request / Response
4	UDP	Request / Response

Router / Dealer is a tightly-coupled request-response style messaging pattern belonging to the ZeroMQ [19] socket API. It allows messages prefixed with a globally unique identifier (GUID) to be routed to a socket, remote or local, which has that same GUID. Each message sent needs to be prefixed with a valid GUID of a node which requires additional initialisation steps in order to acquire this information. Publish / Subscribe is a loosely coupled data distribution style messaging pattern. A publisher publishes a message prefixed with a topic / channel identifier. Only subscribers which have confirmed their interest in messages belonging to this topic / channel get the message routed to them. Scenarios 2 and 3 both use publish / subscribe but they use different transport protocols. Scenario 2 uses standard TCP. TCP is a unicast transport which implies that if a pricing update is to be sent to 300 nodes the Utility node will have to generate and send 300 individually addressed packets (assuming no fragmentation). Conversely, Pragmatic General Multicast [18] (PGM) is an experimental IETF (Internet Engineering Task Force) transport protocol designed to provide reliable multicast communications. In this case, the utility generates only a single packet (again assuming no fragmentation). Whereas TCP and PGM are both reliable transport protocols, UDP is unreliable. It does not have any mechanisms for ensuring reliable delivery but this does mean that it exhibits a lower network overhead.

C. *Link configurations of selected network scenarios*

Table III shows the network condition scenarios used in conjunction with the scenarios shown in Table I and II.

TABLE III - THE PHYSICAL NETWORK RESTRICTIONS.

Network Condition	Description
Ideal	No restrictions on bandwidth (10Gbps nominal – effectively unlimited) or any additional latency or packet loss
Resource Constrained	Bandwidth limited to 250Kb/s, additional 30ms +/- 5ms latency and 30% packet loss.

The resource constrained experiment emulates specific network conditions and represents a hypothetical resource-constrained lossy network on the link from the consumers to the utility such as an IEEE 802.15.4 based solution. Even though this represents two opposite extreme scenarios the results would still support the conclusions made for other network conditions.

 ISBN: 978-1-61208-332-2

Further experimental details are:

- In all experiments, the application layer maximum transmission unit (MTU) was configured for each transport protocol to ensure the packet size on the wire did not exceed 127 bytes. This was done to emulate the larger transport overhead (due to fragmentation) that would be seen when using these transport protocols with data link layers that can only support small packet sizes.
- The virtualised Ethernet bridge interface cards were configured for half duplex communication in order to emulate a half-duplex radio link.
- The payload used was a 1699 byte Extensible Markup Language (XML) string which is compatible with the OpenADR EventState.xsd XML schema [20].
- A price update was issued every 0.15 seconds in the request / response architectures and every 45 (0.15*300 = 45) seconds for the publish / subscribe architectures. This approach produces comparable test results as the fundamental differences in how the data is distributed between request / response and publish / subscribe would otherwise make this difficult. All scenarios achieve the goal of generating the same total number of responses from the consumers.
- All experiments issued price updates for up to 90 seconds and generated 600 responses from the consumers. Tests were allowed to run until all inflight responses were obtained.

The frequency of the Real Time Pricing (RTP) update is higher than any real world application. However, as the number of packets being generated, and hence the congestion, vary linearly with the RTP update frequency, using this frequency simply allows results to be collected easier. The higher frequency has no effect on the conclusions that are made in these experiments.

D. Emulation Test Bed

In order to develop and evaluate the premise that controlling the transport protocol and messaging pattern combinations of an application is an effective approach to manipulating the QoS, a way of allowing the experimental network code to interact with large numbers nodes was required. Emulation was chosen over simulation. Emulation provides a middle ground between a real world trial and simulation. The cost and difficulty of deploying a real world scaled trial is avoided but the ability to produce accurate and detailed data is maintained. There are several advantages to using a fully emulated approach over a simulated or semi emulated one. Firstly the test system is not attempting to approximate another real word system using a simplified model as is the case with a simulated approach. The emulation test bed can be seen as a condensed version of a real life system with all the varying levels of complexity a real world system would have from the standard open source software running on each node down to the physical layer of the network. Network analysis through the use of emulation is not a new area [21-25]. There is a large amount of work that uses network emulation due to the benefits it provides over purely simulation based analysis. More specifically, even the idea of a SG data network has undergone emulation based analysis [24] in the SCORE project which is a SG version of the Common Open Research Emulator (CORE) [26]. However none of the emulators allow manipulation of the transport layer and none allow the customisation of the virtualised node hardware in order to emulate resource constrained devices at the same time as being able to run real application code. Therefore, given the nature of the work it was decided that an emulation based approach would give the necessary flexibility needed and therefore a custom test bed was developed. To implement this test-bed a custom ESXi (VMWare Inc.) bare-metal hypervisor deployment was used.

IV. EXPERIMENTAL RESULTS

The results in this section use the UDP scenarios as a baseline. The raw results are shown in Table IV. This is done due to the UDP scenarios representing the simplest combination being experimented with. By using this scenario as benchmark it is easier to see how the other combinations perform in the given network topology against a well understood, ubiquitous transport protocol.

TABLE IV – THE RAW UDP RESULTS THAT CAN BE USED FOR COMPARISON WHEN RESULTS ARE SHOWN AS PERCENTAGE INCREASES.

	Ideal	**Constrained**
Utility Data	1.424 MBytes	1.424 MBytes
Consumer Data	1.424 MBytes	1.202 MBytes
Overhead	586.080 KBytes	586.080 KBytes
Message Round Trip Delay	2.183 ms	46.814 ms
Message Loss	0 %	39.3 %

The message latency result in Table V show that the publish / subscribe messaging patterns (Exp. 5 and 6) introduce greater latency than the request / response messaging patterns (Exp. 3 and 7).

TABLE V – MESSAGE LATENCY ROUND TRIP DELAY (RTD) AND MESSAGE LOSS (PS: PUBLISH / SUBSCRIBE, RD: ROUTER / DEALER, RR: REQUEST / RESPONSE)

	Message Round Trip Delay (RTD) (ms)	**Message Loss (%)**
1. TCP PS Ideal	345.6	0.00
2. TCP PS Constrained	21500.0	0.33
3. TCP RD Ideal	5.4	0.00
4. TCP RD Constrained	604.3	0.00
5. PGM PS Ideal	363.7	0.00
6. PGM PS Constrained	17040.0	0.33
7. UDP RR Ideal	2.2	0.00
8. UDP RR Constrained	46.8	39.30

Under the lossy, congested network conditions it took on average 17.04 seconds to complete a round trip for the PGM experiment (Exp. 6) and 21.5 seconds for the TCP (Exp. 2). This is extremely high and is due to the way the consumers are responding; the PGM and TCP publish / subscribe consumers use the same TCP request response architecture to respond with. There is no rate limiting which is causing a large amount of congestion. PGM provides an interface to limit the multicast data rate which would be very useful in this case. It can also be seen that even under perfect network conditions, rate unlimited publish / subscribe architectures are not suitable for applications requiring low latency as individual message delays are over 150 times that of the UDP case (Exp. 1 and 5 vs. 7).

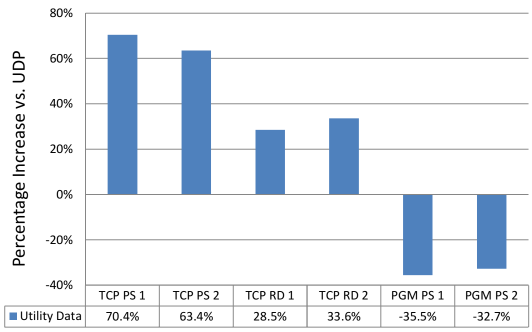

Figure 4 - Data sent by the Utility system compared to the UDP scenario.

The results indicate that the publish / subscribe architectures need a mechanism for rate limiting publishes and responses. The congestion generated when a published pricing update is sent and then responded to by all of the consumers simultaneously quickly overwhelms the resource constrained network generating message losses, which in turn cause retransmissions which contribute to the large amount of data generated. This can be seen in Figure 4 and Figure 5. The TCP publish / subscribe scenario shows this better than the PGM scenario. In this scenario the resource constrained, lossy network test actually performs better than the ideal case as the artificially imposed packet delay is having the effect of limiting the packet rate which even with the 30% packet loss and the retransmissions this would introduce, causes the scenario to generate less traffic than the 'ideal' case. This also indicates that a component in the virtual network is being stressed to the point of packet loss under the high packet rates being generated by the low MTU. Even with acknowledging this it shows that high messages rates with relatively large payload compared to the MTU will cause worst congestion problems than 30% packet loss does.

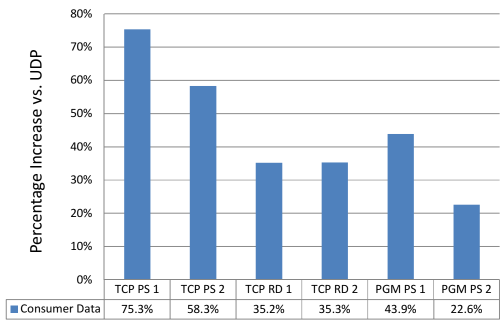

Figure 5 - Data sent by the consumer nodes compared to the UDP scenario.

The latency (Table IV, Exp. 7 and 8) and overhead (Figure 6) results show the UDP experiments outperform both the TCP or PGM equivalents with the TCP. This is not unexpected given the TCP and PGM are both reliable transports, with retransmissions that introduce increased delays against UDP. The notable observation is the performance gap between them. TCP is a generic transport capable of serving many different application requirements quite adequately, but the overhead involved in being so generic is clearly shown in these experiments. There is a clear opportunity to bridge this large gap with a number of UDP based messaging patterns, both unicast and multicast, and apply various application layer reliability mechanisms to them. This would allow applications access to a range of communications service combinations at a higher granularity, so that applications can get a communication service with only the features they need and avoid the general overhead of a one-transport-fits-all approach.

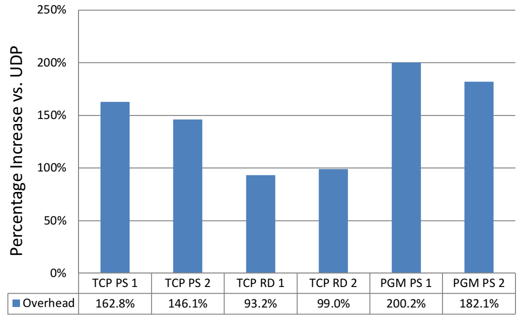

Figure 6 – Protocol overhead compared to the UDP scenario measured as any data that was not the XML payload. (Percentage increase is shown).

Figure 6 shows a large PGM percentage overhead. Given the much lower overall bandwidth consumed using multicast, this is to be expected. Overhead is calculated as any bytes put onto the wire that are not part of the XML payload. In order to generate the 600 responses (the experimental scenario criteria) from the consumers the utility only has to generate 2 PGM packets (ignoring fragmentation due to the low MTU). The overhead is almost entirely due to the TCP request response uplink from the consumer to the utility. Normalising these results against the total data exchanged shows the PGM architectures are in fact the most efficient next to the UDP architectures.

The results show that even though TCP publish / subscribe would appear to be a potential choice for this type of scenario (data distribution with a large fan-out) given that on face value it appears to provide the necessary interface for providing efficient data distribution, it actually performed the worst. TCP-based publish / subscribe involves a high amount of overhead to effectively allow a unicast architecture to emulate services that require a multicast architecture in order to operate efficiently. It provides no network orientated benefit over TCP Router / Dealer. The only benefit it provides is the ability to distribute messages at a more abstract level due to the use of topics / channels. In fact the lack of control on the distribution rate of the messages means that TCP router / dealer is more flexible and consistently generates less overhead and congestion as can be seen in Figures 4-6.

V. CONCLUSIONS AND FUTURE WORK

This paper has presented and validated an argument for exploiting the performance gains achievable by specifically selecting application appropriate transport protocols dynamically at runtime based on specific application requirements. Given the varied network requirements demanded by SG applications and DRE applications in general this approach provides previously inaccessible optimisation opportunities. Furthermore, these gains are achievable without the need to perform costly modifications to any intermediate network infrastructure and would only require modifications to existing networked applications' network interfacing code. The cost of this modification could be mitigated by using a middleware system for managing the transport selection.

To summarise, the results have shown:

- For an ideal RTP update distribution use case PGM publish / subscribe and UDP request / response should be used on the down links and up links respectively for best performance and lowest resource utilisation.
- For the non-ideal case, the unreliable UDP is only viable if the application can suffer lost responses from the consumer – this is a possible scenario. If more reliability is required then another low overhead reliable transport should be used with TCP based options used as a last resort.
- There is a significant gap between the performance of the Reliable TCP / PGM scenarios and the unreliable UDP scenario in terms of overhead and latency. Additional transports are needed to fill the gap.
- TCP based Publish / Subscribe provides no network level benefits.
- Rate unlimited Publish / Subscribe is not viable for applications with a low latency requirement. The packet rate needs to be limited at the point of transmission in order to ensure congestion is not generated.

A large number of additional supported transport protocols, would make it possible for a system to generate custom network interfaces for a much wider range of scenarios in order to improve application performance through manipulation at the transport level. Future work will consider how to automatically manage the large number of potential transport protocol choices which are being suggested using middleware solutions.

VI. References

[1] C. Fraleigh, S. Moon, B. Lyles, C. Cotton, M. Khan, D. Moll*, et al.*, "Packet-level traffic measurements from the Sprint IP backbone," *Network, IEEE,* vol. 17, pp. 6-16, 2003.

[2] IETF. (18/02/14). *RTP: A Transport Protocol for Real-Time Applications*. Available: http://www.ietf.org/rfc/rfc3550.txt

[3] IANA. (18/02/14). *Assigned Internet Protocol Numbers*. Available: http://www.iana.org/assignments/protocol-numbers/protocol-numbers.xhtml

[4] A. Mohsenian-Rad, V. W. S. Wong, J. Jatskevich, R. Schober, and A. Leon-Garcia, "Autonomous Demand-Side Management Based on Game-Theoretic Energy Consumption Scheduling for the Future Smart Grid," *IEEE Transactions on Smart Grid,* vol. 1, pp. 320-331, 2010.

[5] A. Gabaldon, A. Molina, C. Roldan, J. A. Fuentes, E. Gomez, I. J. Ramirez-Rosado*, et al.*, "Assessment and simulation of demand-side management potential in urban power distribution networks," in *Power Tech Conference Proceedings IEEE Bologna*, 2003, p. 5 pp. Vol.4.

[6] M. LeMay, R. Nelli, G. Gross, and C. A. Gunter, "An Integrated Architecture for Demand Response Communications and Control," in *Hawaii International Conference on System Sciences, Proceedings of the 41st Annual*, 2008, pp. 174-174.

[7] S. Chua-Liang and D. Kirschen, "Quantifying the Effect of Demand Response on Electricity Markets," *Power Systems, IEEE Transactions on,* vol. 24, pp. 1199-1207, 2009.

[8] Z. M. Fadlullah, M. M. Fouda, N. Kato, A. Takeuchi, N. Iwasaki, and Y. Nozaki, "Toward intelligent machine-to-machine communications in smart grid," *Communications Magazine, IEEE,* vol. 49, pp. 60-65, 2011.

[9] U. Mutlu, R. Edwards, and P. Coulton, "QoS Aware Bluetooth Middleware," in *Information and Communication Technologies, 2006. ICTTA '06. 2nd*, 2006, pp. 3239-3244.

[10] M. Henning, "The Rise and Fall of CORBA," *Queue,* vol. 4, pp. 28-34, 2006.

[11] Z. Weishan, K. M. Hansen, J. Fernandes, Schu, x, J. tte*, et al.*, "QoS-Aware Self-adaptation of Communication Protocols in a Pervasive Service Middleware," in *Green Computing and Communications (GreenCom), 2010 IEEE/ACM Int'l Conference on & Int'l Conference on Cyber, Physical and Social Computing (CPSCom)*, 2010, pp. 17-26.

[12] K. Young-Jin, L. Jaehwan, G. Atkinson, K. Hongseok, and M. Thottan, "SeDAX: A Scalable, Resilient, and Secure Platform for Smart Grid Communications," *Selected Areas in Communications, IEEE Journal on,* vol. 30, pp. 1119-1136, 2012.

[13] Z. Wenjie, Y. Nanhua, C. Hui, H. Jiajian, L. Chuanjian, and L. Xin Jie, "Providing adaptive QoS for Real-Time Publish-Subscribe Service in SGIOC-HQ," in *Innovative Smart Grid Technologies - Asia (ISGT Asia), 2012 IEEE*, 2012, pp. 1-5.

[14] B. Zieba, M. v. Sinderen, and M. Wegdam, "Quality-constrained routing in publish/subscribe systems," presented at the 3rd International Workshop on Middleware for Pervasive and Ad-hoc Computing, MPAC 2005, Grenoble, France, 2005.

[15] R. E. Schantz, J. P. Loyall, C. Rodrigues, D. C. Schmidt, Y. Krishnamurthy, and I. Pyarali, "Flexible and adaptive QoS control for distributed real-time and embedded middleware," presented at the Proceedings of the ACM/IFIP/USENIX 2003 International Conference on Middleware, Rio de Janeiro, Brazil, 2003.

[16] Z. Wei and A. Feliachi, "Residential load control through real-time pricing signals," in *Proceedings of the 35th Southeastern Symposium on System Theory.* , 2003, pp. 269-272.

[17] H. K. Kalitay and M. K. Nambiar, "Designing WANem : A Wide Area Network emulator tool," in *Communication Systems and Networks (COMSNETS), 2011 Third International Conference on*, 2011, pp. 1-4.

[18] IETF. (21/02/14). *RFC3208 - Pragmatic General Multicast Reliable Transport Protocol Specification*.

[19] ZeroMQ.org. (28/01/14). *Using DEALER and ROUTER Sockets*. Available: http://zeromq.org/tutorials:dealer-and-router

[20] OpenADR. (18/02/14). *OpenADR EventState.xsd XML schema*. Available: http://openadr.lbl.gov/src/EventState.xsd

[21] S. Doshi, U. Lee, R. Bagrodia, and D. McKeon, "Network Design and Implementation using Emulation-Based Analysis," in *Military Communications Conference, 2007. MILCOM 2007. IEEE*, 2007, pp. 1-8.

[22] K. Fall, "Network emulation in the VINT/NS simulator," in *Computers and Communications, 1999. Proceedings. IEEE International Symposium on*, 1999, pp. 244-250.

[23] S. Maier, A. Grau, H. Weinschrott, and K. Rothermel, "Scalable Network Emulation: A Comparison of Virtual Routing and Virtual Machines," in *Computers and Communications, 2007. ISCC 2007. 12th IEEE Symposium on*, 2007, pp. 395-402.

[24] T. Song, S. Wen-Zhan, D. Qifen, and T. Lang, "SCORE: Smart-Grid common open research emulator," in *Smart Grid Communications (SmartGridComm), 2012 IEEE Third International Conference on*, 2012, pp. 282-287.

[25] E. Weingärtner, F. Schmidt, H. V. Lehn, T. Heer, and K. Wehrle, "SliceTime: a platform for scalable and accurate network emulation," presented at the Proceedings of the 8th USENIX conference on Networked systems design and implementation, 2011, pp. 19-19.

[26] J. Ahrenholz, C. Danilov, T. R. Henderson, and J. H. Kim, "CORE: A real-time network emulator," in *Military Communications Conference, 2008. MILCOM 2008. IEEE*, 2008, pp. 1-7.

EnAware: A Comprehensive and Scalable Energy Management and Awareness Solution

António Santos, Carlos Resende, Rui Marques, André Costa Lima
Fraunhofer Portugal AICOS
Email: {antonio.santos, carlos.resende, rui.marques, andre.lima}@fraunhofer.pt

***Abstract*—The raising costs of energy and the increasing consumer awareness with regards to their ecological footprint poses new challenges to the industry. In a domestic context, consumers demand energy efficient products and solutions that can be seamlessly integrated and easy to use. The *EnAware* project, developed by Fraunhofer Portugal AICOS, EFACEC, EFAPEL and Bosch Termotecnologia, addresses these challenges by developing an ecosystem of interconnected smart devices. This ecosystem is composed of a smart energy meter, several smart devices and a home server. The household power meter and the home server communicate via the ModBus protocol over Power Line communication (PLC). Examples of the smart devices are smart sockets, smart switches or heat pumps. The devices communicate with the home ZigBee wireless network using the Home automation profile, which guarantees the smooth integration of third party products into the environment. The core of the *EnAware* system is the home server, which orchestrates the different devices and provides the user interface. The main differentiating factors of this solution are a set of features focusing on raising awareness on energy consumption and the scalability provided by having the service running on an OSGi platform.**

Keywords-Energy Awareness; Energy Monitoring; Smart Grids; Domotics; Smart Homes.

I. Introduction

Nowadays, it is a common understanding that energy is a precious resource and its cost has been dramatically increasing. People use energy that comes from various sources and over time there have been countless social, political and mainly environmental consequences related to the use of energy. One of the main reasons for the recent environment degradation is the unconscious and unsustainable use of energy, which leads to dangerous carbon emissions to the atmosphere. The snow ball effect continues since it is proven that the climate changes are mainly due to these emissions. Unfortunately, humankind is on the end side of all this process, since we are the ones that suffer from the massive and devastating consequences of extreme climatic phenomena.

The recent increase of these climatic phenomena has been leading to several energy awareness campaigns. These campaigns aim not only at improving the Earth environment, but also at saving costs to the end user. However, these campaigns can only have a significant impact if the target audience can measure the effect of the change of behaviour and therefore be motivated to continue this change in order to optimize energy consumption even further. To effectively support these campaigns, energy management and monitoring solutions are effective tools, which should achieve the goal of providing the user with energy consumption reports that can easily provide a clear view of the energy consumption for a given period, either in a global perspective or by equipment or even aggregated by house divisions. This idea goes the same direction as the main standard management axiom stating that ”you cannot manage what you cannot measure”. Every energy management system must provide the means and results for an ongoing motivation and even education, so that the user may keep his behaviour over time. It must be stressed that usually it takes a long time to actually start enjoying the benefits of energy efficiency improvements, whereas the costs associated with them are immediate and most of the times large.

Having all these ideas and concepts in mind, Fraunhofer Portugal AICOS (Assistive Information and Communication Solutions) (FhP), in a joint collaboration with Bosch Termotecnologia S.A., EFAPEL (Empresa Fabril de Produtos Eléctricos, S.A.) and EFACEC, has enrolled in a project whose main goal is to develop a complete energy management system, with a strong energy awareness component. The system is composed of a central device (a home server in the form of a micro-PC), a smart meter developed by EFACEC, a heat pump provided by Bosch and Zigbee-enabled devices (smart plugs, switches, etc.) developed by EFAPEL. EFACEC role is to provide solutions for integrating all domestic side components with the power network, under the perspectives of Network Operation (SCADA/DMS), of Demand Side Management (DSM) and of supply of new services for the consumers. Its role will impact mainly at the level of Demand Response, specifically at Smart Metering level (smart meter, data concentrator and head-end). Bosch Termotecnologia S.A. role is to provide efficient, innovative and sustainable solutions for sanitary hot water, integrating demand response functionality. The role of EFAPEL is to provide domestic Intelligent Electronic Devices, as components for home automation, capable of local metering, improving demand response. Finally, the role of FhP is to provide human computer interaction solutions, namely for the home server, which will integrate all domestic data, interfacing the smart meter and all home components, improving demand response.

This paper describes the overall solution, starting with an introduction in Section I, followed by a detailed review of the state-of-the-art in Section II. The architecture is extensively described in Section III and some initial results are revealed in Section IV. Finally, there is room for conclusions and future work in Section V, being the last Section of the document (VI) just for acknowledgements.

II. State of the art

Energy Management and Home Automation systems are currently somehow easy to find in the market. However, the

 ISBN: 978-1-61208-332-2

currently available products usually do not pack these two features in the same bundle and miss some of the intelligence and interoperability that would allow them to reach optimal results. The interoperability issue is even more severe as it usually means using proprietary communication protocols that limit the user to the products from a restricted group of manufacturers; in what it takes to the Energy Management features, most of the products usually do not take too serious the Energy Awareness topic, missing strong and effective data analysis and mining subsystems that would allow providing suggestions, automated actions, etc., that raises the user's energy awareness levels and reinforces his motivation and engagement in changing his energy consumption behaviour.

The systems provided by Wattvision [1], PowerWatch [2], Tendril [3] and Blue Line Innovations [4] are examples of Energy Management systems, with no focus on Home Automation or Energy Awareness features. This set of systems are capable of processing the energy measurements from the household smart meter or measuring the energy consumption of appliances with metering capabilities, and showing them, through a user friendly interface, to the user. However, the interpretation of the data is completely up to the user. Efergy [5] is another example of an Energy Management system with no major Energy Awareness abilities, but with support for remotely controlled smart sockets.

On the other hand, the systems from Vivint [6], HoneyWell [7] and TaHomA [8] are Home Automation oriented systems that are not capable of measuring energy consumptions.

Savant [9] and Alarm.com [10] systems merge Home Automation and Energy Management features into a single product. However, they are only capable of processing energy measurements from appliances, therefore they have no support for household smart meters. Additionally, and except for the activity pattern feature from [10] that can alert the user to his energy consumptions, either when he is at home or away, there is also no major focus on Energy Awareness and no energy data mining intelligent components.

Regarding Energy Awareness systems, AlertMe [11] is capable of pre-processing the measured consumption data, but has no Home Automation capabilites, having control only over smart meters and heating systems.

With a different perspective, Cloogy [12] is an Energy Management, Home Automation and Energy Awareness system. However, it is not developed with interoperability in mind, lacking the support to open communication standards and is only able to interface with Cloogy devices. The same happens with DEHEMS [13], a Framework Project 7 (FP7) project that use "current cost" devices [14] [15] and focus an important part of its work on HCI and on system intelligence [16].

Finally, Energy@Home [17] is a Home Automation and Energy Management system supporting the ZigBee Home Automation standard, providing Energy Awareness features, i.e., it is a system that comprises all the aforementioned technologies. When compared to the *EnAware* system, it differs in the way the system intelligence is used. For example, while Energy@Home, along with the load shift dependent on demand response and floating tariffs, is more focused on the distance between the current consumption and the contracted power, or the current total energy cost and a predefined maximum value [18], the *EnAware* system is more focused on finding usage patterns and providing suggestions (for instance, suggesting scenes, detecting standby devices, etc.), as well as providing scheduling and rules support based, for instance, on actions on devices or on energy consumption levels.

EnAware, the energy management solution presented in this document, aims at raising energy consumption awareness and provides an innovative modular design that enables the easy inset of new features.

III. ARCHITECTURE

Energy management technologies are evolving quite fast and there are a lot of standards, protocols, platforms and frameworks that need to be considered in order to achieve some pre-defined goals. The mindset behind the design of this solution was set in order to come up with a modular prototype that could easily be adapted to most of the scenarios. The core services were thought to be flexible enough and the peripheral modules to be replaceable or easily adapted. At the end, the system should be easy to set up and use, without compromising the grand objective: reduce energy consumption.

A. Home Server

1) Technical Details: To minimize the impact of having extra equipment at home, the Home Server was designed to run on the smallest device that could be able to meet the requirements of the project. Therefore, the choice fell on CuBox [19]. CuBox is an ARM based micro computer with an approximate power consumption of 3 W, an ARMv7 CPU running at 800 MHz, and 1 GB of RAM. These processing and power consumptions characteristics, although quite limited, were found to be enough for the project. However, the CuBox lacks on the communication interfaces needed by the *EnAware* service, namely ZigBee for the communication with the HAN devices; PLC for the communication with the power meter; and Wi-Fi so it can connect wirelessly to the home gateway (recent versions of CuBox have better specifications, including integrated Wi-Fi). To overcome these issues, three external devices were considered:

- A ZigBee to USB adaptor developed by EFAPEL and Bosch, connected to one of the available USB ports of the CuBox.
- The Linksys PLWK400 [20] was used for the PLC communication with the power meter;
- And the Belkin F7D1102az [21] Wi-Fi to USB adaptor, connected to the other available USB port of the CuBox.

Regarding the operating system, the choice ended up on ArchLinux AMR, a lightweight, flexible and minimalist Linux based operating system for ARM architectures. The *EnAware* service was developed in Java, due to its multi-platform deployment and its association with the OSGi [22] standard, which was used as the development framework. Opting for OSGi as the development framework was due to all the benefits it offers - lightweight, modularity, easy multi-platform deployment and reusable components with a dynamic lifecycle. Several implementations of this standard exist, however the choice fell over Equinox, the reference implementation of the OSGi framework, which implements all of the mandatory and most of the optional features of the OSGi R4 specification. This OSGi specification describes a dynamic module system for Java. These modules are called bundles and are possible to be installed, uninstalled, updated, started and stopped without interrupting the rest of the service or the Java VM.

 ISBN: 978-1-61208-332-2

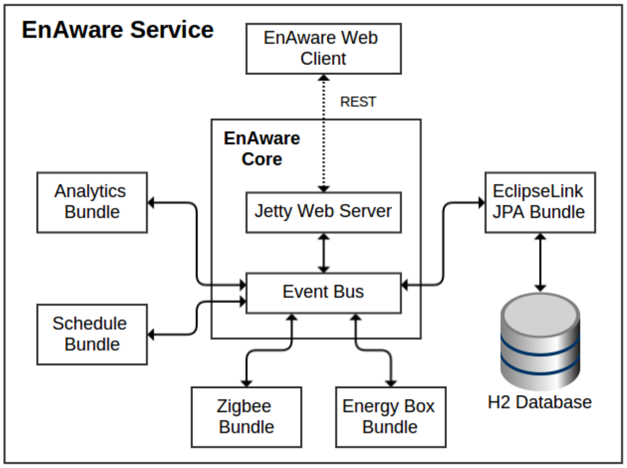

Figure 1: *EnAware* Service Architecture

Figure 2: *EnAware* User Interface

2) Client Application: The *EnAware* client application is a typical AJAX web application built upon a Fraunhofer AICOS internal Javascript framework based on the Dojo Javascript Toolkit, HTML5 and CSS. The *EnAware* web server (Jetty) has a Java Servlet for each data class and the client-server interaction follows the REST architecture (Fig. 1).

The design and implementation of the user interface aimed at meeting two main challenges. One was to provide simple interaction with all the available features and also a minimalistic but informative visualization of a potential large amount of data provided by the smart devices. The other was to have a responsive user interface, which adapts to several hardware platforms, and which might have diverse screen sizes and, at the same time, mouse or touch-based interaction. Web browser based applications are ideal for the latter challenge, as all the target platforms include a web browser and allows for multiple clients to be able to use the system without any per-client software requirements and installation. It is not a simple process to have an application that addresses a multitude of platforms, but not being restricted to a particular system is worth the challenge. When compared with more targeted applications, some compromises have to be made, for example, in interaction responsiveness. An effort was made toward minimizing these shortcomings and take advantage of what web-based applications can provide.

There is also some minimal care about security. The web portal is protected with a simple login mechanism.

The user interface has six main pages (Fig. 2). The ”Home”, or landing page, contains a chart displaying the current total consumption (based on the consumptions read from the household smart meter), as well as the latest system notifications and other shortcuts. The ”Consumption” page allows the user to have a detailed view about all the collected consumptions. Several time intervals can be visualized, from a single day to a whole year, and the consumptions can also be aggregated by house division or by device. The ”Devices” page allows the user to control the devices’ status. Most of the devices can only be turned on or off, but some, like window blinds and light dimmers, can be controlled with a slider. In the ”Scenes” page, the user can activate a scene or add new ones (not to be confused with Zigbee scenes, although they share the same purpose). Scenes act like shortcuts for device usage patterns. For instance, in the morning the user might frequently pull several windows’ electrical blinds up and turn on the radio and the corridor light. Running a Scene executes those actions in one step. Scenes can be configured manually but the system will also automatically send recommendations when it detects frequent patterns. In the ”Goals” page, the user is able to set a monthly consumption goal and check its status. Finally, ”Messages” is a very simple page where the user can check all the system’s notifications.

3) Core Service & Data Analytics: The *EnAware* service shares an event-driven architecture, where all the events are managed and distributed to all modules through an event bus. Most of the analysis performed have immediate results as the system receives data from the devices (stream mining), except cases in which the analysis requires data from an extended time period (for instance, Energy Efficiency Rating and Goals) and immediate results do not make sense. The analytic services were divided into three categories - Energy Measurement Aggregation, Behaviour Analysis and Energy Efficiency Rating.

The aggregation service was created mainly for performance reasons. Device measurements are aggregated into categories, to which they belong as they arrive to the system. Those categories are Timespan (day, week, month and year), House Division, Device and Total. This is an alternative to doing that aggregation when the data is requested. As a simple example, if the whole consumption from the kitchen is needed, that data is already pre-processed, so there is no wait for the whole measurements to be collected and processed.

The behaviour analysis category includes several distinct services - Device Consumption Anomaly (DCA), Goal and Scene. Within the DCA service, summary statistics are computed in a stream-like fashion to evaluate if a device is having an unusual consumption behaviour or some particular behaviour the user may be interested in detecting. One of such behaviours is if a device is standby-capable. The importance of detecting standby-capable devices comes from its non-negligible contribution to the overall consumption. According to IEA projections [23], by 2020, 10% of total appliance electricity consumption in the OECD could be for standby consumption, although there is a ongoing process to force the manufacturers to limit the standby consumption of their equipment to 1W (European Commission Regulation 1275/2008), being this limit lowered to 0.5W four years after the regulation inure. By detecting standby-capable devices and the most likely hours for them to be in standby mode, the system is able to suggest the user a schedule to turn them off. Other behaviours captured by this service are unusually high consumptions, devices which are on or off during an unusual schedule and devices which are on for an unusual amount of time. When such anomaly is detected, a notification is sent to the user and a relevant action upon the related devices is suggested.

The Goal service simply provides means to estimate if the energy consumption cost goal set by the user can be met, based on current and past energy consumption behaviour. The Scene service runs an implementation of the Frequent Pattern Growth algorithm, which allows the system to learn which devices are used within the same time interval. Based on this information, Scenes are suggested to the user, which can then be accepted and available for activation.

The Energy Efficiency Rating service estimates the energy efficiency class of devices based on European regulations. Knowing the efficiency class of the household appliances is another important input for energy awareness. A recent study [24] pointed out that changing to the market's more efficient appliances or the elimination/mitigation of standby consumption can reach up to 48% in savings. This efficiency estimation follows the European regulation as close as possible, but it was found that it is not feasible to strictly apply the energy efficiency calculation formulas due to their numerous parameters like, for example, dimensions or certain technical differences within the same type of appliance, which cannot be derived from the information available to this system. Even if the system would ask the user to fill in these parameters, some would require a detailed knowledge about the appliance's features, something that most users do not share. Energy efficiency is estimated by sampling the average consumption for some time period and extrapolating to the expected time period of each device class. Some device category's formulas might require the device's consumption over a year assuming a well defined usage pattern (e.g., Television, 4 hours per day during 365 days) and others for a specific operation (e.g., Washing Machine, standard washing program under full load). Nevertheless, a more detailed discussion on how this estimation is done is out of the scope of this document.

B. Zigbee Communication stack

The ZigBee technology is currently applied in several fields of the industry, e.g., health care, energy, retailing services or indoor automation to provide improved mechanisms and simplicity to ease the remote control and monitoring of activities or entities. The wide industry support is mainly due to the many standards developed by the members of the ZigBee Alliance to target custom applications. To achieve the goals of *EnAware* and meet its specifications, the Home Automation (HA) public application profile is used to communicate with all devices across the Home Area Network (HAN). The HA profile collects the required cluster set, which is a collection of commands and attributes that together define a controllable interface for a specific functionality, from the ZigBee Cluster Library (ZCL) for the *EnAware* application domains: energy metering, HVAC, switch and level control capabilities.

1) Devices: EFAPEL designed several smart devices, each with a ZigBee radio, in conformance with the ZigBee Standard and the HA profile to integrate the required functionalities. The single and dual-channel devices support one and two actuators or endpoints, respectively, in which independent home appliances can be wired to. The level control devices are capable of producing a variable output to control home appliances susceptible to level and limit control such as, e.g., window shades/blinds or dimmable light fixtures. EFAPEL has also designed a Zigbee-enabled panel with 4 programmable buttons, that is typically installed on the walls and can be used to define and trigger ZigBee Bindings (ZB), i.e. rules in the *EnAware* context, to control multiple devices at the same time. EFAPEL devices firmware also provide support for task scheduling, but do not support ZigBee Scenes (ZS). At the time of this writing, EFAPEL devices operate as ZigBee Routers (ZR), which means that they can communicate with each other by forwarding data. Further developments on device design will include devices backed by batteries that operate as ZigBee End Devices (ZED). With regards to the metering capabilities, EFAPEL devices, due to technical limitations, are not capable of obtaining consumption readings of independent channels on dual-channel devices; each reading is always, given an instant of time or time window, the combined consumption of both channels.

The Bosch heat pump version, developed within the *EnAware* project scope, is an efficient, innovative and sustainable solution for sanitary hot water production, which integrates seamlessly into smart energy technology by providing an integrated energy management system module. Common to the intelligent devices used in the project, the Bosch heat pump is equipped with a ZigBee radio, being capable of reporting energy consumption, reacting to demand response requests and adapt its energy storage conditions according to the energy cost. By using the HA profile, that gives support for the generic features through the mandatory standard clusters and the HVAC clusters - allowing, e.g., temperature set point, operation mode, water heating program - some additional proprietary clusters are used to support the heat pump smart energy integration enabling such features as: power load de-ballasting, load shedding, time shifting and dynamic load adjustment in function of the current tariff. Through these integrated functionalities, fine-grained control over the heat pump consumption, at any time of the day, is attained while keeping uncompromised comfort levels.

2) Home Area Network (HAN): The typical HAN in *EnAware*, depicted in Fig. 3, consists of several smart devices that talk only with the root network node, the ZigBee Coordinator (ZC), in case of ZEDs or with each other in case of ZRs for data relaying. The ZC is a USB dongle, developed by Bosch, using the ZigBee TI EMK [25], based on the TI CC2531 SoC, that is plugged to the home server and is flashed with a customized firmware developed by EFAPEL. EFAPEL additionaly designed a user-friendly case with a button for the USB dongle. The setup procedure of adding a new device to the HAN means unplugging the ZC from the home server, place it very close to the desired device and pressing the button to pair with it. The ZC will work at lower power mode when pairing with a device, otherwise it will pair with undesired devices that belong to networks in the neighbourhood. When the ZC is again connected to the home server, the list of the

HAN devices will be automatically updated.

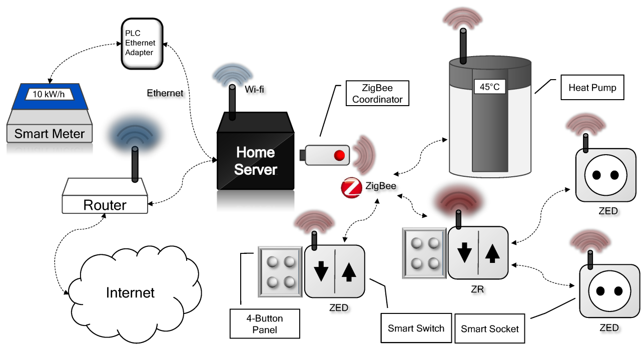

Figure 3: *EnAware* Home Area Network

3) Software layer: The ZigBee OSGi bundle is built in a modular fashion, consisting of four independent layers each implementing a bidirectional interface that enables layer crossing throughout the hierarchy. The modular architecture allows to easily extend the implementation to make the interoperability with other proprietary protocols possible, and to create new types of frames straight from factory methods. Data frames are exchanged with the ZC over a proprietary manufacturer protocol, developed by EFAPEL, on top of a serial-over-USB protocol. To accomplish this, the RXTX open source library API [26] was used as a broker to load and configure an operating system-specific native library implementation, that controls the access to the serial link.

The top layer of the hierarchy is the ZigBee service itself, that directly interfaces with the *EnAware* System core (ES). The messages, which are objects that encapsulate the content of a protocol frame, received from the ZC are processed here, its data extracted and passed on to the ES. Likewise, commands from the ES are assembled into manufacturer-dependent messages and fed to the lower layers. Layer two, the messaging layer, is responsible for managing both synchronous and asynchronous message exchanging. The protocol specifies commands that expect an acknowledge response for validation, and thus requiring synchronization, and others that do not. Moreover, consumption readings, for instance, are asynchronous messages because of their sporadic behaviour. Layer three, the data layer, extracts the payload from the message, builds a manufacturer-dependent frame and converts it to a byte buffer. Furthermore, it assembles messages by parsing raw bytes. Layer four, the serial layer, reads chunks of bytes from the serial connection data stream buffer and passes them on to the data layer.

C. Power Metering

One important part of the Home System architecture is its interface with the household smart power meter (called Energy Box), which was developed by EFACEC specifically for this project. The Energy Box is a smart meter developed having in mind the requirements of smart grids and related technology. Its communication interface is based on the ModBus protocol and uses the TCP [27] frame format.

The Home Server uses the Jamod library [28] (a Java library that supports all the variations of the ModBus protocol) to implement all the communication with the Energy Box, being able to retrieve the total energy consumption, the maximum contracted power and the smart grid demand response events, sending then all this data to upper software layers of the *EnAware* service for appropriate processing.

An important issue in the interface between the Home Server and the Energy Box is the communication medium. The Energy Box, as any other regular power meter, is usually installed outside of the house, possibly in a basement located several meters away from the Home Server. In these conditions, the use of wireless communication protocols is not appropriate, due to the interference and communication range issues that would arise. However, using wired communication may face other difficulties. Some buildings do not offer the needed infrastructures and therefore modifications would need to be done in order to install the link between these two components. To overcome these issues, *EnAware* uses the power line as the communication medium between the Home Server and the Energy Box. More precisely, from the available range of Power Line Communication (PLC) protocols, *EnAware* uses HomePlug AV [29], a PLC communication protocol developed having the home environment in mind.

D. Scalability

The discussion around the scalability of the system must be done in several dimensions. Due to the limited resources of the Home Server, the load scalability is therefore restricted and, although not yet identified but certainly high, there is a limit on the number of monitored devices and the related consumption report period. This can of course be resolved by using a more powerful server, thanks to the multi-platform characteristic of the *EnAware* service. In what it takes to functional scalability, the modular design of the *EnAware* service allows easy integration of new functionalities, which can be implemented as OSGi bundles, possibly being added or removed at any time. This means that not only the system can safely and quickly expand in features but also can easily widen the support for new devices.

IV. Performance and Results

Although the project roadmap includes some field trials, at the time of this writing they did not start yet so it is not possible to reveal some interesting results at this time, namely quantitative data regarding energy savings.

The system, mainly the Home Server, has been subject to intensive stress tests in order to assess its performance and stability. Due to its specifications, the Home Server is quite sensible in terms of CPU usage and memory consumption. The stress tests are composed of several energy consumption readings (30 devices reporting each 5 seconds) and GUI operations per second made by dummy devices. Thanks to a remote monitoring subsystem, we were able to follow to evolution of these indicators and, although they quickly reach the limits of CPU and memory capacity, they kept quite stable over time and the Home Server has been able to run smoothly without any interruption. The major overhead found is related to the persistence layer. The use of JPA with an H2 embedded database is found to be heavy for such a low-power server and if the number of events to be processed is unusually high, the system easily reaches its processing limit. Also, if the number of objects in the database is extremely large, the analytics subsystem also demands a lot of CPU and memory.

 ISBN: 978-1-61208-332-2

These issues, under extreme circumstances, may have impact on user experience while using the web portal since some lag and longer response times will be noticed. Since new versions of CuBox are currently available, which feature better specifications, it is expected that the *EnAware* service can support a higher throughput of events.

V. Conclusion and Future Work

As mentioned in this document, *EnAware*, apart from being an ordinary energy monitoring and management tool, has a special focus on energy awareness by implementing some features in a way that would allow the user to not only have a clear, detailed and intuitive view of the energy consumption, but also to be continuously engaged in optimizing it, reducing costs and helping our planet. Also, the benefits of using the OSGi model, like the support for pluggable data analysis engines, are characteristics that distinguish *EnAware* from the currently available energy management systems.

Some ideas and features were already identified as future work. They can easily be added to the system by taking advantage of using an OSGi model. Top priority goes to completing the support for smartgrids, namely the support for smart tariffs and demand-side management. There is the need to widen the array of supported Zigbee devices, also adding support for more Zigbee profiles and more home automation features. Also, the service is in need of a better user management system, enabling a multi-level access control system that could allow, for instance, having some users restricted to just being able to control a subset of devices or house divisions. The concept of rules is also targeted for expansion, adding support for rules having the energy consumption as trigger. Although the system provides a web portal that can be accessed by any web-enabled device (Desktop/Laptop or any mobile device), it is planned the integration of an LCD display that will be somehow connected to the home server and is also available to display a quick view of the overall energy consumption, as well as some alerts or messages.

VI. Acknowledgements

The Fraunhofer Portugal team would like to acknowledge all the support and intensive work from the all the partners in the consortium, as well as the financial support from COMPETE/QREN/EU.

References

[1] Aerodyno Inc. wattvision - real-time energy monitoring sensors and software for business and homes. http://www.wattvision.com/ [Retrieved: 12-2013].

[2] PowerWatch. Powerwatch - home energy management system hems. http://powerwatch.com/ [Retrieved: 12-2013].

[3] Tendril. Changing the way world uses energy - tendril. http://www.tendrilinc.com/ [Retrieved: 12-2013].

[4] Halifax. Power cost monitor - electricity usage. http://www.bluelineinnovations.com/ [Retrieved: 12-2013].

[5] Efergy Technologies Limited. Home of energy monitors, energy meters and electricity monitors. http://efergy.com/ [Retrieved: 12-2013].

[6] Vivint, Inc. Home automation, home security and solar energy systems. http://www.vivint.com/en/ [Retrieved: 12-2013].

[7] Honeywell International Inc. Honeywell - global technology leader in energy efficiency, clean energy generation, safety & security, and globalization. https://honeywell.com/Pages/Home.aspx [Retrieved: 12-2013].

[8] Somfy Systems, Inc. Tahoma home automation system - affordable home control by somfy. http://www.somfytahoma.com/ [Retrieved: 12-2013].

[9] Savant Systems LLC. Home automation control systems - commercial building automation and intelligent control. http://www.savantsystems.com/ [Retrieved: 12-2013].

[10] Alarm.com. About alarm.com - interactive security systems, home alarm monitoring. http://www.alarm.com/about/about_alarm.aspx [Retrieved: 12-2013].

[11] AlertMe. Alertme. https://www.alertme.com/ [Retrieved: 12-2013].

[12] ISA - Intelligent Sensing Anywhere, S.A.,. Cloogy. http://www.cloogy.com/en/ [Retrieved: 12-2013].

[13] DEHEMS Project. Dehems. http://www.dehems.eu/ [Retrieved: 12-2013].

[14] Current Cost. Current cost - reducing your energy bills so you can live a greener life. http://www.currentcost.com/ [Retrieved: 12-2013].

[15] J. Cooper. D5.3 meter and network plan. http://www.dehems.eu/cms/wp-content/uploads/2009/10/dehems_deliverable_d53.pdf [Retrieved: 12-2013].

[16] K.-M. Chao, N. Shah, A. Matei, T. Zlamaniec, W. Li, C.-C. Lo, and Y. Li, "Intelligent interactive system for collaborative green computing." in CSCWD, W. Shen, J.-P. A. Barthès, J. Luo, P. G. Kropf, M. Pouly, J. Yong, Y. Xue, and M. P. Ramos, Eds. IEEE, 2011, pp. 690–697. [Online]. Available: http://dblp.uni-trier.de/db/conf/cscwd/cscwd2011.html#ChaoSMZLLL11

[17] Energy@home. http://www.energy-home.it [Retrieved: 12-2013].

[18] Energy@home. Energy@home - use cases. www.energy-home.it/Documents/Technical%20Specifications/Energy@Home%20Use%20Cases_v_1_2.pdf [Retrieved: 12-2013].

[19] SolidRun. Cubox. http://www.solid-run.com/cubox [Retrieved: 12-2013].

[20] Linksys. Linksys plwk400 - powerline av wireless network extender kit. http://www.linksys.com/en-mea/products/powerline/PLWK400 [Retrieved: 12-2013].

[21] Belkin International,Inc. N150 micro wireless usb adapter. http://www.belkin.com/uk/support-product?pid=01t80000002K6MVAA0 [Retrieved: 12-2013].

[22] OSGi Alliance. Osgi alliance. http://www.osgi.org/Main/HomePage [Retrieved: 12-2013].

[23] International Energy Agency, Cool Appliances: Policy Strategies for Energy-Efficient Homes, ser. Energy Efficiency Policy Profiles Series. OECD, Organisation for Economic Co-operation and Development, 2003. [Online]. Available: http://books.google.pt/books?id=tSpPAAAAMAAJ

[24] ECEEE 2009 Summer Study : Act! Innovate! Deliver! Reducing energy demand sustainably. European Council for an Energy Efficient Economy (ECEEE), 2009.

[25] "TI CC2531 Evalutation Module Kit," Texas Instruments product. http://www.ti.com/tool/cc2531emk [Retrieved: 12-2013].

[26] T. Jarvi et al. RXTX Library. http://rxtx.qbang.org [Retrieved: 12-2013].

[27] Modbus-IDA.org. Modbus messaging on tcp/ip implementation guide v1.0b. http://www.modbus.org/docs/Modbus_Messaging_Implementation_Guide_V1_0b.pdf [Retrieved: 12-2013].

[28] jamod team. Java modbus library (jamod). http://jamod.sourceforge.net/ [Retrieved: 12-2013].

[29] HomePlug Powerline Alliance, Inc. Homeplug av white paper. http://www.homeplug.org/tech/whitepapers/HPAV-White-Paper_050818.pdf [Retrieved: 12-2013].

Fit for Solar Power - Computer-Assisted Planning of Regional Power Grids

Thomas Rist
Faculty of Computer Science
University of Applied Sciences Augsburg
Augsburg, Germany
e-mail: tr@rz.fh-augsburg.de

Michael Wiest
Faculty of Electrical Engineering
University of Applied Sciences Augsburg
Augsburg, Germany
e-mail: michael.wiest@hs-augsburg.de

***Abstract*—Additional power feed-in by newly installed photovoltaic panels may exceed the available capacity of an existing power grid and can seriously harm grid stability. Grid operators need to consider future scenarios with an increased share of solar power and examine whether the capacity of an existing grid will suffice or whether upgrades or extensions need to be made. In this contribution, we present IPDS, an Interactive Planning tool, which assists a grid planner in making predictions about such future scenarios. Starting from a cadastral map that shows the topological structure of an existing power grid, the tool estimates the yet unused roof surfaces of buildings, which are connected to a certain grid circuit. In turn, this estimation can be used to estimated future additional power feed-in.**

Keywords- Renewable energy - energy source integration models, network/grid plannin, integration of solar technologies.

I. Problem Statement

Due to favorable funding conditions, Germany recorded a rapid increase in photovoltaic (PV) solar power plants during the last years. According to a monthly publication of the German Federal Network Agency [1], the total installed capacity of all funded and registered PV installations reached a value of 34,186 MWp in June 2013. Following forecasts of the German energy industry [2], this development will continue even in case of expected declining public funding.

The increasing share of solar power generation contributes to the implementation of the politically encouraged energy turnaround from fossil fuels to renewable sources. On the other hand, this development challenges grid companies and stake-holders who need to decide on investments into a technically reliable and economically well-dimensioned grid infrastructure.

The capacity of an existing power grid is limited by the maximum load capacity of its components (such as cables and transformers). As a consequence, additional feed-in by newly installed PV panels may exceed available grid capacity and cause severe problems. In order to prepare for counter measures, grid operators need to consider future scenarios with an increased share of solar power and examine whether the capacity of an existing grid will still sufficiently meet the requirements of power quality [3], and if not, which components should be upgraded (e.g., larger cable sizing), or in which way a grid needs to be extended (e.g., additional lines and transformers).

Calculations of this kind are usually performed with the assistance of a dedicated power system analysis software, such as PowerFactory [4]. Typically, such programs rely on schematic maps to display local wiring together with other grid components, such as transformers, of a concrete power grid. However, calculations are performed on the provided input data. Since the result of such capacity calculations impacts decision making on whether or not to invest into a certain grid infrastructure, it is decisive that capacity calculation are performed on the basis of accurate assumptions about possible future scenarios. Unfortunately, it is difficult to predict in detail where additional PV sources may appear in the future as the installation of new PV panels is to a great extend in private hands (especially in rural areas by the owners of houses and land). In the sequel, we describe a tool which assists grid operators in making accurate predictions.

II. IPDS - An Interactive Tool for Planning Grid Extensions

Interviews with grid operators revealed that the amount of not yet used roof surfaces is a good starting point for making assumptions where new PV panels may appear in the future. Therefore, a grid planner may consult two sources of information: i) a cadastral map that shows the boundaries and ownership of land parcels and buildings together with the topological structure of an existing power grid, and ii) an aerial view (e.g., a street map or an aerial image) for finding out where are roof surfaces which are suitable for additional PV panel installations. Such working practices motivate the development of an interactive tool -called IPDS- that supports the following essential use cases:

A. *Use Case 1: Interactive examination of existing circuits*

As a basic data source we rely on cadastral maps that show the topological structure of an existing power grid. In order to gain some independence from proprietary data representation formats of an underlying drawing tool, or a certain geographic information system, we assume that such maps are available as ready-to-print PDF-documents. To this end, our IPDS tool comprises a PDF-scanner module which turns a static cadastral map into an interactive map view. More precisely, IPDS extracts from the PDF source document both cadastral information, such as buildings and streets, as well as grid components, such as power lines and transformers which may belong to different circuits.

Figure 1. An interacive cadastral map showing boundaries of land parcels and buildings, as well as the structure of an existing power grid.

After completion of the scan process the cadastral map with all grid circuits is displayed in the map-view panel of the IPDS tool (cf. Fig. 1). In addition, IPDS comprises an interactive tree-view showing a hierarchical view of all objects that have been extracted during the scan process. These are the individual power lines circuits (showed in different colors), and all buildings connected to those lines.

There is a cross-reference between the graphically displayed objects in the map-view and the named objects entailed in the tree-view. Since in IPDS's graphical user interface all objects are selectable (via mouse click), the grid planner can easily select or deselect a certain circuit in order to get an overview of the grid structure to be examined. To handle larger maps, the map-viewer also supports zooming-in/out as well as panning.

B. *Use Case 2: Estimation of PV energy gain*

Starting from the schematic representation of buildings as given by the cadastral map, an estimation of potential additional power feed-in through PV panel installations can be obtained by summing up the map area covered by buildings. However, in a power grid examination task the grid planner should be enabled to easily include and exclude buildings so that a variety of possible scenarios can be considered. To this end, IPDS supports a number of options for selecting objects (i.e., buildings and outhouses):

- *Selection of all buildings* which are connected to a certain grid circuit. This selection can be done using the tree-viewer. There, single circuits appear as colored top-level tree nodes. If selected, all sub-ordinate nodes (i.e., buildings) are selected, too.
- *Selection of a certain building*. In this case, the user simply clicks on the graphical representation of the building in the map view (cf. Fig. 2).
- *Selecting a group of adjacent buildings*. If interested in a certain neighborhood, IPDS allows the user to draw a rectangular frame (or likewise a closed polygonal line) around objects of interest shown in the map-view (cf. Fig. 2). All enclosed objects are considered as selected in proceeding calculations.
- *Exclusion of objects from a group*. In addition, it is possible to exclude a certain object from a group of selected objects. For instance, think of a church or a listed heritage building which should be excluded as an installation of PV panels is very unlikely.

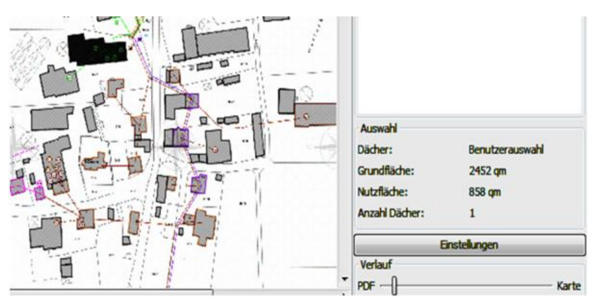

Figure 2. Selection of a single building (marked in black)

Figure 3. Selection of a group of buildings for estimating available roof surface area for additional PV panel installations.

PV-generated energy gain depends on a number of factors, including:

- Surface Area (in square meters) of installed PV panels.
- E_{PV_Type} - the specific solar module efficiency value of a certain PV panel type. This value is usually given in "Watt Peak" measured under standard test conditions (1000 W/m² global radiation, 25 °C module temperature, and air mass factor of 1.5). Sunlight conversion rates of current commercial PV panels typically range between 18% and 21%.
- $E_{Installation}$ - a factor to modify the E_{PV_Type} value based on the specific characteristics of a PV panel installation (i.e., latitude of location, cardinal direction, slope angle). In addition, an accurate value for solar radiation at a certain location can be computed by means of a solar position algorithm (SPA) [5].

For the power grid planner potential maximum values are of highest interest. As a consequence, the maximum available surface area for PV panel installations becomes the crucial variable to be estimated.

As soon as an object selection has been made, IPDS calculates the occupied ground area A_{object} (label "Grundfläche" in Fig. 2 and Fig. 3) of all selected buildings and estimates a potential usage area A_{usage} (label "Nutzfläche" in Fig. 2 and Fig. 3) for additional PV installations. Note that already installed PV-panels are known as they are part of the power grid and thus have reference points on the cadastral map.

We calculate A_{usage} as a certain percentage of A_{object}, i.e., the summed ground area given by the set of marked objects in the map-view. Assuming that in (German) rural villages gable roofs are dominant, and due to skylights, dormers and chimneys, parts of a roof area cannot be used for PV panel installation. Using the tool showed that a default value for A_{usage} may be chosen between 35% and 60%, which is in accordance with values reported in [6]. However, a grid planner is free to adjust this value (IPDS parameter "Nutzfaktor"), e.g., to account for actual roof types and cardinal direction of buildings.

To obtain an estimate for power feed-in generated by PV panels, the estimated usage area must be multiplied by a factor that accounts for the installed PV module efficiency. IPDS allows the grid planner to make specific adjustments of the above mentioned parameters E_{PV_Type} and $E_{Installation.}$ In most cases, for the sake of comparability, a grid planner may leave default settings unchanged when exploring different scenarios.

III. Improving IPDS

The current version of IPDS supports grid planners who examine grids in order to identify necessary grid extensions which are needed for an additional feed-in from anticipated future PV sources. The tool is actually used in a project supported by a regional power company in southern Germany that also maintains power grids of a number of rural villages.

However, feed-back from grid planners who worked with IPDS suggests several improvements and extensions to be made, some of which are briefly discussed in the sequel.

A. Support for data exchange between IPDS and power system analysis software

An obvious inconvenience in the current grid planning workflow is the fact that a grid planner needs to work with two software systems, a power system analysis tool on the one hand, and IPDS on the other hand. While IPDS can be used to estimate the potential future feed-in through additional PV panels for an existing grid circuit, a power system analysis tool is needed to calculate the impact of the additional feed-in on the existing grid circuit.

Vice versa, if changes to the grid structure are made in the power system analysis tool, one would need to incorporate these changes into the grid representation used by IPDS. While technically less challenging, the issue of data exchange between IPDS and commercial power system analysis tools can only be solved in cooperation with the developers of such tools.

B. Improving estimates for PV-usable roof area by means of an aerial image

As elaborated in Section 2, IPDS performs a rough estimate of available roof surface area for additional installations of PV panels. The roof surface area is estimated on the basis of the ground area covered by buildings as indicated in a cadastral map.

Thanks to available map services on the Internet, such as the Google Maps API [7], it is possible to use an aerial image of a region as background in the IPDS map-view panel and superimpose a cadastral map together with a topological power grid visualization. For the purpose of demonstration, we added a Google Maps import function to IPDS that allows to import a map section.

By means of a slider, the transparency value of the satellite image can be continuously adjusted between 100% transparency, i.e., only the cadastral map is visible (left-hand frame of Fig. 4), and 0%, which displays the background map fully saturated. As shown by the right-hand frame of Fig. 4, a value of 20% renders the background map partly visible so that roof shapes are well recognizable.

Such an overlay provides a grid planner with additional information about the actual roof shapes, and thus allows

 ISBN: 978-1-61208-332-2

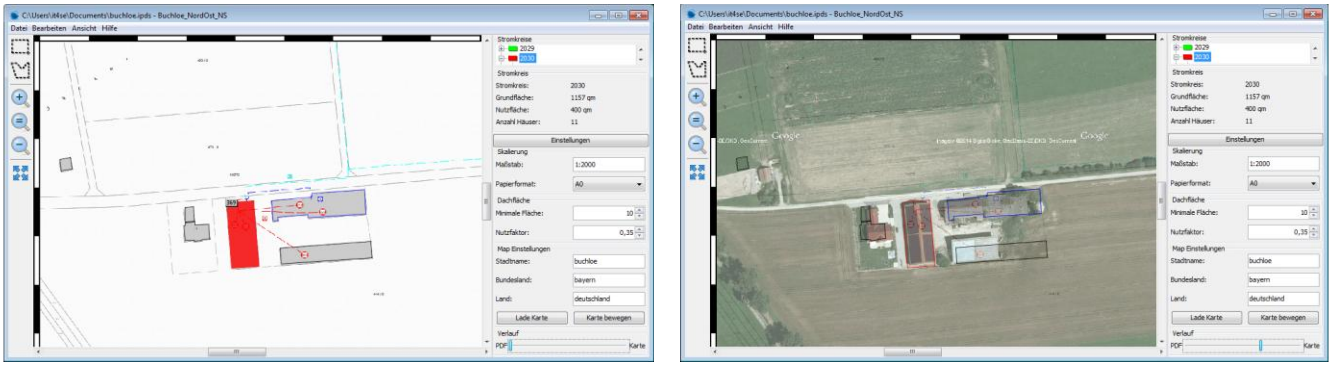

Figure 4. An interactive cadastral map with power lines (left) can be overlaid on a satellite image (retrieved via the Google Maps API).

her/him to set the parameter "Nutzfaktor" (cf. Use Case 2) more accurately.

C. *Improving estimates for PV-usable roof area by means of image analysis techniques*

In case of an aerial image is available, one can even go a step further and try to deploy image analysis techniques for an automated classification of roof shapes. To this end, we experimented with some simple edge detection filters, such as a Sobel discrete differentiation operator. Fig. 5 shows an aerial view of two gable roofs. After edge detection, roof princes appear as edges that run parallel to a pair of edges which belong to a roof's side boundaries.

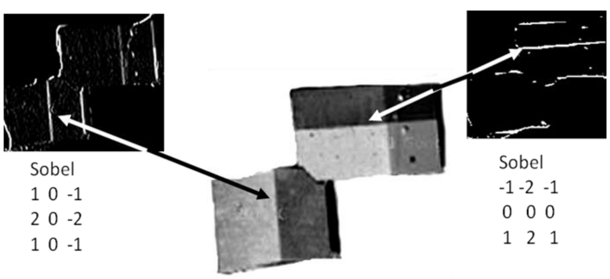

Figure 5. Automated detection of gable roofs by means of Sobel discrete differentiation operators for edge detection. After edge detection, roof princes stand out clearly in the image.

IV. CONCLUSIONS

In this contribution, we have presented IPDS, an interactive tool that assists a power grid planner in making accurate assumptions about potential additions of PV panel installations to an existing grid. In its current version, the tool supports several core functions, foremost the estimation of roof surface area of buildings, which are connected to a certain grid circuit and thus are candidates for additional PV panel installations.

While the tool is currently used and evaluated in the context of rural grid planning tasks, we also sketched some ideas for extension, such as the deployment of image analysis techniques for automated roof shape recognition.

A more accurate estimate of PV-usable roof area constitutes an upper bound for potential additional PV capacity. However, other factors, such as PV panel costs, public funding conditions, and community-specific building regulations are further parameters to be considered, too.

ACKNOWLEDGMENT

IPDS is has been developed in the context of the BMBF-funded German-New Zealand research collaboration IT4SE, grant number 01DR12041. Further information on IT4SE can be found under www.it4se.net

The development of IPDS has been motivated by scenarios as examined in the research project "Network Planning under Uncertainty", which is conducted by the Faculty of Electrical Engineering at the Univ. of Applied Sciences Augsburg in cooperation with the Technical University of Braunschweig.

The presented screenshots have been taken from an IPDS prototype, which was implemented by computer science students during the summer semester 2013. We would like to thank the group members for their dedication and perseverance.

REFERENCES

[1] Bundesneztagentur (German Federal Network Agency). [Online]. Available from: www.bundesnetzagentur.de 2014.02.14.

[2] M. Reichmuth, "Die EEG-Mittelfristprognose bis 2016". Energiewirtschaft-Recht-Technik und Umwelt, Section: Energiewirtschaftliche Tagesfragen. [Online]. Available from: www.et-energie-online.de/Zukunftsfragen/ 2014.02.14.

[3] VDE Association for Electrical, Electronic & Information Technologies, "VDE application rule: VDE-AR-N 4105". [Online]. Available from: www.vde.com/de/fnn/arbeitsgebiete/seiten/n4105.aspx 2014.02.14.

[4] PowerFactory is a product by DIgSILENT GmbH.

[5] I. Reda, A. Andreas, A. Solar Position Algorithm for Solar Radiation Applications. NREL Report No. TP-560-34302, 2003, revised version published January 2008. [Online]. Available from: www.nrel.gov/midc/spa/ 2014.02.14.

[6] G. Kerber "Aufnahmefähigkeit von Niederspannungsverteilnetzen für die Einspeisung aus Photovoltaikkleinanlagen" PhD Dissertation, Technical University Munich, 2011.

[7] Google Developers, "Google Maps API". [Online]. Available from: https://developers.google.com/maps 2014.02.14.

Evaluating Power-Performace Benefits of Data Compression in HPC Storage Servers

Konstantinos Chasapis, Manuel F. Dolz, Michael Kuhn, Thomas Ludwig
Department of Informatics
University of Hamburg
Hamburg, Germany
E-mails: {konstantinos.chasapis,manuel.dolz,michael.kuhn,thomas.ludwig}@informatik.uni-hamburg.de

***Abstract*—Both energy and storage are becoming key issues in high-performance (HPC) systems, especially when thinking about upcoming Exascale systems. The amount of energy consumption and storage capacity needed to solve future problems is growing in a marked curve that the HPC community must face in cost-/energy-efficient ways. In this paper we provide a power-performance evaluation of HPC storage servers that take over tasks other than simply storing the data to disk. We use the Lustre parallel distributed file system with its ZFS back-end, which natively supports compression, to show that data compression can help to alleviate capacity and energy problems. In the first step of our analysis we study different compression algorithms with regards to their CPU and power overhead with a real dataset. Then, we use a modified version of the IOR benchmark to verify our claims for the HPC environment. The results demonstrate that the energy consumption can be reduced by up to 30 % in the write phase of the application and 7 % for write-intensive applications. At the same time, the required storage capacity can be reduced by approximately 50 %. These savings can help in designing more power-efficient and leaner storage systems.**

Keywords-power consumption; high performance computing; parallel distributed file system; storage servers; compression;

I. Introduction

As we progress towards Exascale systems, the economic cost of energy consumption and the pressure exerted by power dissipation on cooling equipment are rapidly becoming major hurdles to the deployment of new HPC facilities. As of today, the most energy-efficient HPC supercomputers deliver close to 4.5 GFLOPS/W (10^9 floating-point operations per second, per watt) [1]. Simple arithmetic shows that building an ExaFLOPS system based on this technology would require about 220 MW, yielding this approach economically unfeasible. Even if we can maintain the considerable improvements experienced by the most energy efficient systems, the goal of building a 20 MW Exascale system will still be largely exceeded. If we want to continue enjoying the significant advances enabled by scientific computing and supercomputers during these past decades, a holistic investigation is needed to improve energy efficiency of HPC hardware and software.

High performance I/O is also a major stumbling block to reach the ExaFLOPS barrier. Considering that large-scale computations and simulations conducted to solve complex problems in several scientific domains increasingly produce large amounts of data, the ExaFLOPS performance goal is still far away. While CPU speed and HDD capacity have increased by roughly a factor of 1,000 every 10 years [2], HDD speed only develops at a slow pace: only a 300-fold throughput increase over the last 25 years. Minimizing the amount of data that is being stored in the storage subsystem can also help to improve application performance. Thus, data compression can be used for this purpose. Furthermore, due to the increasing electricity footprints, energy used for storage represents an important portion of the Total Cost of Ownership (TCO) [3]. For instance, the German Climate Computing Center's (DKRZ) HDD-based storage system is using 10,000 HDDs to provide a 7 PB file system for earth system science. Assuming a power consumption of 5–10 W for typical HDDs, this results in energy costs of 50,000–100,000 € per year for the HDDs alone. To make the problem worse, the growth of HDD capacity has recently also started to slow down, requiring additional investment to keep up with the increasing processing power.

In this paper, we analyze the usefulness of HPC storage servers that compress file data in order to reduce the required amount of storage hardware and the overall energy consumption, and thus the total cost of storage hardware usually found in supercomputers. As current-generation CPUs provide ample performance for data processing such as compression or encryption, we provide a case for turning on compression by default to reduce the number of required storage devices. In particular, the paper includes the following contributions:

1) We review the architecture of parallel distributed file systems for the special case of Lustre and present a power-performance profiling and tracing environment for these platforms.
2) We evaluate the performance of different compression algorithms using different synthetic and real scientific sets of input data.
3) We employ the parallel I/O Performance Benchmark (IOR) [4] and real scientific data to assess different compression algorithms of the Lustre parallel distributed file system with its Zettabyte File System (ZFS) [5] back-end.
4) We evaluate the impact of slowly-spinning, energy-aware HDDs used in storage servers as a way to reduce energy consumption in large storage servers.

The paper is structured as follows: In Section II, we present the parallel distributed file system Lustre and our power-performance measurement framework used to analyze the benefits of our proposed approach. Section III contains an extensive evaluation of our proposal to compress all data produced by HPC systems. We present related work in Section IV. Finally, we draw conclusions and discuss further work in Section V.

 ISBN: 978-1-61208-332-2

II. Background

In this section, we explain the power-performance framework that we use to conduct our evaluation. First, we describe the distributed parallel file system and the power-performance measurement framework we used. Parallel distributed file systems are used in HPC to aggregate storage devices existing in several storage servers. File systems provide an abstraction layer between the applications and the actual storage hardware such that application developers do not have to worry about the organizational layout or technology of the underlying storage hardware. To meet the high demands of current high performance computing applications, they distribute data across multiple storage servers. Moreover, they are designed to allow parallel access to multiple clients and cooperatively work with the same data concurrently. On the other hand, power measurements are needed to design new energy-aware approaches and meet consumption constraints. In this case, wattmeters attached to the server nodes are leveraged in a profiling and tracing framework that allows developers to correlate performance and power data.

A. *Parallel distributed file systems*

To perform our evaluation we use Lustre [6]. Lustre is an open-source parallel distributed file system and is widely used on current supercomputers. Lustre powers around half of the TOP100 supercomputers and almost one third of all TOP500 supercomputers [7]. In contrast to other proprietary solutions such as the General Parallel File System (GPFS) [8], it is possible to adapt, extend and improve Lustre since it is licensed under the GPL (version 2).

Lustre distinguishes between file system clients and servers. While all clients are identical, the servers can have different roles. First, the *object storage servers* (OSSs) manage the file system's data and are connected to possibly multiple *object storage targets* (OSTs) that store the actual file data. Second, the *meta-data servers* (MDSs) handle the file system's meta-data, such as directories, file names and permissions. Each *MDS* is connected to possibly multiple *meta-data targets* (MDTs) that store the actual meta-data. The general architecture of Lustre is illustrated in Figure 1 on this page. Both *MDTs* and *OSTs* use an underlying back-end file system to store their data. This file system is responsible for block allocation and hiding the internal disk layout from Lustre. Past versions of Lustre supported only a modified version of `ext4` [9] called `ldiskfs`. Current versions of Lustre also support ZFS [10].

ZFS is a local file system that offers a rich set of advanced features. Among others, it provides advanced storage device management, data integrity, as well as transparent compression and deduplication. It supports several compression algorithms: zero-length encoding `zle`, `lzjb` (a modified version of `LZRW1` [11]), `lz4` [12] and `gzip` (a variation of LZ77 [13]). Performing the compression on the servers has several advantages: the compression is completely transparent to other applications (including Lustre), no modifications of the client libraries or operating systems are necessary, and computations of the clients are not influenced by the CPU overhead of the compression.

B. *Power-Performance measurement framework*

To analyze power and performance metrics of the Lustre storage servers, we employ a version of the integrated framework presented in [14] that works in combination with

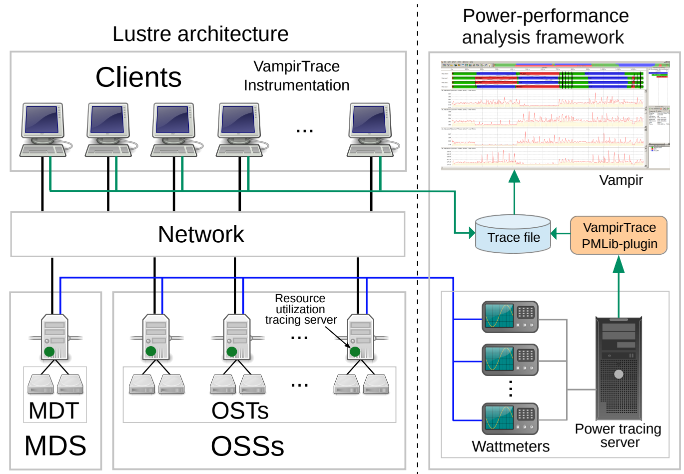

Figure 1. Lustre architecture in combination with the power-performance analysis framework.

VampirTrace and Vampir, which are profiling/tracing and visualization tools, respectively. The left part of the Figure 1 on the current page offers a graphical representation of the Lustre architecture; the right depicts the tracing and profiling framework. To use our approach, the application running on the clients is compiled using the VampirTrace compiler wrappers, which automatically instrument the code. Next, the application is run on the client nodes, thus consuming a certain amount of energy, both on the client and server nodes, due to the I/O instructions. The server nodes are connected to power measurement devices that account for the consumed energy and send the power data to the tracing server. The attached VampirTrace `pmlib` plugin employs the client API that sends start/stop primitives in order to gather captured data by the wattmeters onto the tracing server, where an instance of the `pmlib` server is running. Once the application's run has finished, the VampirTrace `pmlib` plugin receives the power data from the tracing server. The instrumentation post-process generates the performance trace files and the `pmlib` plugin inserts the power data into them.

In addition to the power measurements, we also account for the resource utilization values of the Lustre servers: CPU load, memory usage and storage device utilization. We run special `pmlib` server instances on the server nodes that retrieve these values from the `proc` file system (leveraging the `psutil` Python library). Thus, `pmlib` plugin instances running with the instrumented application connect with the `pmlib` servers and also receive resource utilization data of the Lustre servers. Finally, using the Vampir visualization tool, the power-performance traces can be easily analyzed through a series of plots and statistics.

III. Evaluation

A. *Methodology*

We leverage Lustre 2.5 with its ZFS back-end file system for our power-performance analysis of data compression in HPC storage servers. As the first step, we evaluate the compression algorithms that ZFS supports. For this purpose, we choose a real data-set that contains input and output data of the Max Planck Institute Ocean Model (MPI-OM), which is the ocean/sea-ice component of the MPI Earth System Model [15]. The data-set has a size of 540 GB, consisting of around 73 % binary data and 27 % NetCDF data [16]. In this evaluation, we also collect system statistics such as the average CPU load

 ISBN: 978-1-61208-332-2

and the execution time of the tests. To extract the compression ratio of the data-set we use the `zfs get compressratio` command that reports the actual compression ratio of the file system. Next, we evaluate the compression algorithms that are suitable for our needs with repeated and random patterns in order to understand the behavior in extreme cases. Using above information, we pick the best suitable algorithms in terms of compression ratio, CPU load and execution time to run our Lustre experiments. For now, we do not consider decompression because the evaluated compression algorithms are expected to have higher overheads when compressing than when decompressing.

In order to evaluate data compression in HPC storage servers we use the IOR, which is a commonly-used, configurable benchmark for parallel distributed file systems. We employ IOR to simulate realistic write activity of the Lustre storage servers and collect the execution time and throughput. To fulfill our needs we have modified IOR in the following ways: *(i)* we have inserted a sleep operation to simulate the computation phase, *(ii)* we have added a read phase in the initialization of the program in order to fill the write buffer with part of our scientific data-set, and *(iii)* we have instrumented it using the VampirTrace compiler wrappers to obtain power-performance-related traces.

To simulate the application's computation phase, we introduce a sleep instruction into our modified IOR. This simulates a common scenario in real-world applications, which alternate between computation and I/O phases. Consequently, the I/O systems see bursts of activity during the I/O phases and are otherwise relatively idle. For our cases, we choose to sleep approximately four times more than the time spent during write operations, which is reasonable for write-intensive applications. The data used to initialize the write buffer corresponds to the first 10 GB of an appropriate file from our scientific data-set with compression ratios of 1.7 and 2.53 for the `lz4` and `gzip` algorithms, respectively.

B. Environment Setup

The experimental setup includes a cluster with ten client and ten server nodes. Each client node is equipped with two Intel Xeon Westmere X5650 processors (6 cores each) running at 2.66 GHz, 12 GB of RAM and one 250 GB Seagate Barracuda HDD. Each storage server node has one Intel Xeon Sandy Bridge E31275 processor (4 cores) running at 3.40 GHz, 16 GB of RAM, one 160 GB Intel SSD and three 2 TB Western Digital Caviar Green HDDs. An energy-saving feature of these HDDs is the IntelliPark technology (also referred to as `idle3`) [17], which positions the read/write heads unloaded in a parking position and turns off certain drive electronics after a pre-defined inactivity time (8 seconds by default).

The client and servers nodes are interconnected using Gigabit Ethernet. Furthermore, a Lustre file system is provided by the ten server nodes. All ten nodes are configured as OSSs and use a ZFS pool containing one of their HDDs. Additionally, one node also fulfills the role of the MDS, using a ZFS pool containing the SSD; we use the SSD to exclude the influence of meta-data operations. We configure Lustre to stripe file data among all OSSs to get the best possible performance. The total amount of data that we write is 600 GB, that is, almost four times more than the total amount of RAM that the the OSSs are equipped with. When simulating client-side computation, we write 150 GB per I/O phase and then issue a sleep command to the clients for 600 seconds, which is approximately four times the length of the I/O phase.

TABLE I. COMPARISON OF DIFFERENT COMPRESSION ALGORITHMS.

Comp. Algorithm	Comp. Ratio	Avg. CPU Util. (%)	Runtime Ratio
`off`	1.00	23.7	1.00 (2:11 h)
`zle`	1.13	23.8	1.04
`lzjb`	1.57	24.8	1.09
`lz4`	1.52	22.8	1.09
`gzip-1`	2.04	56.6	1.06
`gzip-2`	2.05	62.3	1.05
`gzip-3`	2.02	71.9	1.11
`gzip-4`	2.08	73.0	1.07
`gzip-5`	2.06	80.2	1.21
`gzip-6`	2.04	84.7	1.88
`gzip-7`	2.05	85.1	2.36
`gzip-8`	2.06	86.8	4.79
`gzip-9`	2.08	83.1	13.66

We have observed that the best results were delivered with private files per process and consequently use this configuration. These results are to be expected because there is no locking overhead implied from the file system's POSIX semantics. We configure IOR with a block size of 256 KB for the following reasons: *(i)* it is aligned to the file system stripe size, *(ii)* it is large enough to minimize performance penalties and *(iii)* it fits the guidelines found in [18], which describes energy-efficient best practices for file system operations. Moreover, to avoid any caching effects on the client side, we allocate 85 % of their main memory. This helps achieving realistic measurements because real-world applications usually allocate as much memory as possible for computations.

Finally, to account for the power consumption we employ four external, calibrated ZES ZIMMER LMG450 devices connected between the power supply unit of the I/O servers and the electrical outlets. For the experiments we set the sampling frequency to 20 Hz.

C. Experimental Results

To evaluate the different compression algorithms supported by ZFS we copied the data-set into an uncompressed ZFS pool and set up another separate ZFS pool with the compression algorithm that we want to test. After copying the data-set into the compressed ZFS pool, we measured the compression ratio as well as the CPU utilization and runtime. Table I on this page shows a comparison of different compression algorithms supported by ZFS. As can be seen, disabling the compression yields the best runtime and lowest CPU utilization. The zero-length encoding (`zle`) adds negligible runtime and CPU overhead and already achieves a compression ratio of 1.13, which reduces our data-set to 470 GB. Both `lzjb` and `lz4` increase runtime and CPU utilization slightly but offer a significant boost for the compression ratio; even though the average CPU utilization is actually reduced with `lz4`, more CPU time is consumed due to the increased runtime. `lzjb` and `lz4` compress our data-set to 343 GB and 354 GB, respectively. The `gzip` compression algorithm additionally allows the compression level (1–9) to be specified. Even the lowest `gzip` compression level (`gzip-1`) further increases the compression ratio to 2.04 and compresses our data-set to 267 GB. While the runtime overhead is negligible, the CPU utilization more than doubles. The higher `gzip` compression levels do not significantly improve the compression ratio but increase the runtime and CPU utilization. Consequently, we choose `lz4` and `gzip-1` as the compression algorithms for our further analysis.

TABLE II. COMPARISON OF DIFFERENT DATA PATTERNS WITH SELECTED COMPRESSION ALGORITHMS.

Pattern	Comp. Algorithm	Comp. Ratio	Avg. CPU Util. (%)	Runtime Ratio
repeated	off	1.00	23.7	1.00 (11:13 min)
	lz4	126.96	15.8	1.28
	gzip-1	126.96	23.3	1.24
random	off	1.00	23.5	1.00 (11:21 min)
	lz4	1.00	24.1	0.97
	gzip-1	1.00	66.1	1.03

To have a better understanding of the algorithms that we use in our evaluation of the parallel distributed file system, we examine them with different data patterns. First, we use easily-compressible (repeated data obtained using the `yes` utility) and non-compressible (random data from the `frandom` RNG) data patterns. In this case, we have used a pre-generated 50 GB file and proceeded as in the previous evaluation. Table II on the current page shows a comparison of different data patterns using the previously selected compression algorithms `lz4` and `gzip-1`. As can be seen, compressing the repeated data incurs computational overhead. Both algorithms achieve the same compression ratio and almost the same runtime ratio. However, the `lz4` algorithm uses significantly less CPU time than the `gzip-1` algorithm (around 68 %). Both algorithms have a lower CPU utilization than the default configuration without compression; we believe that this is due to the fact that the dramatically reduced amount of data almost completely eliminates the overhead incurred by ZFS's checksumming and copy-on-write. When using the random data, again, both algorithms need approximately the same runtime and achieve the same compression ratio. In this case, the difference in CPU utilization is even more pronounced with the `lz4` algorithm using it only around 36 % as much as the `gzip-1` algorithm.

While the previous tests only used a single machine to get a basic understanding of the different compression algorithms and their associated overheads, the following parallel tests use the underlying Lustre file system. Table III on the following page shows the results for our write-only IOR benchmark, which does not include any computation. As can be seen, the measurement without compression (`off`) was used to set the baseline. The CPUs were moderately utilized with 116 %, while the HDDs were heavily utilized with around 87 %. The benchmark achieved a maximum throughput of 723 MB/s, which led to a runtime of 852 s. All ten I/O servers combined consumed 530 W during the benchmark run, which resulted in an energy consumption of 125 Wh. When using the `lz4` and `gzip-1` compression algorithms, the written data could be reduced significantly; `lz4` achieved a compression ratio of 1.92, while `gzip-1` even managed a ratio of 2.49. While `lz4` did not increase CPU utilization significantly (153 %), `gzip-1` produced a much higher computational overhead (402 %). In both cases, this reduced the HDD utilization and, at the same time, tremendously increased the throughput to 1,114 MB/s, which is the maximum our cluster can theoretically achieve. As a result, both benchmarks using compression finished much faster (553 s). However, in `gzip-1`'s case, the high power consumption of 848 W increased the energy consumption to 130 Wh. In `lz4`'s case, the moderate CPU and HDD utilization led to a slight increase in the power consumption (592 W). Combined with the shortened runtime, this resulted in significant savings regarding the energy consumption (91 Wh). Overall, `lz4` manages to reduce the total energy consumption to roughly 70 %.

TABLE IV. PARALLEL I/O BENCHMARK WITH SELECTED COMPRESSION ALGORITHMS (WRITE AND COMPUTATION).

Comp. Algorithm	Time (s)	Avg. Power (W)	Energy (Wh)	idle3 Timer
off	3,212.15	456.03	406.90	Disabled
lz4	2,951.01	461.72	378.48	
gzip-1	2,950.54	503.63	412.77	
off	3,181.35	436.18	385.46	Enabled (8 s)
lz4	2,951.11	437.33	358.50	
gzip-1	2,950.60	484.55	397.14	

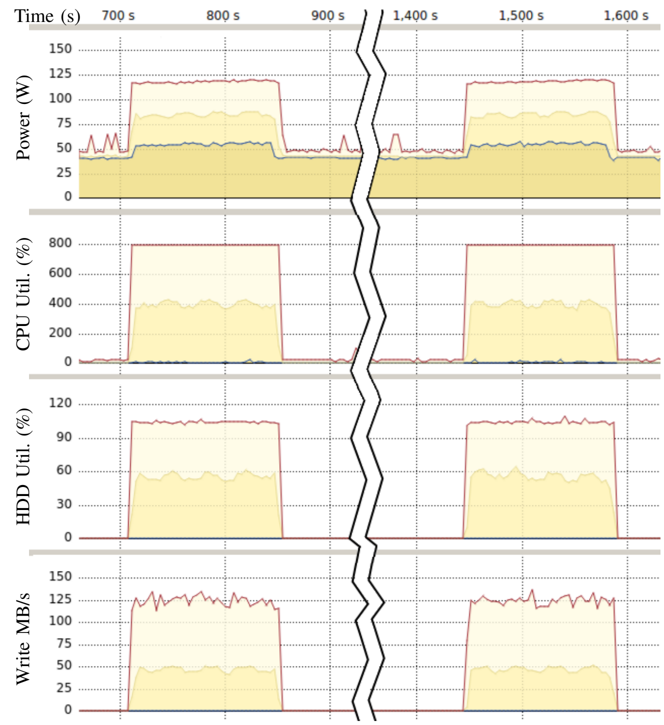

Figure 2. Vampir trace with the modified version of IOR benchmark using the `ondemand` governor and the `gzip-1` algorithm.

Next, we present our final experiments that are closer to real applications. Figure 2 on the current page shows a Vampir trace as recorded on one of the I/O storage servers while running our modified version of IOR benchmark. In this scenario the server alternates between idle and I/O state that correspond to the I/O activity and computation phase on the clients side.

The top of the figure shows the trace's timeline. As can be seen, the trace excerpt starts at around 670 s into the benchmark run and ends at approximately 1,620 s. The spikes on the left and right of the figure show I/O activity. To be able to focus on these important parts, we do not show most of the idle period in the middle. Our `pmlib` server collected data for the power consumption, CPU utilization, as well as HDD utilization and throughput. As can be seen, while the I/O server is being heavily utilized during the I/O phases, it is mostly idle during the computation phases. Since Vampir is not able to show all recorded values due to the lower screen resolution, it draws three lines representing the minimum, the average and the maximum for each of the values.

Table IV shows the results for our modified IOR benchmark, which simulates client-side computation; the rest of the configuration was not changed in any way. Due to the resulting idle phases on the I/O servers, we omit the HDD utilization and throughput. Additionally, we evaluated the IntelliPark (`idle3`)

TABLE III. PARALLEL I/O BENCHMARK WITH SELECTED COMPRESSION ALGORITHMS (WRITE-ONLY).

Comp. Algorithm	Comp. Ratio	Avg. CPU Util. (%)	Avg. Dev. Util. (%)	Throughput (MB/s)	Time (s)	Avg. Power (W)	Energy (Wh)
off	1.00	115.5	87.0	723.40	851.95	529.88	125.39
lz4	1.92	152.6	74.8	1114.03	553.29	592.07	90.99
gzip-1	2.49	402.3	55.8	1113.71	553.50	847.63	130.32

power saving feature supported by our Western Digital HDDs. As in the previous case, we repeated the benchmark using different compression algorithms; a run without compression (off) was used as the baseline.

The first set of results was conducted with the IntelliPark feature disabled. As can be seen, with 3,212 s, the benchmark now takes significantly longer to complete. While the average power consumption decreases to 456 W, the increase in run time leads to a total energy consumption of 407 Wh. For the lz4 and gzip-1 compression algorithms, the results look similar to the write-only test: Both algorithms help to significantly reduce the overall runtime by increasing the I/O throughput while, at the same time, increasing the average power consumption. In gzip-1's case, the power consumption of 504 W leads to a higher energy consumption of 413 Wh. lz4, however, only slightly increases the power consumption to 462 W, which results in an energy consumption of 378 Wh. While these energy savings are not as significant as the ones on the write-only test, lz4 still manages to reduce the total energy consumption to 93 %.

For the second set of tests, we enabled the IntelliPark mechanism and set the timeout to the default of 8 s. While IntelliPark helps minimizing power consumption of unutilized hard disk drives by intelligently parking the head of the disk, it also has its drawbacks: The disks may break sooner due to the high frequency of parking and unparking the disk head. Additionally, the warm up period required to resume operating from the park state adds extra latency to disk accesses. However, in HPC the hard drives are either fully utilized or idle for longer periods, which means that the IntelliPark technology fits perfectly for this use case. We measured the power consumption of a single server by enabling and disabling the IntelliPark technology on a single hard drive and observed a difference of approximately 1 W. The results of our experiment with IOR confirm that we can decrease the power and energy consumption by about 5 % across the board. Apart from that, we do not observe any other meaningful change.

IV. RELATED WORK

The energy efficiency of parallel I/O has been examined in several papers. Rong Ge et al. [18] provide an evaluation of parallel I/O energy efficiency using PVFS and NFS. They evaluated different I/O access patterns and block sizes, concluding that accessing larger data sets is more energy efficient. Moreover, they examined the potential benefits of leveraging DVFS in the compute nodes for I/O-intensive applications where lowering the clock frequency to half reduced energy consumption by around 25 % without compromising application performance. Our benchmarks are configured according to this paper. The authors in [19] propose Sera I/O, a portable and transparent middleware to improve the energy efficiency of I/O systems. Sera I/O creates a table from profiling information of micro-benchmarks. At run time, Sera I/O analyzes the I/O pattern of the application using the pre-created table and applies DVFS techniques accordingly. Takafumi et al. [20] propose an energy-aware I/O optimizer for check-pointing and restart on a NAND flash memory system. Based on profiles, the optimizer applies DVFS and controls the I/O processes, resulting in energy savings.

Mais Nijim et al. [21] integrate flash memory in the storage architecture to improve the energy efficiency of the disk subsystem. Using flash memory as a cache to keep the frequently used data sets, they are able to serve most of the I/O from the cache and put the majority of the disk drives into standby mode to save energy. The authors of DARAW [22] added write buffer disks to the system architecture to minimize energy consumption. Using a set of disks to temporarily store I/O accesses allows them to spin down storage disks for longer periods and reduce the power consumption. The impact of the disk speed on energy consumption is also presented in [23], [24]. In contrast to these papers, in our evaluation we investigate the impact of techniques that can be implemented in the disk controller to reduce energy consumption and do not require any modification in any other part of the system. However, above techniques can be combined with our proposal to improve even more the energy savings.

The authors of [25] explore the impact of compression in the storage servers by examining several file formats. However, our work is complementary to this and we are specially targeting HPC storage servers. Previous studies [26] have shown that scientific data can achieve high compression ratios depending on the used algorithm. Moreover, in comparison to our approach, which is server side compression, the authors in [27] propose compression on the client side, minimizing the amount of transmitted data between compute and storage nodes by sacrificing CPU cycles on the compute nodes.

Our previously conducted deduplication study for HPC data already showed great potential for data savings, allowing 20–30 % of redundant data to be eliminated on average [28]. However, deduplication can be very expensive in terms of memory overhead. The advantage of compression in comparison of deduplication is that is does not require any kind of lookup tables and is thus much cheaper to deploy because no additional hardware is required.

V. CONCLUSIONS AND FUTURE WORK

Our evaluation shows that data compression in HPC storage servers can be used to save energy and improve I/O performance. On the one hand, less HDDs are required to store the same amount of data due to the compression. On the other hand, it is also possible to achieve a higher throughput by storing more data in the same amount of time; this is especially relevant in I/O-intensive cases. These two advantages lead to lower procurement and operational costs.

However, it is important to carefully choose compression algorithms due to their inherent CPU overhead. Using expensive algorithms will increase power consumption and can theoretically even decrease performance. We have identified lz4 as a suitable compression algorithm for scientific data and will use it for further analysis in the future. Lustre's ZFS backend provides a convenient possibility to leverage these compression algorithms without modifying or influencing client applications. In more detail, using a real world data-set we demonstrated that we can achieve compression ratios of more than 1.5 without any significant increase of CPU utilization. Additionally, we observed a reduction in energy consumption of 30 % during the write phases and 7 % in write-intensive applications.

In the future, we plan to extend our evaluations by including additional factors such as energy-efficient SSDs, different CPU and network configurations as well as the overhead caused by decompression. Additionally, we want to take a closer look at other technologies, such as encryption, which are suitable to be handled by the storage servers. With current parallel distributed file systems, it is not possible to shut down individual storage nodes during operation; we will also investigate this possibility to enable additional energy savings.

Acknowledgments

We would like to thank the German Helmholtz Association's LSDMA project, the EU's EXA2GREEN project (FP7-318793) and the EU's SCALUS (FP7-PEOPLE-ITN-2008-238808) project for their support. We also wish to thank Petra Nerge from our group for providing access to the data-set used for the experiments.

References

[1] The Green500 Editors, "Green500," http://www.green500.org/, 12 2013, last accessed: 2013-12.

[2] R. Freitas, J. Slember, W. Sawdon, and L. Chiu, "GPFS scans 10 billion files in 43 minutes," IBM Advanced Storage Laborator. IBM Almaden Research Center. San Jose, CA, vol. 95120, 2011.

[3] M. L. Curry, H. L. Ward, G. Grider, J. Gemmill, J. Harris, and D. Martinez, "Power Use of Disk Subsystems in Supercomputers," in Proceedings of the Sixth Workshop on Parallel Data Storage, ser. PDSW '11. New York, NY, USA: ACM, 2011, pp. 49–54. [Online]. Available: http://doi.acm.org/10.1145/2159352.2159364

[4] Lawrence Livermore National Lab, "The IOR benchmark," http://sourceforge.net/projects/ior-sio, 2008, last accessed: 2013-12.

[5] J. Bonwick, M. Ahrens, V. Henson, M. Maybee, and M. Shellenbaum, "The Zettabyte File System," 2003.

[6] P. Schwan, "Lustre: Building a file system for 1000-node clusters," in Proceedings of the 2003 Linux Symposium, vol. 2003, 2003.

[7] D. Fellinger, "The State of the Lustre® File System and The Lustre Development Ecosystem," http://www.opensfs.org/wp-content/uploads/2013/04/LUG_2013_vFinal.pdf, 04 2013, last accessed: 2013-10.

[8] F. Schmuck and R. Haskin, "GPFS: A Shared-Disk File System for Large Computing Clusters," in Proceedings of the 1st USENIX Conference on File and Storage Technologies, ser. FAST '02. Berkeley, CA, USA: USENIX Association, 2002, pp. 231–244. [Online]. Available: http://dl.acm.org/citation.cfm?id=1083323.1083349

[9] A. Mathur, M. Cao, S. Bhattacharya, A. Dilger, A. Tomas, and L. Vivier, "The new ext4 filesystem: current status and future plans," in Proceedings of the Linux Symposium, vol. 2. Citeseer, 2007, pp. 21–33.

[10] Lawrence Livermore National Lab, "Native ZFS on Linux," http://zfsonlinux.org, last accessed: 2013-12.

[11] R. Williams, "An extremely fast Ziv-Lempel data compression algorithm," in Data Compression Conference, 1991. DCC '91., Apr 1991, pp. 362–371.

[12] Yann Collet, "LZ4 Explained," http://fastcompression.blogspot.com/2011/05/lz4-explained.html, 2008, last accessed: 2013-12.

[13] J. Ziv and A. Lempel, "A universal algorithm for sequential data compression," Information Theory, IEEE Transactions on, vol. 23, no. 3, May 1977, pp. 337–343.

[14] S. Barrachina et al., "An Integrated Framework for Power-Performance Analysis of Parallel Scientific Workloads," 3rd Int. Conf. on Smart Grids, Green Communications and IT Energy-aware Technologies, 2013, pp. 114–119.

[15] J. H. Jungclaus et al., "Characteristics of the ocean simulations in the Max Planck Institute Ocean Model (MPIOM) the ocean component of the MPI-Earth system model," Journal of Advances in Modeling Earth Systems, vol. 5, no. 2, 2013, pp. 422–446. [Online]. Available: http://dx.doi.org/10.1002/jame.20023

[16] R. Rew and G. Davis, "Data Management: NetCDF: an Interface for Scientific Data Access," IEEE Computer Graphics and Applications, no. 10-4, 1990, pp. 76–82. [Online]. Available: http://dx.doi.org/10.1109/38.56302

[17] Western Digital, "Western Digital Green," http://www.wdc.com/en/products/products.aspx?id=780, 12 2013, last accessed: 2013-11.

[18] R. Ge, "Evaluating Parallel I/O Energy Efficiency," in Proceedings of the 2010 IEEE/ACM Int'l Conference on Green Computing and Communications & Int'l Conference on Cyber, Physical and Social Computing. IEEE Computer Society, 2010, pp. 213–220.

[19] R. Ge, X. Feng, and X.-H. Sun, "SERA-IO: Integrating Energy Consciousness into Parallel I/O Middleware," in Cluster, Cloud and Grid Computing (CCGrid), 2012 12th IEEE/ACM International Symposium on, May 2012, pp. 204–211.

[20] T. Saito, K. Sato, H. Sato, and S. Matsuoka, "Energy-aware I/O Optimization for Checkpoint and Restart on a NAND Flash Memory System," in Proceedings of the 3rd Workshop on Fault-tolerance for HPC at Extreme Scale, ser. FTXS '13. New York, NY, USA: ACM, 2013, pp. 41–48. [Online]. Available: http://doi.acm.org/10.1145/2465813.2465822

[21] M. Nijim, A. Manzanares, X. Ruan, and X. Qin, "HYBUD: An Energy-Efficient Architecture for Hybrid Parallel Disk Systems," in Computer Communications and Networks, 2009. ICCCN 2009. Proceedings of 18th Internatonal Conference on, Aug 2009, pp. 1–6.

[22] X. Ruan, A. Manzanares, S. Yin, Z. Zong, and X. Qin, "Performance Evaluation of Energy-Efficient Parallel I/O Systems with Write Buffer Disks," in Parallel Processing, 2009. ICPP '09. International Conference on, Sept 2009, pp. 164–171.

[23] S. Gurumurthi, A. Sivasubramaniam, M. Kandemir, and H. Franke, "DRPM: dynamic speed control for power management in server class disks," in Computer Architecture, 2003. Proceedings. 30th Annual International Symposium on, June 2003, pp. 169–179.

[24] Q. Zhu, Z. Chen, L. Tan, Y. Zhou, K. Keeton, and J. Wilkes, "Hibernator: Helping Disk Arrays Sleep Through the Winter," vol. 39, no. 5. New York, NY, USA: ACM, Oct. 2005, pp. 177–190. [Online]. Available: http://doi.acm.org/10.1145/1095809.1095828

[25] R. Kothiyal, V. Tarasov, P. Sehgal, and E. Zadok, "Energy and Performance Evaluation of Lossless File Data Compression on Server Systems," in Proceedings of SYSTOR 2009: The Israeli Experimental Systems Conference, ser. SYSTOR '09. New York, NY, USA: ACM, 2009, pp. 4:1–4:12. [Online]. Available: http://doi.acm.org/10.1145/1534530.1534536

[26] N. Hübbe and J. Kunkel, "Reducing the HPC-Datastorage Footprint with MAFISC - Multidimensional Adaptive Filtering Improved Scientific data Compression," in Computer Science - Research and Development, Executive Committee. Hamburg, Berlin, Heidelberg: Springer, 2012.

[27] B. Welton, D. Kimpe, J. Cope, C. M. Patrick, K. Iskra, and R. Ross, "Improving I/O Forwarding Throughput with Data Compression," in Proceedings of the 2011 IEEE International Conference on Cluster Computing, ser. CLUSTER '11. Washington, DC, USA: IEEE Computer Society, 2011, pp. 438–445. [Online]. Available: http://dx.doi.org/10.1109/CLUSTER.2011.80

[28] D. Meister, J. Kaiser, A. Brinkmann, T. Cortes, M. Kuhn, and J. Kunkel, "A Study on Data Deduplication in HPC Storage Systems," in Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis, ser. SC '12. Los Alamitos, CA, USA: IEEE Computer Society Press, 2012, pp. 7:1–7:11. [Online]. Available: http://dl.acm.org/citation.cfm?id=2388996.2389006

Fostering Energy Awareness in Residential Homes using Mobile Devices

Florian Reinhart*, Kevin Schlieper*, Michael Kugler*, Elisabeth André*, Masood Masoodian† and Bill Rogers†
*Human Centered Multimedia
Augsburg University
Email: andre@informatik.uni-augsburg.de
†Department of Computer Science
The University of Waikato
Email: {masood, coms0108}@waikato.ac.nz

***Abstract*—There is considerable global effort being made towards identifying ways of reducing energy consumption to cope with growing demands. Although there is potential for energy saving in many sectors, our focus is on reducing energy consumption in residential homes. We have developed a system which combines home automation and energy usage monitoring technologies. The system offers a range of tools designed for mobile devices to assist users with monitoring their energy usage and provides mechanisms for setting up and controlling home appliances to conserve energy. In this paper, we describe our system and a user study we have conducted to evaluate its effectiveness. The findings of the study show the potential benefits of this type of mobile technology.**

***Keywords*–Ecological feedback technologies; user awareness of energy consumption; user interfaces; information visualization.**

I. Introduction

People require energy every day, for heating their water, for cooking their food, for cooling their houses, etc. Energy is used not only for basic human needs, but also for making people's lives more enjoyable and comfortable. Since we are using more and more energy for maintaining our life styles, the worldwide energy demand is steadily growing. In general, there are two trends that affect the energy supply industry: growing demand and sustainable generation. While worldwide energy consumption has been increasing over the past few decades [1], so has our capability and willingness to generate energy in a sustainable manner using renewable sources like solar or wind energy. For instance, in the United States alone the number of homes with photovoltaic installations grew by 40% in 2009 and this trend is likely to continue [2]. Although using renewable energy reduces the output of greenhouse gasses [3], these sources of energy are less reliable in terms of generation, and therefore require better management in terms of consumption.

Statistics show that about one-third (29%) of the energy used in the European Union [4] is consumed by residential homes (see Figure 1), which is very similar to other developed regions. Although, there are other sectors like industry which consume a lot of energy, private households offer a great potential for saving energy.

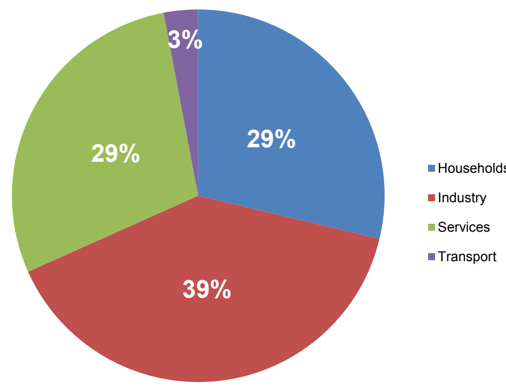

Figure 1: Electricity usage in the European Union

Here, we propose two strategies for conserving energy in private households. The first strategy is to use an automated system that controls home appliances, so that energy-wasting behavior by users can be partly mitigated (e.g., automatically turning off the heating when windows are left open). The second strategy is to use mobile technology to assist residents to save energy and encourage them to change their energy consumption behavior. To do this, however, the users must first know how they can improve their usage behavior, and one prerequisite for this is to be are aware of how much energy they are consuming.

In general, the process of a non-automated approach to energy saving in private households can be broken down into two main parts. First, the residents must identify saving potentials in their household, which requires them to be aware of the energy consumption of individual appliances. Second, once they know how and where energy is actually being used, they need to be assisted and persuaded to change their behavior to reduce their energy usage. Unfortunately, however, there are some challenges in achieving both of these two parts. In relation to the first part, usually people know their overall energy consumption, because that is what their energy providers bill them for. Nevertheless, finding out how much energy each device actually uses or how much of

 ISBN: 978-1-61208-332-2

the overall consumption each individual person in a house has consumed, is very difficult. As challenging as solving this essentially technical problem may be, overcoming the challenge of changing people's behavior to save energy while still living a comfortable life is even harder.

In this paper, we introduce a system, called Ubiquitous Smart Energy Management (USEM), which not only provides an automated solution for reduction of electricity usage, but also caters for the non-automated approach by providing detailed energy consumption information to users, and incorporates mobile tools to assist and encourage users to change their energy consumption behavior. The focus of this paper is mainly on these non-automated components of USEM designed for mobile devices.

We start this paper with a review of related literature on technologies for energy saving in general (Section II), and at residential homes in particular (Section III). We then present the mobile user interface components of USEM (Section IV), and discuss a laboratory-based study we have conducted to evaluate the usability of these mobile tools (Section V). Finally, we briefly provide the results of an analysis of the capabilities of USEM against existing guidelines for the design of persuasive technology (Section VI), and draw some conclusions (Section VII).

II. Encouraging Energy Saving

Although one might assume that people would naturally be inclined to save energy, not only because this is good for the environment, but also because it leads to saving money, this does not always seem to be the case. Therefore, in recent years several technologies have been proposed to encourage energy saving by fostering awareness of energy use. These technologies can be categorised into those that utilize games or social media, and those that rely on ambient devices.

A. Games and Social Media

Several systems have been developed to encourage people to conserve energy and increase their energy use awareness through games and social media. The Power Explorer [5] game tries to help teenagers save energy. This mobile phone game takes into account the changes in energy consumption at home by the players. There are different game elements: habitat, pile and duels. The habitat shows the user's avatar in a virtual climate environment, in which energy usage causes CO_2 clouds to appear, which is bad for the avatar. In the pile view, players can see how they are ranked compared to other players, and in the duels players compete directly against each other. The goal of the duels is to increase the energy awareness about the appliances, since players have to adjust their household energy consumption to win. A study of Power Explorer showed that a group players consumed about 20% less energy than a reference group of non-players.

Other research has focused on integration of home energy feedback into social networks. For instance, Mankoff and colleagues [6] demonstrated integration of energy usage feedback to the MySpace social network to motivate people to conserve energy. Similarly, Foster and colleagues [7] have developed a Facebook application, and have shown in a study that energy consumption can be reduced through social encouragement and competition. Petkov and colleagues [8] expand the idea of social comparison with their social application EnergyWiz in which users can compare their energy usage with their own history and that of others.

Midden and Ham [9], on the other hand, performed a laboratory-based experiment in which participants could save energy while using a simulated washing machine. This study showed that social feedback provided by an embodied agent was more effective than just factual feedback about the energy savings made.

These types of social network related systems rely on surveillance and self-monitoring techniques. However, they generally only provide feedback at the household level and not at the individual user's level.

B. Ambient Devices

Kim and colleagues [10] outline design requirements for ambient devices to create effective persuasion. In a study they identify ten stages from raising awareness to behavior change and the maintenance of behavioral changes. Based on their findings, they then propose several persuasion methods including virtual and visually attractive rewards, personalized and tailored feedback, self-monitoring, and use of social media to compare one's own performance with others.

The Energy AWARE Clock [11] is an example of an ambient device that visualizes current and past energy usage of a household. The three design principles of complexity, visibility and accessibility are used to reduce the complexity of consumption data, make visible "hidden" or not directly obvious electricity consumers, and have the consumption data easily accessible. A three month user study of nine households showed that the users developed a better awareness of their energy use, and thought about changing their behavior to save energy.

Other ambient devices have been developed to help users save other resources such as water. Examples of these include UpStream [12] and Shower Calendar [13]. Studies of these systems have shown that these systems lead to reduction in water consumption.

Ham and colleagues [19] conducted a study to see if ambient technology has the capability of persuading people subconsciously. In this study, the participants were asked to rate the energy usage of three devices. The three groups of participants either received supraliminal feedback (150 ms), subliminal feedback (25 ms) or no feedback at all on their given answers. The feedback was given in the form of smiley faces directly after rating the consumption of a device. The results indicated that both groups with feedback gave more correct answers on average than the group without any feedback. Furthermore, the subliminal feedback group gave comparable answers to the supraliminal feedback group, and they also stated that they had not consciously seen any feedback.

III. Energy Saving at Home

Technologies available to residential users to assist them with saving energy at home can be divided into two broad

 ISBN: 978-1-61208-332-2

TABLE I: Comparison of existing energy usage monitoring and home automation systems

	Energy usage monitoring	Controlling appliances	Manual task scheduling	Automatic task scheduling	Support for off-grid renewable energy	Support for off-the-shelf devices	Ease of installation	Extensibility
Solar Home [14]	yes	yes	no	no	yes	yes	difficult	no
AIM [15]	yes	yes	yes	no	no	no	difficult	no
Energy Aware Smart Home [16]	yes	no	no	no	no	no	moderate	no
Intel Home Energy Management [17]	yes	yes	yes	no	no	yes	easy	yes
HESS [18]	yes	yes	no	no	no	yes	difficult	no

categories: energy use measurement and home automation systems.

Energy use measurement systems are able to record the energy usage for either individual home appliances or the entire household. Energy consumption statistics can then be generated and analyzed by the user to identify where energy savings could be made. Examples of such systems are Current Cost [20], the Energy Detective [21] and Wattson [22]. However, these systems only show users how much energy they have consumed in the past, and at best make some general suggestions about how to reduce energy usage in the future. These types of systems cannot actively control appliances to put energy saving tips into practice.

Home automation systems, on the other hand, can actively control energy consuming devices in a household, usually using a range of sensors that react to their environment (e.g., an air-conditioning system can be automatically switched on/off based on temperature sensor data). Examples of such systems include HomeMatic [23], Gira [24] and Intellihome [25]. These types of system however are not very widely used.

One of the reasons why home automation systems are not used is because there are currently no universally adopted standards for home automation. Although standards such as KNX [26] exist, the majority of home appliances do not conform to them. Without these standards, controlling off-the-shelf devices is usually limited to switching them on/off. Although this can be sufficient for controlling simple devices (e.g., lights), it is not useful for controlling more complex appliances (e.g., a washing machine).

There are several energy related prototype systems for residential homes that aim to deal with these problems. These system provide either or both of home automation and energy usage monitoring capabilities. Table I gives a summary of the capabilities and limitations of these systems against a number of important criteria. As can be seen, none of the systems satisfies all the identified requirements.

IV. USEM Prototype System

We have developed USEM [27] to overcome some of the limitations of the existing systems, as summarized in the previous section. Like some of these systems, USEM supports residential users in conserving energy by combining the functionalities of home automation and energy usage monitoring systems.

To make people aware of their energy usage, USEM provides detailed statistics about the household's past energy consumption. For example, the overall consumption can be displayed for individual rooms, devices, or occupants of the house. These statistics allow the users to analyse their consumption history and, thus, identify saving potentials. In some cases USEM might also be able to suggest actions that would lead to a decrease of energy consumption. Furthermore, USEM supports the user to put theoretical energy saving ideas into practice. For instance, the user can schedule appliances to run when varying energy rates are the cheapest, or USEM can switch off devices when they are not needed (e.g., turning off the printer whenever the PC is switched off). By providing an intelligent scheduling for energy consuming tasks USEM is also able to use energy when it is available in off-grid cases (e.g., when the solar panels are generating electricity). Using a combination of these techniques USEM attempts to ensure that energy usage peaks are avoided, maximum renewable energy is used when available, and overall power usage reduced in an intelligent manner without necessarily reducing comfort levels.

To interact with USEM we have developed three different user interfaces (1) a web interface for performing more complex tasks such as managing manual and automatic task scheduling, (2) a tablet interface acting as a control unit that could be used from around the house, and (3) a mobile phone interface that can be used to interact with USEM while on the go.

The web interface has not specifically been tailored towards mobile technology. As mentioned, it provides higher-level access to the functionality of the system, and allows configuration, scheduling, and visualization of relevant information. The

Figure 2: Home screen of the tablet interface of USEM

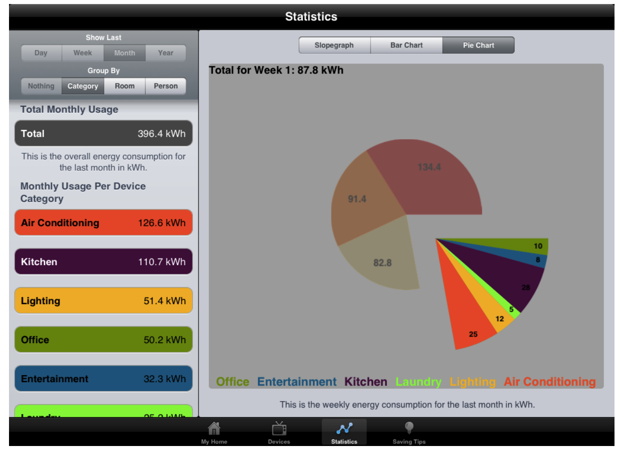

Figure 3: Usage information screen of the tablet interface showing a pie chart

web interface, along with the more intelligent components of USEM required for task scheduling, etc. have been discussed elsewhere [28]. The focus of this paper, on the other hand, is on the tablet and mobile phone components of USEM.

A. Tablet Interface

The tablet interface acts as a control and access unit for USEM. Although the current application has been developed for an Apple iPad, it is envisaged that it could also in the future be installed in flat-panel displays incorporated into furniture, picture frames or walls to act as an ambient device interface.

Figure 2 shows the home screen of the tablet interface which is visible when the device is not being controlled by the user. This allows the users to have a constant view of the most important information about their household, which encourages them to monitor their energy use. The home screen is customizable with several widgets to display information such as a list of currently running devices, up-to-date energy prices when available, etc. This screen also shows the energy usage target set by the user, the current usage level, to motivate the user to keep their usage below their set target. If the target is being threatened, for instance when the user turns on a device, they get a warning from the system giving them the option of turning off a device which is not being used (if any) or not going ahead with their scheduled activity.

This interface also allows users to view their energy consumption information over the past year, month, or day. Figure 3 illustrates one of the energy usage visualization. This information can be viewed in several different chart formats, and in various categories such as for the entire house, different rooms, all users, different users, all devices, different category of devices, etc. This is another important element of the user interface in terms of encouraging energy usage awareness.

Furthermore, the tablet interface gives energy saving recommendations, based on the past and current energy consumption data, to help users reduce their energy use. Figure 4 presents an example energy saving tips screen. On this screen the system suggests actions that would decrease the household's energy

Figure 4: Energy saving recommendations and their consequences if applied

consumption, as well as calculating the savings that could be made if the advice is followed.

B. Mobile Phone Interface

The mobile phone interface, developed for Apple iPhone, can be used to retrieve the status of home appliances or to interact with them remotely. It also notifies the user about energy usage events that occur while the user is away (e.g., a scheduled task cannot be undertaken because there is not enough renewable energy available). In such cases, the mobile phone interface provides suggestions (Figure 5) about how the problem could be resolved and gives the user the opportunity to decide what to do (e.g., cancel a scheduled task, or turn off another device).

Of course, the mobile phone interface can also be useful while the user is at home. For instance, it can be used to remotely access any of the devices connected to USEM

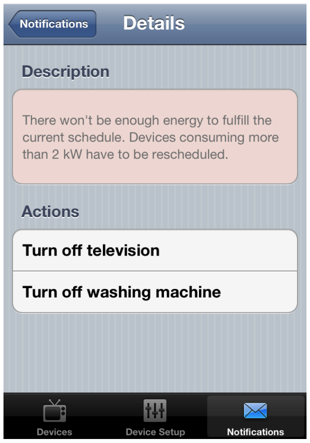

Figure 5: Event notification screen of the mobile phone interface

Figure 7: The set-up used for the user evaluation

Figure 6: USEM devices (left) and new device set-up (right) screens of the mobile phone interface

(Figure 6, left). Users can also directly interact with home appliances by scanning the unique barcode that is attached to each device when they are added to USEM. For example, this allows creating a new task for the washing machine right after the user has put the laundry into the machine. The mobile phone application also simplifies the initial set-up procedure for new devices, by mapping the barcode tags attached to power sockets with those of devices connected to them through a simple scanning process (Figure 6, right).

V. User Study of USEM

We conducted a user study to evaluate the usefulness of USEM and gauge if people would actually use a system like USEM to save energy if they had access to it. In the following sections, we describe this study and discuss some of its main findings.

A. *Methodology*

The study was conducted at a usability lab, where the participants performed a series of tasks using the web, tablet, and mobile phone interfaces. To do the tasks the participants were provided with a laptop, an iPhone, an iPad, and model of a dryer and a computer as two home appliances (each with a barcode attached), as shown in Figure 7. We also attached a barcode to a power socket to make it recognizable by the mobile phone interface of USEM.

Each session started with a tutorial, which included some sample tasks similar to the actual tasks that the participants performed after the tutorial. The study session took about an hour in total, and was divided into three parts covering the use of the three interfaces of USEM. The sessions started with the web interface, as this was the most general part of the system and gave the user a comprehensive overview of USEM (this part of the study is not relevant to this paper and is not discussed here). This was followed by tasks that users performed using the mobile phone and tablet interfaces.

At the end of each task the participants answered several questions related to the task and the tool they had just used. Further, at the end of the session the participants completed a final questionnaire covering some questions about the users' overall impression of USEM.

B. *Study Participants*

Twenty participants took part in this study. They were between 20 and 62 years old, with an average age of 35. Five of them were female and 15 male; 11 were students, two researchers, two managers, two office administrators and one housewife. Thirteen of the participants (65%) had previous experience using a multi-touch screen; 11 (55%) owned a smartphone and 4 (20%) owned a tablet device. All of the participants used a computer daily, none had any previous knowledge of USEM. Table II shows a summary of the participants' demographic data.

 ISBN: 978-1-61208-332-2

TABLE II: Demographic of the study participants

	No. of Participants	Percentage
Gender		
Male	15	75.00 %
Female	5	25.00 %
Occupation		
Student	11	55.00 %
Researcher	2	10.00 %
Other	7	35.00 %
Experience		
Multi-touch screen	13	65.00 %
Daily PC usage	20	100.00 %
Own device		
Smartphone	11	55.00 %
Tablet	4	20.00 %

Figure 8: Usage information screen of the tablet interface showing a slopegraph

C. Study Tasks

As mentioned earlier, the study participants had to perform specific tasks using each of the different USEM mobile device interfaces. We asked the user to perform the following tasks on the mobile phone:

1) Adding a new home appliance to USEM.
2) Controlling a home appliance remotely.
3) Creating a new scheduled task to let USEM execute it at a later time.
4) Sending users a demo notification and asking them to react accordingly.

The participants then performed the following tasks using the tablet interface:

1) Controlling a home appliance remotely.
2) Exploring energy consumption statistics using a bar chart visualization.
3) Exploring energy consumption statistics using a pie chart visualization.
4) Exploring energy consumption statistics using a slope-graph visualization.

Figure 8 shows an example of type of slopegraph [29], Park [30] used in this study.

D. Task Questionnaires

As mentioned earlier, after the completion of each task the participants were asked to answer a questionnaire. Two questions were common to all the task questionnaires. These are:

1) How easy was it to perform this task?
2) How useful would it be to have this functionality?

The participants answered these questions using a Likert scale of 1-7, with 1 being the least positive and 7 the most positive.

E. Final Questionnaire

To evaluate the participants' overall impression of USEM we asked them the following questions as part of a final questionnaire:

1) How likely do you think it would be that you would decrease your energy consumption with the help of the visualizations on the iPad?
2) Would you want visualizations like on the iPad to be a permanent part of your home?
3) How often would you use your mobile phone to control your appliances remotely while you are away from home?
4) How often would you like to be notified on your mobile phone about what is going on in your household in terms of energy consumption?
5) Would you adapt your daily routine in order to use more renewable energy? (e.g., start cooking dinner an hour later?)
6) How useful do you find the overall system with regard to efficient energy usage?

F. Study Results

1) Results for the mobile phone interface: Figure 9 provides a summary of the average ratings given by the study participants for each of the two questionnaire questions for each of the 4 tasks performed using the mobile phone interface. As the results show, the participants generally found the tasks easy to perform, and the functionality provided by the interface useful.

Task 3 was the only task that was considered as being slightly more challenging than the other tasks. However, the functionality needed to perform this task was still rated as being useful. It is also important to note that this task was the most abstract task, which relied on the intelligent scheduling components of USEM.

2) Results for the tablet interface: Figure 10 shows a summary of the average ratings given by the study participants for each of the two questionnaire questions for each of the 4 tasks performed using the tablet interface. Since the difficulty of the tasks steadily increased, the ease-of-use rating for the tasks decreased slightly from Task 1 to Task 4.

However, in general all task have been rated as easy to perform with average ratings ranging from 6.30 to 6.75. The ratings given to the usefulness of the functionality provided by the tablet interface for performing each of the tasks has a

 ISBN: 978-1-61208-332-2

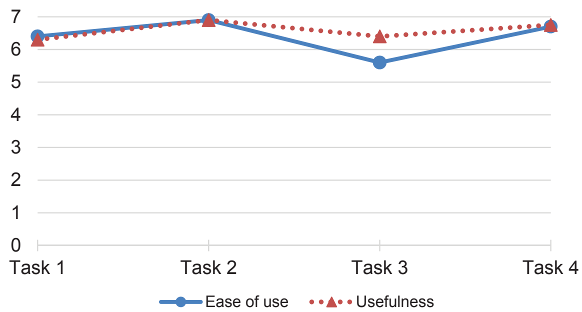

Figure 9: The average ratings given by the participants for the tasks performed using the mobile phone interface

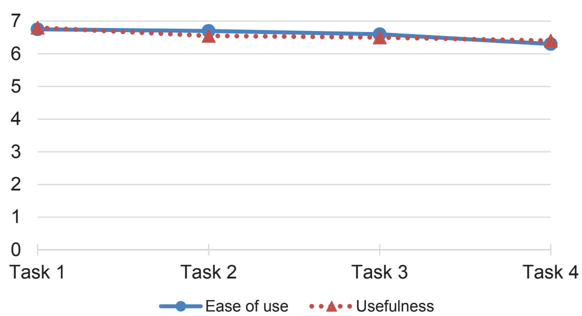

Figure 10: The average ratings given by the participants for the tasks performed using the tablet interface

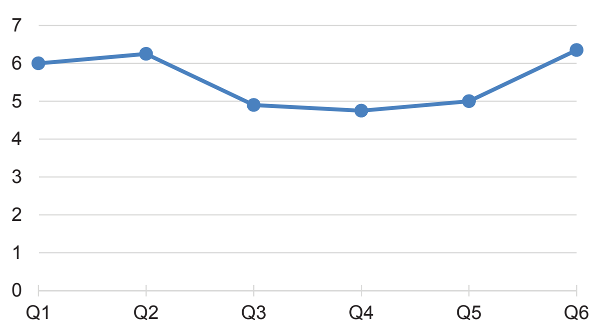

Figure 11: The average ratings given by the participants for the questions of the final questionnaire

similar trend to the difficulty of the tasks. Once again, overall the participants found the functionality provided very useful.

3) Results of the final questionnaire: The results of the final questionnaire are summarized in Figure 11. Perhaps the most important finding we can conclude from the final questionnaire is that the participants believe that USEM would be useful in helping them use energy more efficiently (Question 6). It is also important to note that the participants believe they would change their daily routines in order to use more renewable energy (Question 5).

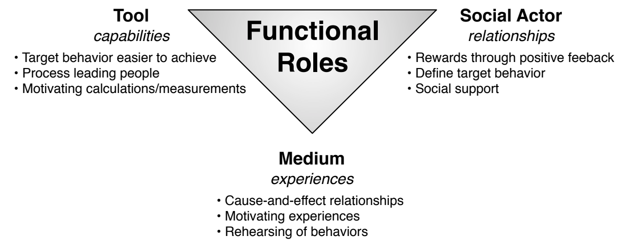

Figure 12: The roles computer technology can play in persuasive context, as defined by Fogg [32]

The following is a summary of some of the main points made by the study participants in their comments to questions of the final questionnaire.

- Many of the participants stated that they would like to be able to determine what events they should be notified about. They feared that they would get annoyed or distracted by notifications if they did not have some control over the notifications sent to their mobile phone.
- Most of the participants especially liked the possibility of controlling all the appliances using a single interface, rather than using a variety of different user interfaces for controlling different home appliances.
- The participants confirmed that they do not know how much energy each of their appliances uses. They rate their energy usage awareness as relatively low. They liked the energy usage visualizations provided by USEM, and thought that these would assist them to better understand the power usage of their appliances.
- Several participants stated that they do not have a good understanding of the kWh measurement unit. Instead they would prefer some kind of visualization, which is easier to understand and does not require any technical knowledge. They also suggested to display dollar amounts, and setting the saving goal in dollars as well.
- One participant commented that he would like recommendations for a saving goal. In this participant's opinion, it is difficult to set a saving goal, since it might be hard to determine a realistic energy consumption limit. So, the system could provide a recommendation for a feasible saving goal based on previous usage data.

VI. Persuasive Aspects of USEM

It is important to note that tools and technologies, such as USEM, which aim to assist people with changing their behavior need to be "persuasive" in their approach. The idea of *computers as persuasive technology* was introduced by Fogg [31] to deal with the question of how interactive computer technology could be used to persuade people to change their behavior or attitude. The functional triad, as defined by Fogg [32], is a framework that illustrates three roles computers can play in a persuasive context (Figure 12). These roles are categorized as tool, medium and social actor.

In the context of the work presented in this paper, we are mainly concerned with computers as persuasive tools.

Tools increase capabilities by making the desired behavior easier to achieve, by guiding people through processes, or by calculations and measurements that motivate people to reach their goals. There are seven different categories of persuasive technology tools [32], which can be combined together in a single system or application.

1) **Reduction:** People can be persuaded by reducing complexity. A good example of a reduction is the *one-click buy* functionality provided by Amazon [33], which reduces the ordering process to a simple button click.
2) **Tunneling:** This is the process of leading a user step-by-step through a specific procedure. An example of tunneling is the ordering process of online shopping sites. Such a guided process can provide opportunities for persuasion along the way. For instance, an online shopping site can suggest other items of interest to the buyer during the ordering process.
3) **Tailoring:** This approach persuades through customization, by providing only the type of information which is relevant and interesting to the user. An example of this is customized newsletters sent to users offering them products that match their buying profiles.
4) **Suggestion:** This means providing suggestions at the right moment. An example of this is advertisements along a highway, that for instance place an advert for a restaurant near its physical location and not miles away.
5) **Self-Monitoring:** People like to control themselves and check whether they have reached a predetermined goal. An example of this is a heart rate monitor that can be used to monitor the heart rate during exercise.
6) **Surveillance:** People tend to change their behavior when they know that they are being observed. An example of this is messages like "How am I driving?" at the back of some delivery trucks, to ensure that the drivers know people can complain about their bad driving, so they drive more carefully.
7) **Conditioning:** Giving positive, or negative, reinforcement can have a persuasive effect. An example of positive reinforcement is being on the high scorers list in a computer game, which can persuade people to play the game longer to improve their placement on the list.

To measure the success of a system as a persuasive technology clearly requires a long-term study of the use of the system in real-life settings to see if it indeed assists its users with changing their behavior. Although we are yet to conduct such a study of USEM, we have attempted to analyze the ways in which USEM might be able to play the role of persuasive technology listed above. Below, we provide a summary of this analysis.

1) **Reduction:** USEM reduces the complexity of the large volume of energy usage data, collected for many devices over an extended period, by categorizing it, and allowing the user to view it in a variety of forms.
2) **Tunneling:** USEM provides step-by-step guidance for dealing with the process of adding new appliances to the system, dealing with notifications, managing energy saving targets when they are breached, etc.
3) **Tailoring:** Energy usage information provided is tailored to individual users (i.e., their personal data), energy saving recommendations provided are tailored to each specific USEM installation and are always relevant to the context.
4) **Suggestion:** When USEM warns users about missing their targeted saving goals, it suggests what actions could be taken, for instance by giving a list of devices that could be turned off. Also, when USEM sends notifications to the mobile phone interface when scheduled tasks cannot be undertaken, it provides a list of suggestions that the user can select from.
5) **Self-Monitoring:** By measuring energy usage of each individual (when possible), USEM allows them to monitor their own current performance against targeted saving goals, as well as allowing them to monitor their past usage history in various statistical visualization forms.
6) **Surveillance:** Due to the fine granularity of energy usage data that USEM collects, the user knows (even when living in a house with others) that their consumption behavior is recorded and can be tracked by others when allowed.
7) **Conditioning:** By allowing users to compare their own energy usage behavior to others, as well as their set targets, USEM provides users with positive or negative reinforcements depending on their performance.

VII. Conclusion

This paper has presented the mobile interface components of the USEM system, which aim to support the inhabitants of residential homes with the process of monitoring their energy usage, and making energy savings possible without necessarily reducing their comfort levels. USEM allows its users to connect and control their home appliances, as well as analyze and understand energy consumption information by those appliances using a range of scheduling, notification, and visualization tools specifically developed for mobile devices.

Our laboratory-based user evaluation of USEM has shown the potential benefits of its interface components for mobile devices in providing the necessary means of assisting people with saving energy, as well as encouraging them to monitor and change their energy use behavior.

We have also briefly analyzed the capabilities of USEM as a persuasive technology, by examining some the features of its mobile interface components against existing guidelines for the design of persuasive technology. Although this analysis shows that USEM satisfies these guidelines, it is important to conduct a more formal real-life user evaluation of the persuasive capabilities of USEM.

We are therefore planning to deploy USEM in residential homes to carry out a long-term study to evaluate the effectiveness of USEM in changing users' behavior in terms of making energy conservation.

Acknowledgements

The user study presented in this paper was approved by the Ethics Committee of the Faculty of Computing and Mathematical Sciences, The University of Waikato. We would like to gratefully acknowledge the contributions of our study participants.

This research has been supported by the IT4SE project, funded by the German Federal Ministry of Education and Research (Grant number NZL 10/803 IT4SE) under the APRA initiative. More information about the IT4SE project can be found at http://www.it4se.net.

REFERENCES

[1] U.S. Energy Information Administration, World Energy Use and Carbon Dioxide Emissions, 1980-2001, April 2004.

[2] L. Sherwood, "U.S. solar market trends," in Proceedings of the Solar Conference 2010, 2010.

[3] J. K. Casper, Greenhouse Gases: Worldwide Impacts, ser. Global Warming. Infobase Publishing, December 2009.

[4] E. Commission, Energy: Yearly statistics 2008. Publications Office of the European Union, June 2010.

[5] A. Gustafsson, M. Bång, and M. Svahn, "Power explorer: A casual game style for encouraging long term behavior change among teenagers," in Proceedings of the International Conference on Advances in Computer Enterntainment Technology, ser. ACE '09, October 2009, pp. 182–189.

[6] J. Mankoff, D. Matthews, S. R. Fussell, and M. Johnson, "Leveraging social networks to motivate individuals to reduce their ecological footprints," in Proceedings of the 40th Hawaii International Conference on System Sciences, ser. HICSS 2007, January 2007, pp. 87–96.

[7] D. Foster, S. Lawson, M. Blythe, and P. Cairns, "Wattsup?: Motivating reductions in domestic energy consumption using social networks," in Proceedings of the 6th Nordic Conference on Human-Computer Interaction: Extending Boundaries, ser. NordiCHI '10, October 2010, pp. 178–187.

[8] P. Petkov, F. Köbler, M. Foth, and H. Krcmar, "Motivating domestic energy conservation through comparative, community-based feedback in mobile and social media," in Proceedings of the 5th International Conference on Communities and Technologies, ser. C&T '11, June 2011, pp. 21–30.

[9] C. Midden and J. Ham, "The persuasive effects of positive and negative social feedback from an embodied agent on energy conservation behavior," in Proceedings of the AISB 2008 Symposium on Persuasive Technology, ser. AISB 2008, vol. 3, April 2008, pp. 9–13.

[10] T. Kim, H. Hong, and B. Magerko, "Design requirements for ambient display that supports sustainable lifestyle," in Proceedings of the 8th ACM Conference on Designing Interactive Systems, ser. DIS '10, December 2010, pp. 103–112.

[11] L. Broms et al., "Coffee maker patterns and the design of energy feedback artefacts," in Proceedings of the 8th ACM Conference on Designing Interactive Systems, ser. DIS '10, December 2010, pp. 93–102.

[12] S. Kuznetsov and E. Paulos, "UpStream: Motivating water conservation with low-cost water flow sensing and persuasive displays," in Proceedings of the 28th international conference on Human factors in computing systems, ser. CHI '10, April 2010, pp. 1851–1860.

[13] M. Laschke, M. Hassenzahl, S. Diefenbach, and M. Tippkämper, "With a little help from a friend: A shower calendar to save water," in CHI '11 Extended Abstracts on Human Factors in Computing Systems, ser. CHI EA '11, May 2011, pp. 633–646.

[14] N. Banerjee, S. Rollins, and K. Moran, "Automating energy management in green homes," in Proceedings of the 2nd ACM SIGCOMM workshop on Home networks, ser. HomeNets '11, 2011, pp. 19–24.

[15] A. Capone, M. Barros, H. Hrasnica, and S. Tompros, "A new architecture for reduction of energy consumption of home appliances," in Proceedings of the European conference TOWARDS eENVIRONMENT, March 2009.

[16] M. Jahn et al., "The energy aware smart home," in Proceedings of the 5th International Conference on Future Information Technology, ser. FutureTech '10, May 2010, pp. 1–8.

[17] Intel Corporation, "Intelligent home energy management proof of concept," http://www.intel.com/embedded/energy/homeenergy/323157.pdf, [retrieved: February, 2014].

[18] K.-S. Choi et al., "Architectural design of home energy saving system based on realtime energy-awareness," in Proceedings of the 4th International Conference on Ubiquitous Information Technologies & Applications, ser. ICUT '09, December 2009, pp. 1–5.

[19] J. Ham, C. Midden, and F. Beute, "Can ambient persuasive technology persuade unconsciously? Using subliminal feedback to influence energy consumption ratings of household appliances," in Proceedings of the 4th International Conference on Persuasive Technology, ser. Persuasive '09, April 2009, pp. 29:1–29:6.

[20] Current Cost Ltd., "Current Cost," http://www.currentcost.com, [retrieved: February, 2014].

[21] Energy Inc., "The Energy Detective (TED)," http://www.theenergydetective.com, [retrieved: February, 2014].

[22] DIY Kyoto, "Wattson," http://www.diykyoto.com, [retrieved: February, 2014].

[23] eQ-3 AG, "HomeMatic," http://www.homematic.com, [retrieved: February, 2014].

[24] Gira Giersiepen GmbH, "Gira," http://www.gira.de, [retrieved: February, 2014].

[25] Intellihome Automatisierungstechnik GmbH, "Intellihome," http://www.intellihome.com, [retrieved: February, 2014].

[26] KNX Association, "KNX Standard," http://www.knx.org/knx-standard/standardisation, [retrieved: February, 2014].

[27] M. Kugler et al., "Architecture of a ubiquitous smart energy management system for residential homes," in Proceedings of the 12th Annual Conference of the New Zealand Chapter of the ACM Special Interest Group on Computer-Human Interaction. New York, NY, USA: ACM, 2011, pp. 101–104.

[28] ——, "Assisting inhabitants of residential homes with management of their energy consumption," in Proceedings of the The 4th International Conference on Sustainability in Energy and Buildings, ser. KES series in Smart Innovation, Systems and Technologies, vol. 22. Springer Verlag, 2013, pp. 147–156.

[29] E. Tufte, "Slopegraphs for comparing gradients: Slopegraph theory and practice," http://www.edwardtufte.com/bboard/q-and-a-fetch-msg?msg_id=0003nk&topic_id=1, [retrieved: February, 2014].

[30] C. Park, "Edward Tufte's "Slopegraphs"," http://charliepark.org/slopegraphs, [retrieved: February, 2014].

[31] B. J. Fogg, "Captology: The study of computers as persuasive technologies," in CHI '97 Extended Abstracts on Human Factors in Computing Systems: Looking to the Future, ser. CHI EA '97, March 1997, p. 129.

[32] ——, Persuasive Technology: Using Computers to Change What We Think and Do. Morgan Kaufmann, 2003.

[33] P. Hartman, J. P. Bezos, S. Kaphan, and J. Spiegel, U.S. Patent 5960411: Method and system for placing a purchase order via a communications network. Amazon.com, Inc., September 1999.

Smart Houses for Energy Efficiency and Carbon Dioxide Emission Reduction

Jean-Nicolas Louis, Antonio Caló, Eva Pongrácz
Thule Institute, NorTech Oulu
University of Oulu
Oulu, Finland
e-mails: jean-nicolas.louis@oulu.fi, antonio.calo@oulu.fi, eva.pongracz@oulu.fi

***Abstract*—Understanding the energy consumption of households is a cornerstone for improving residential energy efficiency. In addition, the CO_2 emission profile of energy consumption must be fully understood, to achieve the decarbonisation of energy sector in Europe. Smart houses incorporated into smart grids allow the survey and control of household energy consumption. In this article, the electricity consumption and its related CO_2 emissions are studied for a typical Finnish household. A model detached house is used to simulate the effect of home automation, designed to optimize energy usage, on the CO_2 emissions of this household. Hourly electricity production data are used with an hourly electricity consumption profile generated using fuzzy logic. CO_2 emissions were obtained from the monthly and weekly electricity generated data. The CO_2 emissions related to the use of electric appliances represent around 543 kg_{CO2}/y per dwelling when considering the electricity generated only, and 335 kg_{CO2}/y when balancing the emissions with the exported and imported electricity. Home automation reduced the CO_2 emissions by 13 %. Part of emission reduction was achieved through peak shifting, by moving energy consumption load from daytime to night time. This paper aims at highlighting the role of home automation in reducing CO_2 emissions of the residential sector in the context of smart grid development.**

Keywords-CO_2 emissions; Home Automation; Load shifting; modelling;

I. INTRODUCTION

In December 2013, the European Commission has set clear goals in its Energy Roadmap 2050 (COM(2011)885/2), to achieve a decarbonised society. Decarbonisation in this context means reducing greenhouse gas emissions to 80 % - 95 % below 1990 levels by 2050. This will provide considerable challenges for electricity production, consumption and management. Smart grids represent one tool for achieving this target. Smart grids aim at increasing the energy efficiency of the network, peak load shaving, load shifting, and reduction of energy consumption. Smart buildings are expected to be an integral part of smart grids, with smart meters as the gateway allowing the entrance of smartness into the building. Smart meters receive and send information to and from the building for use such as in Home Area Networks, and Grid handling.

The massive deployment of smart meters under way in Europe facilitates digital measurements, and will allow a consequent access to energy consumption data to energy companies and authorities. European Union (EU) Member States have the obligation of implementing smart meters covering 80 % of consumers by 2020 at the latest [1]. In contrast to the European Energy Efficiency Directive (2012/27/EU) [1], the Finnish Electricity Market law 588/2013 and its application Act 2009/66 on the electricity supply in the survey and measurement sets the deadline of 2014 [2]. Legal obligations to increase energy efficiency also provide a motivation to the deployment of renewable energy sources (RES) as a vector for energy production, both electrical and heating, in a large scale as well as in the buildings. Home energy management can have a significant role in contributing to energy efficiency and cutting or shifting peak load. This can be achieved through an active collaboration of energy consuming systems and the information network on a local level [3]. Putting together the different factors, smart grids, smart building, RES-based heat and electricity and energy efficiency, involve the development of a smart energy network (SEN), capable of managing the energy system through constant monitoring.

The impact of energy efficiency on emissions from the residential sector has been a subject of much research [4]. It has been shown that electric load shifting from the residential sector may reduce air pollution in urban areas [5]. To this effect, developing mathematical tools that are able to anticipate and cut emissions through the deployment of smart systems and home automation is of major importance.

This article aims at exploring the significance of home automation and its impact on the carbon emissions of dwelling, and the possible ways home automation can contribute to decarbonisation. In the first section of the paper, a description of the CO_2 emissions from the production and the use of electricity in Finland will be presented. The second section presents the methodology used for translating hourly carbon emissions to single households will be described. The third section shows and details the results from the simulations carried out on two chosen type of dwelling which will be described and analysed.

II. RELATED RESEARCHES

Researches on smart houses and their development have been going on for quite some times. Smart homes can be broadly seen as a building monitored and controlled for multiple purposes [6]. The energy management feature taken in charge by the smart home is one aspect that has

 ISBN: 978-1-61208-332-2

been developed. Algorithms for generating electricity consumption load profile have been developed on an hourly and half-hourly basis [7], but also with a finer grid on a minute-basis [8]. These algorithms can be further used to emphasise the potential of energy smart houses and their roles in improving the energy efficiency, reducing the energy consumption, and reducing the carbon emissions from the energy used. More elaborated algorithms have been developed where the integration of each appliance within the dwelling have been modelled [9], [10]. Finally, the management of appliances within the dwelling may as well be implemented in simulation for optimizing their usage and enhancing demand-side management [11], [12].

Previous studies have attempted to measure the impact of energy efficiency measures on the CO_2 emissions from the residential sector [4]. Detailed algorithms for evaluating the CO_2 emissions from appliances usage have been proposed [13]. One of the main drawbacks of the previous methods is that the carbon emissions are based on a fixed coefficient, thus limiting the understanding of the CO_2 emissions mechanism. A more dynamic model has been elaborated for estimating the CO_2 emissions and their impact on demand response [14]. Although the last research has based its dynamism on real dataset of energy production on an hourly basis for various countries, the CO_2 emissions related to the production of electricity is based on the IEA annual report on CO_2 emissions [15]. Consequently, studies on segmented electricity production, related CO_2 emissions, and the impact of home automation on the emissions are lacking.

III. Electricity consumption and Carbon Emissions in Finland

The role of the residential sector in reducing carbon emissions is paramount in the development of the future smart grid [16]. In terms of CO_2 emission reduction in the residential sector, the largest effort should be made in retrofitting buildings. The average renewal time of the residential sector is estimated to be around 70 years [4]. The influence of technology on CO_2 emissions needs to be highlighted. Consequently, technology upgrading can greatly influence the total CO_2 emissions of the residential sector. In Finland, lighting consumes over 30 % of the total electricity used in households for appliances [17]. The upgrade of lighting technology is one way for impacting energy consumption [9], but also for reducing carbon emissions [18]. Furthermore, home energy management systems will continue to play a role for increasing energy efficiency, reducing energy consumption [19] and allow load shifting.

Electricity generation and consumption is being constantly surveyed, recorded and reported by the Official Statistics of Finland. In 2012, household appliances consumed 8 072 GWh of electricity [17]. At the same time, 2 579 781 households were registered in Finland [20], resulting in an average consumption of 3 129 $kWh/y.dw^{-1}$. There can be considerable deviation from this average value, if the households is in an apartment building or a detached house [21]. Furthermore, the total electricity production in Finland was around 67.7 TWh in 2012, while the total consumption of electricity was around 82.9 TWh, and a total of 8.4 $MtCO_2$ were emitted. Therefore, it can be estimated that the share of electricity using devices in the total CO_2 emissions from electricity production and use are 1001 tCO_2/GWh_{pro} or 817 tCO_2/GWh_{cons}.

IV. Methodology

A. Electricity consumption profile

The house electricity demand profile is drawn on an hourly basis using different components for evaluating the electricity consumption from appliances, without primary and secondary electric heating systems. Two dwellings will be studied: one having home automation and one without home automation and the difference in their CO_2 balance will be evaluated.

The modelled house contains twenty-one appliances, all of them labelled A or B [9]. The house being in Finland, one of the appliances is an electric sauna stove. The sauna stove used for the modelled house was set to be 6 kW. The overall electricity consumption of appliances in this modelled house is 4 501 kWh/y, which correlates with the findings of the European ODYSSEE MURE project and that of the Sähkötohtori Analysis [21]. The measured data were obtained from detached houses in Oulu, Finland, which were equipped with a 10 kW sauna stove.

B. Hourly electricity generation and emissions

Data acquisition consisted of analysing the electricity generation of all power plants in Finland and the categories of power plants on an hourly basis. Secondly, the carbon dioxide emissions associated with the aforementioned categories were calculated on an hourly basis. Monthly CO_2 emissions are available from July 2007 to October 2013 [22]. It is then possible to evaluate the CO_2 emissions on an hourly basis by associating both elements, the primary energy source for electricity generation and the associated monthly CO_2 emissions.

The electricity generated in Finland on an hourly basis is reported by the Finnish Transmission Service Operator – Fingrid since 2004 [23]. The data is split into two groups: the electricity generated from the power plants and the electrical load on the network taking into consideration the import and export of electricity. Moreover, the Finnish Industry Association (Energiateollisuus) recorded the weekly electricity generated from 1990 [24], which has been broken down between the energy technology used for producing the electricity: wind, hydropower, nuclear, CHP Industry, CHP district heating, conventional and gas turbine power plant. Finally, Fingrid informs in real-time the state of the network using the same categories as mentioned above. Thus, for building up the hourly electricity generation by categories for the year 2012, the weekly average electricity production by category is used in parallel with the hourly electricity generated countrywide. The exported electricity is

 ISBN: 978-1-61208-332-2

considered in the electricity generated and in the corresponding CO_2 emissions. The imported electricity is considered as a share of the CO_2 emissions from electricity consumption in Finland. In order to include the imported electricity into the overall emissions from electricity consumption in Finland, it is necessary to know the energy mix for producing the electricity of the country from which Finland is importing. The hourly electricity generated from a particular energy source in the primary country is evaluated using (1). The notation *h, w*, and *m* designate the hourly, weekly, and monthly time step respectively, *i* is the energy technology used for producing the electricity, and *tot* stands for the total amount of a unit countrywide.

$$P_{h,i} = \frac{P_{w,i}}{P_{w,tot}} \cdot P_{h,tot} \qquad (1)$$

Where $P_{h,i}$ is the electric energy generated by a given technology per hour [MWh/h], $P_{w,i}$ is the electric energy generated on a weekly basis by a given technology [MWh/w], $P_{w,tot}$ is the total amount of electricity produced in Finland per week [MWh/w], and $P_{h,tot}$ is the total amount of electricity produced per hour [MWh/h].

Once the hourly electricity generated by technology has been defined, it is possible to evaluate the hourly emissions from the power plant park.

C. The CO_2 emissions from power plants

As part of the its legal obligation, Finland is publishing the CO_2 emissions from power plants, and every energy intensive industry reports its expected and measured CO_2 emissions for each site [25]. The Finnish Industry Association estimates monthly specific emissions related to electricity production, based on the type of fuel used by the energy industry [24]. By knowing the hourly electricity production from each sector, it results in estimating the CO_2 emissions for each hour countrywide using (2) to (5).

$$E_{w,i} = a \cdot \left(\frac{P_{w,i}}{P_{m,i}} \cdot \frac{\delta_m}{7} \right) \qquad (2)$$

Where, *a* is evaluated using (3) if the full week is within the same month *n*, or (4) if the full week is between two months, *n* and *n+1*.

$$a = \frac{7E_{m,n}}{\delta_m} \qquad (3)$$

$$a = \left(\delta_w \cdot \frac{E_{m,n}}{\delta_m} \right) + \left(\frac{E_{m,n+1}}{\delta_m} \cdot (7 - \delta_w) \right) \qquad (4)$$

Thus, the hourly emission is given by,

$$E_{h,i-gen} = \frac{P_{h,i}}{P_{w,i}} \cdot E_{w,tot} \qquad (5)$$

Where $E_{h,i-gen}$ is the total emissions from the electricity generated hourly [ktCO_2/h], and $E_{w,tot}$ is the total weekly emissions for all power plants [ktCO_2/m], δ_w is the day number within a week where Monday is 1 and Sunday is 7, δ_m is the number of days in the studied month, $E_{m,n}$ is the monthly CO_2 emissions for the month *n*. Fig. 1 (a) illustrates the energy generated and its corresponding CO_2 emissions on an hourly basis for the year 2012 in Finland. It can be noticed that, although there is a strong correlation of CO_2 emissions to electricity generation, emissions may decrease even though the energy generation increases, due to the fact the energy mix is changing Fig. 1 (b).

The emissions due to the electricity imported are further implemented to the primary emissions from the electricity generated within the country. The CO_2 emissions from the electricity generated dedicated to the export is further subtracted from the hourly emissions $E_{h,c1}$. In order to account the total CO_2 emissions from the electricity load in the country, the emissions from each country with which Finland is trading electricity are evaluated, meaning Norway, Sweden, Russia and Estonia. As the hourly energy mix is not known for each country, a general coefficient of CO_2 emissions has been considered for each country named

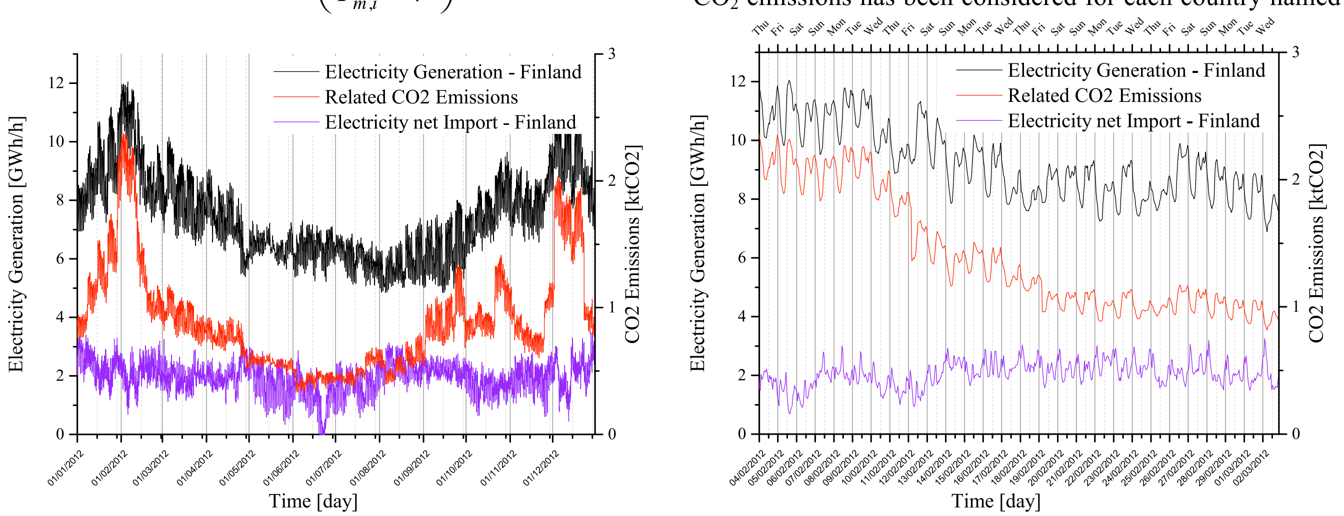

(a) (b)

Figure 1. Hourly electricity generation, net import and their related CO_2 emissions in (a). 2012, and (b). from 4.02-02.03.2012.

previously, respectively 30, 17, 384 and 1 014 $kgCO_2/MWh_{pro}$ [15].

The share of CO_2 emissions coming from each trading country is evaluated using (6).

$$E_{h,c_n} = \sum_{n=2}^{n} \frac{P_{h,net-c_n}}{P_{h,load}} \cdot E_{c_n} \quad (6)$$

Where, $E_{h,cn}$ is the hourly emissions for each participating country to the electricity trade [$kgCO_2/h$], $P_{h,load}$ is the hourly electric load on the Finnish network [MWh/h], $P_{h,net-cn}$ is the net balance of electricity traded between Finland and the country n [MWh/h] in case of export or the difference between the electricity generated and the electricity exported in the case of Finland, and E_{cn} is the coefficient of CO_2 emissions for the corresponding country [$kgCO_2/MWh$]. In case $P_{h,net-cn}$ is negative, the coefficient of CO_2 emissions is equal to $E_{h,i-gen}$ as the emissions from the Finnish production is exported as well, otherwise, E_{cn} takes the value defined by the IEA.

Finally the hourly emissions E_h are determined as the sum of the hourly emissions for each participating country to the electricity trade $E_{h,cn}$ as shown in (7).

$$E_h = \sum_{n=1}^{n} E_{h,c_n} \quad (7)$$

The emission data in Fig. 1 are then translated to a single household where the hourly electricity consumption profile has been previously generated using (8).

$$E_{h,house} = \sum_{j} \frac{P_{j,house}}{P_{h,tot}} \cdot E_h \cdot 10^3 \quad (8)$$

Where $E_{h,house}$ is the hourly emissions from the household [$kgCO_2/h$], and $P_{j,house}$ is the total hourly electricity consumed by the household excluding the electric heating [kWh/h]. Two cases are differentiated: CO_2 levels towards the production of electricity within the primary country, the net CO_2 emissions level considering the import and export as presented in this section. In the first case, P takes the value of the total electricity produced in the primary country $P_{h,tot}$. In the second case, P takes the value of the total load on the electric grid of the primary country $P_{h,load}$.

The results give an estimate of CO_2 emissions related to the electricity consumption in a private household on an hourly basis. This model is then applied to the previously modelled dwelling in order to estimate the daily CO_2 emissions from an average Finnish dwelling.

V. RESULTS AND DISCUSSION

The model showed that the CO_2 emissions are highly dependent on electric consumption levels. Depending on the energy mix for electricity production at a given time, CO_2 emission levels may be lower at peak hours and thus not proportional to consumption levels. Two models have been developed. In the first case, the CO_2 emissions from the dwelling are accounted relatively to the electricity production only. In the second case, the CO_2 emissions from the dwelling are balanced with the electricity exported and imported. Fig. 3 represents the energy consumption for the two modelled dwellings with home automation (Fig. 3 (a)), and without home automation (Fig. 3 (b)). The electricity consumption shown was extracted for a randomly selected week in May 2012, starting on Monday, the 23rd of May.

A. *Case 1: Emissions related to the electricity production*

These dwellings are similar in their characteristics e.g. number and types of appliances, number of inhabitants, dwellings dimensions, or users' habits. The CO_2 emission levels vary from 0.06 to 0.20 $kgCO_2.kWh^{-1}$. The levels depend on the energy mix of Finland's electricity generation. Consequently, the emissions from the dwelling, on an hourly basis, peak at 1.93 $kgCO_2.h^{-1}$ for the dwelling without home automation and 1.81 $kgCO_2.h^{-1}$ for the dwelling with home automation. In the first peak emission case, the related energy demand was 10.03 $kWh.h^{-1}$ and, in the second peak emission case, 9.42 $kWh.h^{-1}$. The maximum electricity consumption in the first case is 12.33 $kWh.h^{-1}$, and 10.16 $kWh.h^{-1}$. The emission peaks are somewhat related to the level of electricity consumption but also to the energy mix for electricity generation at the same time. The use of home automation may reduce the instantaneous peak of CO_2 emissions. The daily electricity profile of the dwellings and CO_2 balance between the two dwellings are represented in Fig. 3 (c).

The difference in the profile of the two modelled dwellings result in a 592 $kWh.y^{-1}$ reduction of total electricity consumption. In terms of CO_2 emissions, the dwelling that is not equipped with a home automation emits 543 $kgCO_2.y^{-1}$, while the house with home automation emit 473 $kgCO_2.y^{-1}$. The amount of CO_2 saved represents 12.78 % of original emissions.

The home automation shifted some of the electricity consumption from the evening peak to the night. It resulted in a decrease of the CO_2 emissions in the evening down to 37 % from the original level, and an increase of 51 % of CO_2 emissions overnight (Fig 3. (d)). Considering, however, that the emissions overnight are about 0.1 $kgCO_2.h^{-1}$ on average,

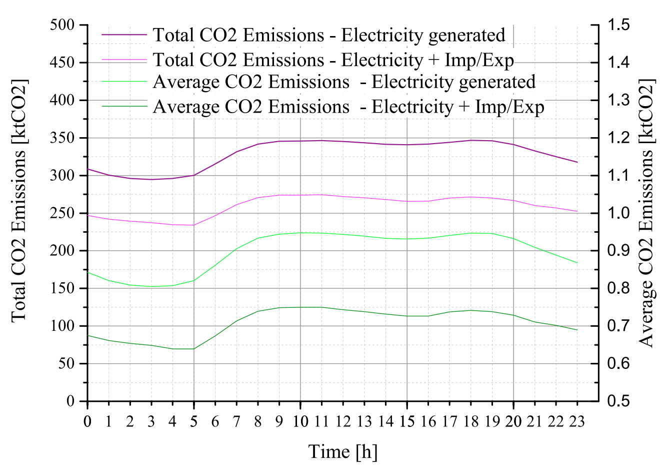

Figure 2. Total and average daily profile of the carbon dioxide emission in 2012, Finland

this can be regarded as relatively small cumulative amount. The emissions increased overnight by 3 to 5 $kgCO_2$, and reduced by 17 $kgCO_2$ on average over the whole year during the evening.

While the home automation was not optimised for reducing CO_2 emissions but for cutting the building peak load consumption, it resulted in the decrease of CO_2 emissions that relates to the electricity consumption. Notwithstanding, it is to be seen that the emissions related to electricity generation countrywide vary throughout the day. Fig. 2 represents the summed CO_2 emissions per hour on the left axis and the hourly average profile of CO_2 emissions on the right axis for the year 2012 from the electricity produced in Finland.

The CO_2 emissions during the peak hours are 0.95 $ktCO_2.h^{-1}$ on average, and correspond to a total of 346 kt_{CO2} between 6 and 7 pm. The lowest point on the daily plot of CO_2 emissions occurs around 2 and 3 am, with an average emission of 0.8 $ktCO_2.h^{-1}$ and a corresponding emission for this particular hour throughout the year of 294 $ktCO_2$.

B. Case 2: Emissions related to the net load

The CO_2 emissions in this 2nd case were found much lower than in the first case. Firstly, the total CO_2 emissions factor $E_{h,i\text{-}gen}$ has slightly decreased. This can be interpreted as an improvement in the global CO_2 emissions from the electric load at the country level. This is explained by the fact that Finland is mostly importing its electricity from Sweden and in second place from Russia. Norway and Estonia represents a small share of the electricity trade. On the one hand, Sweden has an average emission factor around 7 times smaller than the one of Finland, and on the other hand, Finland is exporting electricity with a high emission factor. Also, as the emissions from the Finnish electricity has been calculated for every hour, it introduce some peaks of CO_2 emissions while the electricity from the surrounding countries are applied a constant factor, thus bring a bias result. Nevertheless, the exchange of electricity is beneficial for Finland in terms of CO_2 emissions. The same dwellings that the one mentioned in case 1 were found to have 335 $kgCO_2.y^{-1}$ in case the home automation was not simulated, and 293 $kgCO_2.y^{-1}$ when home automation was running. For both dwellings, the difference between case 1 and 2 is around 38 %. This shows that the CO_2 emissions can be interpreted very differently depending on whether the

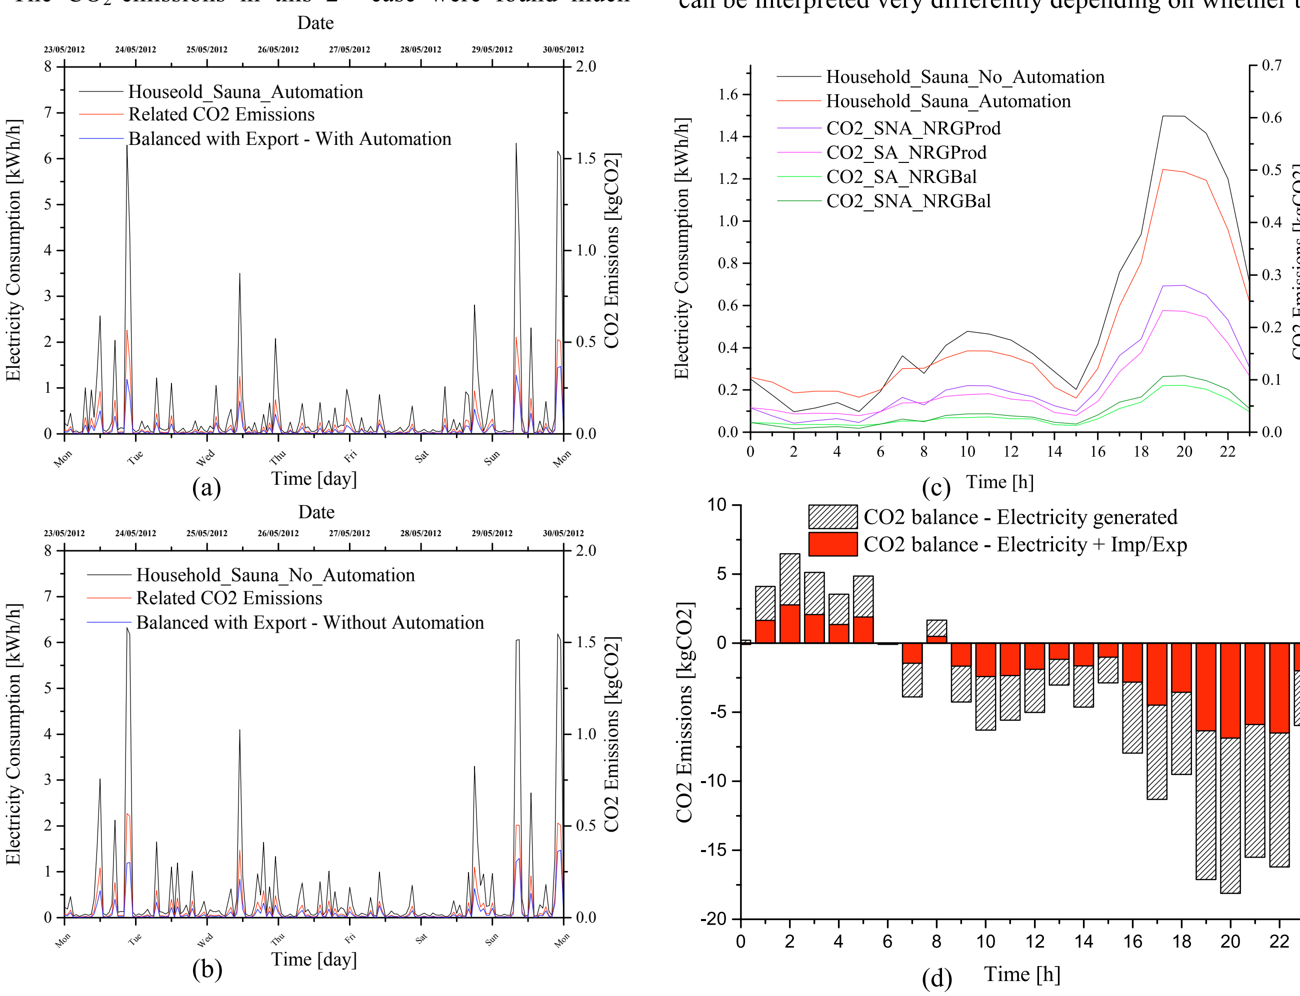

Figure 3. Electricity consumption with its related CO_2 emissions for (a) a dwelling with automation and (b) a dwelling without automation, (c) Daily electricity consumption profiles and (d) its related CO_2 emissions balance between 2 dwellings.

produced electricity, dedicated to the export, is subtracted from the overall electricity consumption of the country or if it should be included into the total CO_2 emissions.

Similarly, the peak of CO_2 emissions due to the electricity consumption in the dwellings are reduced compared to the Case 1 by 24 %. In the Case 2, the peak of CO_2 emissions for the dwelling without home automation reach 1.46 $kgCO_2$/h, and 1.37 in case the home automation is running in the dwelling.

At the system level, the total and average hourly CO_2 emissions have decreased as well. In case the exported and imported electricity are accounted in the total emissions, the low peak occurs between 4-5 am with an average emissions of 0.64 $ktCO_2$ and the high peak period occurs between 10-11 am with an average emissions of 0.75 $ktCO_2$ and cumulates 275 $ktCO_2$ for the same hour.

Regarding the shift of CO_2 emissions due to the home automation device and the feedback strategies used for informing the private consumers, it has decreased by 6 $kgCO_2$ in the evening and has risen by 2.7 $kgCO_2$ in the night time. The quantities of CO_2 shifted, presented in Fig. 3 (d), are different from Case 1 to Case 2 as the CO_2 emission profile for both cases are different, as shown in Fig. 2.

C. Discussion

Both cases showed that load shifting can contribute to 12.7 % decrease in CO_2 emissions. However, there is a difference depending if the balance of import and export is considered.

As well, consumer awareness and their willingness to comply is also a factor in the potential for reducing CO_2 emissions. Table 1 summarises the results from the CO_2 emissions and the electricity consumption from both dwellings.

It is necessary to point out the importance of methods evaluating emissions on the results. It is paramount that the countries involved use the same methodology for their CO_2 evaluation. In this study, Finland is mostly importing electricity from Sweden and Russia and exporting to Norway and Estonia. For Sweden, it results in importing “polluted” electricity and exporting cleaner electricity to Finland. Consequently, for Finland, the shifting of CO_2 emissions is greater when relating the emissions to the gross electricity production. Multi-objective algorithms will need to be developed for optimising electricity consumption and/or CO_2 emissions. In addition, an added level of complexity is if export/import net emissions are considered or not.

TABLE I. CO_2 EMISSIONS SUMMARY FOR THE TWO STUDIED CASES

	CO_2 emissions relative to		
	Electricity produced	***Net electricity consumed***	***Unit***
Min. $E_{h,i\text{-}gen}$	0.06	0.04	$kgCO_2$/kWh
Max. $E_{h,i\text{-}gen}$	0.20	0.19	
Max $E_{h,house}$ SA	1.81	1.37	$kgCO_2$/h
Max $E_{h,house}$ SNA	1.93	1.46	
Max $P_{j,house}$ SA	12.33	12.33	kWh/h
Max $P_{j,house}$ SNA	10.16	10.16	
Total $E_{h,house}$ SA	543	335	$kgCO_2$/a
Total $E_{h,house}$ SNA	473	293	
Max Average $E_{h,i\text{-}gen}$	0.95	0.75	$ktCO_2$
Min Average $E_{h,i\text{-}gen}$	0.8	0.64	
Max Sum $E_{h,i\text{-}gen}$	346	275	
Min Sum $E_{h,i\text{-}gen}$	294	234	

VI. CONCLUSIONS

The article detailed the CO_2 emissions of electricity generation in Finland. Firstly, monthly and weekly data of electricity generation were used to calculate corresponding CO_2 emissions into hourly data. This was used to evaluate the CO_2 emission profile of households. The model was based on hourly electricity load profiles previously built.

Secondly, the CO_2 emissions associated with imported and exported electricity generation were accounted as well. Both cases show the same peak distribution in their daily profile. Notwithstanding, emissions will depend on the fuel used at a particular hour. Therefore, the relationship between electricity production, import and export is not straightforward. The cumulated carbon emission overnight from the electricity produced in Finland stands at around 290 $ktCO_2.h^{-1}$, while the peak reaches 345 $ktCO_2.h^{-1}$. Considering the import and export of electricity, and their related CO_2 emissions, the peak dropped to 230 $ktCO_2.h^{-1}$ overnight, and the high peak at 275 $ktCO_2.h^{-1}$, respectively.

Although the home automation has not been optimised for reducing the CO_2 emissions from the modelled household, the CO_2 emissions from the electricity consumption are somehow proportional to electricity consumption levels. The study showed that home automation might reduce the carbon emission by 12.7 % while influencing the private consumers’ everyday routine. The CO_2 emissions have been reduced most substantially during the evening peak, by 18 $kg_{CO2}/h.y^{-1}$ in the first case and by 6 $kg_{CO2}/h.y^{-1}$ in the second case, while the emissions at night time have increased from 3 to 5 $kg_{CO2}/h.y^{-1}$ on average. Although the CO_2 emissions related to electricity consumption from appliances are strongly correlated, the energy mix for producing this electricity needs to be considered and thus optimised for reducing the carbon footprint of households.

Consequently, smart buildings within a smart grid may not only participate to load shifting, increase energy efficiency or decrease in electricity consumption, but can also contribute significantly to the reduction of CO_2 emissions. It will, in turn, impact the total CO_2 emissions of the country and will assist in achieving the decarbonisation goal of the EU.

This limitation of this research is that there was no information available on the variation of the energy mix from the exporting countries, therefore, the import electricity had to be considered with a yearly constant CO_2 emission factor. Secondly, in the case of Finland, a more detailed estimation would require knowing the energy mix hour-by-hour, rather than estimating it from the monthly average.

Further research will investigate the impact of private consumers in correlation with home automation for reducing the CO_2 emissions of households. In addition, a full assessment considering district-heating systems shall be done, in order to achieve full integration of smart buildings in a smart energy network. Finally, the multi-objective algorithms will have to be further developed.

ACKNOWLEDGEMENT

The authors thank the Thule Institute Doctoral Programme for financing this research.

REFERENCES

[1] Publications Office, Directive 2012/27/EU of the European Parliament and of the Council of 25 October 2012 on energy efficiency, amending Directives 2009/125/EC and 2010/30/EU and repealing Directives 2004/8/EC and 2006/32/EC Text with EEA relevance. 2012, pp. 1–56.

[2] TEM, Valtioneuvoston asetus sähköntoimitusten selvityksestä ja mittauksesta. 2013, pp. 1–6.

[3] D. Long Ha, S. Ploix, M. Jacomino, and M. Hoang Le, “Home energy management problem: towards an optimal and robust solution,” in Energy Management, no. 5, F. Macia Perez, Ed. InTech, 2010, pp. 77–106.

[4] H. Hens, G. Verbeeck, and B. Verdonck, “Impact of energy efficiency measures on the CO_2 emissions in the residential sector, a large scale analysis,” Energy & Buildings, vol. 33, Jan. 2001, pp. 275–281.

[5] N. Gilbraith and S. E. Powers, “Residential demand response reduces air pollutant emissions on peak electricity demand days in New York City,” Energy Policy, vol. 59, no. C, Aug. 2013, pp. 459–469.

[6] L. C. De Silva, C. Morikawa, and I. M. Petra, “State of the art of smart homes,” Engineering Applications of Artificial Intelligence, vol. 25, no. 7, Oct. 2012, pp. 1313–1321.

[7] M. Stokes, “Removing barriers to embedded generation: a fine-grained load model to support low voltage network performance analysis,” Institute of Energy and Sustainable Development de Montfort University, Leicester, 2005.

[8] A. Grandjean, J. Adnot, and G. Binet, “A review and an analysis of the residential electric load curve models,” Renewable and Sustainable Energy Reviews, vol. 16, no. 9, Dec. 2012, pp. 6539–6565.

[9] J.-N. Louis, A. Caló, K. Leiviskä, and E. Pongrácz, “Home Automation for a sustainable living – Modelling a detached house in Northern Finland,” Proceedings of the 7th International Conference on Energy Efficiency in Domestic Appliances and Lighting EEDAL’13, Sep. 2013, pp. 1–11.

[10] N. Arghira, L. Hawarah, S. Ploix, and M. Jacomino, “Prediction of appliances energy use in smart homes,” Energy, vol. 48, no. 1, Dec. 2012, pp. 128–134.

[11] R. Missaoui, H. Joumaa, S. Ploix, and S. Bacha, “Managing Energy Smart Homes according to Energy Prices: Analysis of a Building Energy Management System,” Energy & Buildings, vol. 71, Mar. 2014, pp. 155–167.

[12] P. Chavali, P. Yang, and A. Nehorai, “A Distributed Algorithm of Appliance Scheduling for Home Energy Management System,” IEEE Trans. Smart Grid, vol. 5, no. 1, Jan. 2014, pp. 282–290.

[13] R. Huang, M. Itou, T. Tamura, and J. Ma, “Agents based Approach for Smart Eco-home Environments,” presented at the The 2010 International Joint Conference on Neural Networks (IJCNN), 2010, pp. 1–8.

[14] P. Stoll, N. Brandt, and L. Nordström, “Including dynamic CO_2 intensity with demand response,” Energy Policy, vol. 65, no. C, Feb. 2014, pp. 490–500.

[15] IEA, International Energy Agency, “CO_2 Emissions From Fuel Combustion Highlights 2013,” Oct. 2013, pp. 1–158.

[16] S. J. Darby, J. Stromback, and M. Wilks, “Potential carbon impacts of smart grid development in six European countries,” Energy Efficiency, vol. 6, no. 4, May 2013, pp. 725–739.

[17] Official Statistics Finland, Ed., Energy Consumption in households by energy source in 2012, 2012. [Online]. Available: http://www.stat.fi/til/asen/2012/asen_2012_2013-11-13_tau_002_en.html. [Accessed: 05-Dec-2013].

[18] E. Tetri, A. Sarvaranta, and S. Syri, “Potential of new lighting technologies in reducing household lighting energy use and CO_2 emissions in Finland,” Energy Efficiency, Dec. 2013, pp. 1–12.

[19] M. G. Ippolito, E. R. Sanseverino, and G. Zizzo, “Impact of building automation control systems and technical building management systems on the energy performance class of residential buildings: An Italian case study,” Energy & Buildings, vol. 69, Feb. 2014, pp. 33–40.

[20] Official Statistics Finland, Ed., Household-Dwelling units by number of persons 1960 - 2012. [Online]. Available: http://www.stat.fi/til/asas/2012/asas_2012_2013-05-22_tau_001_en.html. [Accessed: 05-Dec-2013].

[21] V. Rouhiainen, “Decomposing Electricity Use of Finnish Households to Appliance Categories,” presented at the Energy efficiency in domestic appliances and lighting, Proceedings of the 5th international conference EEDAL '09, Berlin, 2010, vol. 2, pp. 438–455.

[22] Finnish Industry Association, “Power generation in Finland - fuels and CO_2-emissions,” Energiateollisuus, Helsinki, Nov. 2013.

[23] Fingrid Oyj, Ed., Load and generation. [Online]. Available: http://www.fingrid.fi/en/electricity-market/load-and-generation/Pages/default.aspx. [Accessed: 04-Dec-2013].

[24] Finnish Industry Association, “Electricity net production, imports and exports, in Finland,” no. 258. Energiateollisuus, Helsinki, 28-Nov-2013.

[25] Energiamarkkinavirasto, “Emission Allowance Balance.” EMV, Helsinki, 28-Nov-2013.

Energy-Efficient Heterogeneous Cluster and Migration

Jukka Kommeri and Tapio Niemi
Helsinki Institute of Physics
CERN, Geneva, Switzerland
Email: jukka.kommeri@cern.ch

***Abstract*—A recent trend of server consolidation using virtualization has allowed datacenters to improve their server utilization rate and total energy efficiency. Virtualization technologies are used to share physical hardware between multiple services. This has led to another challenge: how to place the virtual servers into the physical servers? Especially, this is important in cases in when the workload of virtual servers is not constant.**

In this paper, we study the overhead of virtual server migration in physics computing on energy efficiency with an emphasis on quality of service. Our method is based on the standard migration technique that allows us to move virtual machines between physical machines without significant interference on the service running on the virtual machine.

Our results indicate that by utilizing dynamic resource sharing among the virtual servers and load balancing between heterogeneous physical machines, it is possible to improve energy efficiency of online cloud services.

Keywords-energy-efficiency, virtualization, migration, heterogeneous hardware, physics computing

I. INTRODUCTION

The energy consumption of the information and communications technology (ICT) sector has been rapidly growing for the past decade. This has received a lot of attention and concerns [1], [2]. Much of this is due to over provisioning of hardware to serve the peak loads and provide high service availability. According to many studies, the average utilization rate of a server is around 15% of maximum but it depends much on the service and it can be even as low as 5% [3], [4]. These values are much lower than in the other fields of industry, although peak loads do not only exist in the IT sector.

New technologies, both software and hardware related, have been studied and adopted to face this increase in energy consumption. The two main approaches in this field are 1) energy-proportional hardware, and 2) service consolidation through virtualization. Developing energy-proportional hardware has appeared to be very challenging, especially developing such a technology for the memory has not so far progressed much. On the other hand, just manufacturing computers takes both money and energy, thus keeping servers idle waiting for peak loads is not a profitable approach even with possible future energy-proportional hardware [5]. Instead, server consolidation applying virtualization could solve the problem in some cases. This is already visible, for example Koomey [6] indicates that server virtualization has already decreased the growth in the server market.

Though virtualization is already used to improve the energy

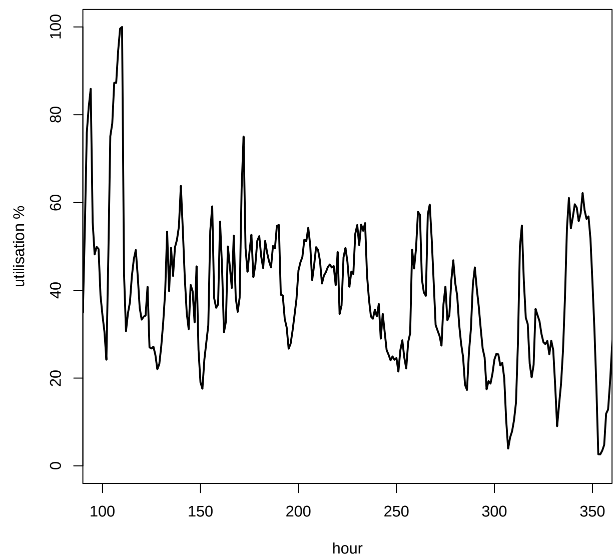

Figure 1. Cluster utilization rate.

efficiency of physical servers, there is still room for improvement. Datacenters need to serve their customers and provide capacity for the peak loads. As our random sample containing a subset of servers in CERN, Conseil Europen pour la Recherche Nuclaire, datacenter log in Figure 1 indicates, the peak loads are surrounded by less busy periods. The cluster needs to adapt for these changes. This can be achieved by dynamically moving virtual machines between physical hosts depending on the load of the virtual and physical machines. There are also big differences among different physical hardware solutions since some of them are better or more suitable for heavy computation while others work more efficiently in light-weight computing.

In this study, we focus on energy efficiency of migration and how it affects the quality of service of a virtualized service. Our hypothesis is that the energy consumption of the system improves when dynamically placing virtual servers to different hardware based on their current load. We validate our hypothesis by studying the energy consumption of virtualized servers with realistic physics analysis software while migrating the virtual servers between physical servers. Our test set up consists of two types of hardware: one low-performance energy-efficient server to host mostly idle or even dormant virtual machines and one powerful server to host heavily

 ISBN: 978-1-61208-332-2

loaded virtual machines. We used an open source virtualization platform Kernel-based Virtual Machine (KVM) and the Compact Muon Solenoid Software framework (CMSSW) [7] as a realistic test case. The same software is used at CERN to analyze collision data produced by the CMS experiment. Our results show that heterogeneous hardware can be used to improve both the energy efficiency of the cluster and optimize the use of floor space.

The paper is organized in the following way. In Section II, we have related work, which is followed by Section III that gives more detailed description on migration. Then, we have our test methodology in Section IV and describe our test environment in Section V. These are followed by test results in Section VI and conclusion in Section VII.

II. Related work

Virtualization itself is not an optimal solution since logical sharing of physical resources with virtualization presents some overhead that depends on both the placement of virtual servers and their workload [8]. This makes the sharing of physical resources between virtual machines challenging, especially, when the load of virtual servers is not constant. A lot of research has been put into virtual machine scheduling, the way how virtual machines are placed in a cluster of physical machines. Many different variations of bin-packing algorithms have been proposed [9], [10], [11]. Some of these algorithms also do load balancing dynamically by moving virtual machines between physical resources. This technique is called migration. Migration gives more freedom in choosing the physical hardware as the virtual machine is not bound to single hardware. It also gives flexibility when managing physical resources.

Some studies already point that exploiting heterogeneous hardware in a virtualized cluster can be beneficial. Hirofuchi et al. [12] use dedicated servers for virtual machines with less load and other dedicated servers for running heavily loaded virtual machines. Virtual machines are moved between these two types of physical servers in function of their processing needs. In their work, the only difference between virtual machine pool server and execution server is the amount of memory. Profiting from energy-efficient hardware was also studied by Verma et al. [9], but in their study the implementation was left at the level of power models where different hardware had different power models. Also, the preference of allocating virtual machines to more energy-efficient hardware has been studied [10], [13].

We extend this thinking by using an energy-efficient server for the virtual machine pool server and a power-efficient server as an executing node. As we have found in our earlier work [8], virtual machines do not consume physical resources when they are idle. This makes it possible to store as many idle virtual machines on a pool server as one can fit in its memory. As the idle servers do not need much computing power, one can choose a more energy-efficient hardware for that purpose. On the execution server one does not need that much memory, but more computing power per virtual machine.

III. Migration

Migration is a technique used to move virtual machines from one physical server to another. In practice, this means that the hard disk, memory contents and processor state is moved from one physical server to another physical server. There are several ways to perform a migration. Choosing the proper one depends much on the way computing environment is set up and what the current use case is. Migration can be done either online or offline. In offline migration, the virtual machine is shutdown for the duration of the state and memory transfer as in online migration this is minimized and the virtual machine does not experience a notable down time. Online migration is also known as live migration. For live migration to work properly, one has to set up a shared virtual machine disk image or a root partition in both physical servers. This can be achieved with various network-attached storage (NAS) services; NFS, iScsi, AoE, Ceph, etc. Having a NAS generally speeds up the migration process. Without a dedicated NAS one would also have to copy the contents of the virtual machine disk image from the original host to the destination host. This takes much more time than just moving the contents of the memory of the virtual machine to a new location.

When making a migration decision, it is good to evaluate the possible parameters that affect migration and also how migration affects other virtual machines or applications running on the servers. Liu et al. [14] have studied how different parameters affect the migration time and energy consumption. They have come up with a model for predicting migration time, downtime and energy consumption. This model is defined by Equations (1,2,3). Most important parameters are memory dirtying rate (D), rate of transmission (R) and virtual machines memory size (V_{mem}). If rate of transmission, i.e., network bandwidth is smaller than memory dirtying rate, live migration is not possible. From these three it is possible to calculate migration time (T_{mig}).

$$\lambda = D/R \tag{1}$$

$$n = log_{\lambda} \frac{V_{thd}}{V_{mem}} \tag{2}$$

$$T_{mig} = \sum_{i=0}^{n} T_i = (V_{mem}/R) \times \frac{1-\lambda^{(n+1)}}{1-\lambda} \tag{3}$$

These equations are mainly indicative as many parameters vary during the migration process and due to the choice of the migration type. Usually both V_{mem} and V_{thd} are configured statically. For the live-migration to succeed λ needs to be bigger than than zero. Equation (3) is valid for both online and offline migration where T_0 is the time it takes to migrate the initial memory set. Equation (3) shows the iterative nature of migration process where n is the number of iterations. Configuring V_{thd} can have a big impact on T_{mig} as then the dirtying rate has a bigger impact [15].

As mentioned earlier, virtualization makes it possible to control the load of physical resources by migration of virtual machines. Cluster management systems are based on heuristics that use data gathered from both physical machines and virtual machines. In the simplest form, they make their decision based on how physical CPUs are loaded while the complex ones take into consideration also other issues such as memory usage, network traffic, service level agreements (SLA), server energy consumption, server thermal state, virtual machine intercommunication, etc. [11]. Even the simplest algorithm improves energy efficiency of a cluster that has a varying workload.

Development of the most optimal algorithms have gotten a lot of attention and several different solutions have been proposed. In this section, we explore a few of them. These algorithms vary much in complexity and as there exist no standard way for testing, their comparison is difficult. Also, the fact that most algorithms have been tested with different simulators does not make the comparison any easier. As Srikantaiah et al. [16] write, development of an algorithm is a compromise between time and performance. The algorithms that produce the most optimal results tend to take longer to calculate, e.g., the constraint programming algorithm by Hermier et al. [17]. But even the simplest algorithms such as the one by Shrikantaiah [16], that picks the best server by computing a simple Euclidean distance of the addition of new workload, can improve cluster energy-efficiency as then the system is reacting to the change of resource requirements. However, this Euclidean distance algorithm like many other algorithms, assumes that the requirements of the new task are known. The basic idea of all these management systems is to minimize the use of physical machines and this is done by packing already loaded servers more efficiently.

IV. METHODOLOGY

In many services, the workload fluctuates a lot as a function of time. For example, the workload can be near zero during nights and weekends. This is very much the case in our data as Figure 1 illustrates. During off-peak periods, the request rate is just around 1/5 of the peak periods, thus also less computing power is needed during off-peak periods. Based on this observation our hypothesis is:

> Energy efficiency can be improved without decreasing the quality of service by moving virtual servers to an energy-efficient low power server during off-peak periods.

Naturally, we assume that the energy-efficient server is not powerful enough to host virtual servers during the peak hours. Additionally, to simplify our theoretical model, we assume that all virtual servers are always hosted by the same server. Now, we can form our model as follows: Let the idle power of the energy-efficient server A be A_{idle} and full power A_{peak} and for the power server B, B_{idle} and B_{peak}, respectively. Now the upper limit for the energy consumption of running the system x hours will be:

$$E = A_{peak} \times x \times p_A + A_{idle} \times x \times (1 - p_A) + B_{peak} \times x \times (1 - p_A) + B_{idle} \times x \times p_A + n \times E_m$$

Where p_A is the portion of time that the server is hosted on the server A, n is the number of migrations, and E_m is the extra cost of one migration operation.

In our test setting (see Section VI), the values for the parameters are as follows: $A_{idle} = 25W$, $A_{peak} = 78W$, $B_{idle} = 101W$, $B_{peak} = 246W$, and $E_m = 1Wh$.

We compare two cases: 1) the virtual service is hosted all the time on Server B and we do not have Server A, and 2) during off-peak periods, the service is migrated to Server A from Server B. The energy consumption can be computed as follows:

Case 1: $E_1 = B_{average} \times x$
Case 2: $E_2 = A_{peak} \times x \times p_A + B_{average} \times x \times (1 - p_A) + n \times E_m$

We assume that the load level is high enough to keep Server A almost in its peak power all the time when the service is run on it. Further, we assume that even off-peak load is still high enough to make Server B to consume 75% of its peak power. These assumptions are reasonable as illustrated in Figure 7. It is quite straightforward to see that Case 2 is more energy-efficient if the workload is not very high. Upper and lower bound power values of the servers A and B are measured from real hardware. An estimate of the cost of migration is determined by collecting a large sample of migration results while having realistic workload on the virtual machine.

V. TEST ENVIRONMENT

Our tests aimed at measuring the energy consumption and overhead of virtualization and migration. As a workload in virtual machines we had real physics analysis applications, that are used at CERN, and data that is produced at CERN particle accelerator, Large Hadron Collider (LHC). We measured how performance was affected by migration.

In our tests, we had separate hardware for basic migration test and hardware comparison tests. In the migration tests, we wanted to have identical servers so that the effect of migration could be measured. In the migration test, we used Dell 210 single processor servers with Intel X3430 processors and 12GB of memory as hypervisors. In the hardware comparison tests, we had a Dell 210 II server with energy-efficient Intel E31260L 2.4 GHz processor and a Dell 415 server with a high efficiency AMD Opteron 4276HE 2.6 GHz processor and 24GB of memory. Servers were connected with a gigabyte network. Power usage data of the servers was collected with a Watts up? PRO meter via a USB cable as it is illustrated in Figure 2. Power usage values were recorded every second.

In migration tests, the hypervisors were connected with a secondary network, which was dedicated to migration traffic. Also, the hypervisors were configured to use all the bandwidth of the migration network. In both test cases the virtual machine

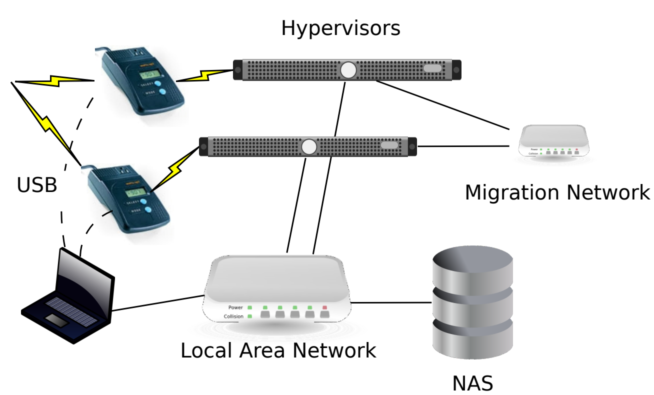

Figure 2. Illustration of test setup.

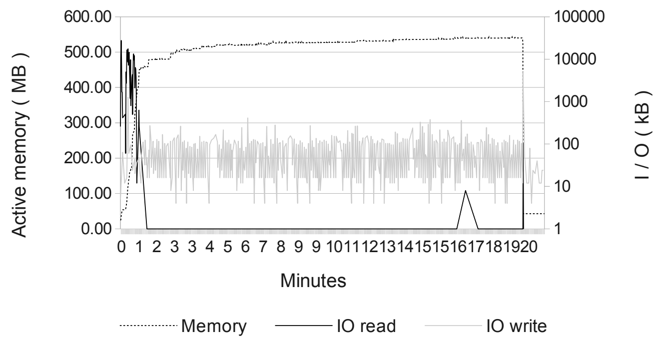

Figure 3. CMMSW task memory and IO curves.

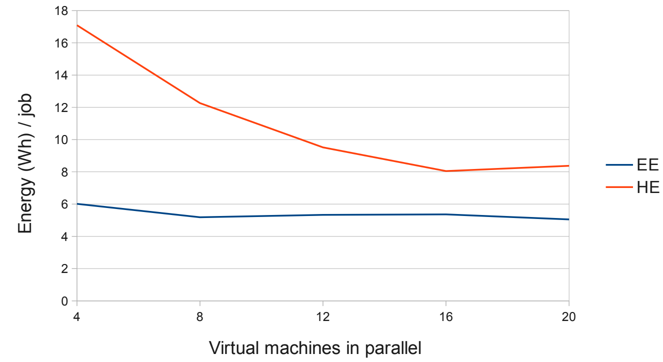

Figure 4. Energy consumption per physics analysis job when running them in parallel and in dedicated virtual machines.

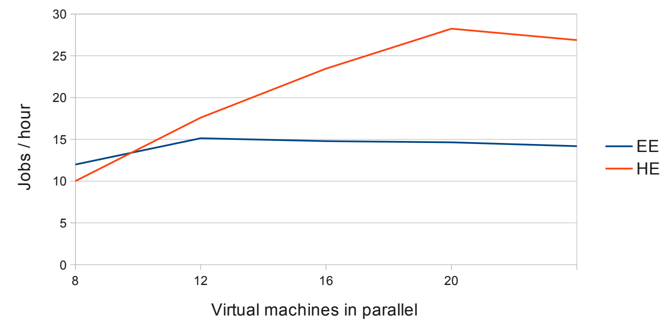

Figure 5. Server throughput when physics jobs in parallel and in dedicated virtual machines.

images, hard drives, are on a powerful remote server HP Proliant DL585 G7. A separate network storage server, Dell PowerEdge T710, was used to serve physics analysis software and data. Both network storage servers were connected to the local area network, which is separate of the migration network.

A. *Test Applications*

The operating system used in all machines, virtual or real, was a standard installation of 64 bit Linux, Ubuntu. Physical servers were installed with Ubuntu version 13.04 as the virtual machines had a version 10.04 of Ubuntu. KVM was used as virtualization technology. Virtual machine images were shared over iSCSI from a network attached storage.

Virtual machines were installed with a separate root that contained Scientific Linux at Cern 5 (SLC5) installation and CMSSW version 4.2.4. From this separate root, CMSSW was run using chroot system call. The separate root was chosen as the SLC5 was not as performant Ubuntu as a virtual machine operating system. For the CMSSW tests real data files produced by the CMS experiment were used. These data files were shared to the virtual machines over a network file system, NFSv4. Execution of the analysis was sped up by limiting the number of analyzed events to 300. Depending on the hardware the execution of single analysis take from 10 to 19 minutes.

Behavior of our workload is shown in Figure 3. These statistics were collected with a Linux tool called Vmstat and measured inside the virtual machine, which was running a CMSSW analysis. The CPU curve was left out as it was almost a straight line from the beginning.

B. *Test cases*

We divide our tests into two categories. In the first part, we conducted a parallelism test on two different type of servers and in the second part we conducted migration overhead tests with varying workload. In both test cases, the workload used was the same as in previous section. The hardware for the migration test was different as for this test a homogeneous environment is more suitable. Migration overhead was tested by doing the migration at different points of the execution of the analysis software. As the execution on hardware chosen for migration test took about 19 minutes, the migration points 1,5,9,13 and 17 minutes were chosen. Only one migration per one analysis task was run and the test was repeated ten times for all combinations. Also, the same tests were repeated with different background loads. Higher background load was achieved by starting three extra virtual machines on both hypervisors and running two CMS analysis on each one. Making the total load per hypervisor six jobs in addition to the virtual machine that is being migrated.

In the hardware comparison tests, we run the physics analysis task on a different number of virtual machines each running a single analysis task. The same tests with varying load were performed on both types of hardware, the energy-efficient server (EE) and the high efficiency server (HE). Here the energy efficiency means that the server can do more work with less energy as the high efficiency means that the server is able to more work in less time. Test servers were initialized with virtual machines and then the workload was started remotely on each virtual machine at the same time and test would capture the duration from the termination of the last workload. As with all our tests, between every test the servers were rebooted to reset the test environment.

VI. Test Results

At first we measured how different our energy-efficient server is from the high efficient one. Figure 4 illustrates how energy consumption per job changes when parallelism is increased. The two CPU Opteron server never reaches the energy-efficiency of the single CPU server, but its energy efficiency improves when running on higher load. In Figure 5, we have the total throughput of the physical servers with different loads. It shows how the 2 CPU server can handle much more load than the energy-efficient server.

 ISBN: 978-1-61208-332-2

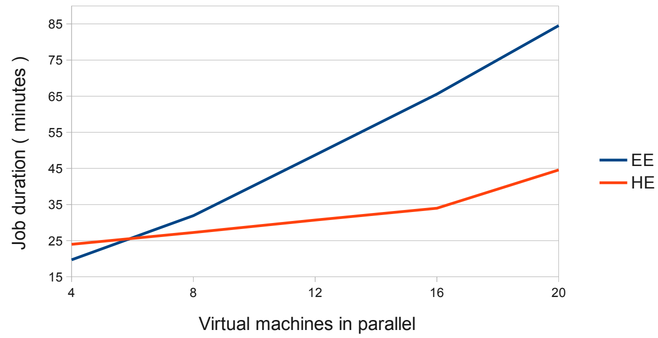

Figure 6. Job duration in minutes with different number of virtual machines running single job each.

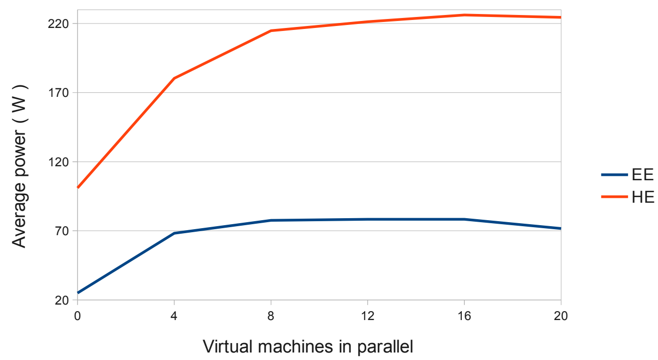

Figure 7. Average power in Watts with different number of virtual machines running single job each.

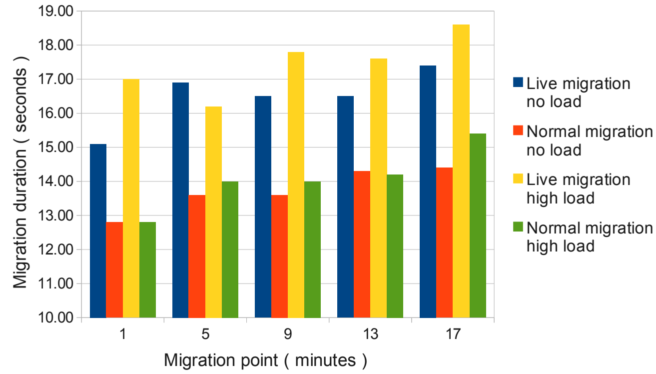

Figure 8. Duration of one migration on different parts of the CMS analysis task.

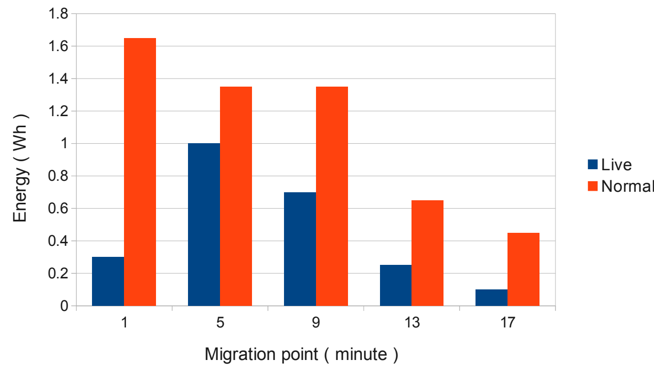

Figure 9. Energy overhead of one migration.

Another way of measuring performance is quality of service. Figure 6 illustrates better how the quality of service deteriorate when more parallelism is added. The cores of the energy-efficient server are more efficient. They are better both in efficiency as in energy efficiency. As the load increases the higher core count of the high efficient server provides better performance.

There are big differences in energy consumption of different hardware types. Figure 6 shows that the energy consumption of an idle single CPU energy-efficient server can be almost three time that of the high efficient two CPU server. Also, the power range of the energy-efficient server is narrower. The two processor server requires some load to become more energy-efficient. In our case more than eight parallel jobs.

Finally, we measured how big an impact one migration has on the CMS analysis task. Migration time was measured at different parts of the CMS analysis task. In Figure 8, we have the results of these migration measurements. As described earlier, the migration time depends much on the variation of load in the virtual machine. The duration of a normal migration rises steadily as the virtual machine expands its memory space. In all tests, the virtual machine runs one CMS job and then it shuts down. Virtual machine accumulates memory contents and as it has plenty of memory, it is never released.

Although the migration times of live migrations are longer, the energy consumption can still be lower. Figure 9 illustrates the difference of energy consumption overhead of live and normal migration. Although the difference between the two types is not significant, there are still benefits in choosing the correct one. The benefit of live migration is that it can continue the execution even though its been transferred from host to another. Thus, even though the migration lasts longer its energy overhead due to the continuous execution is less. On the other hand, if we double the load of the virtual machine, live migration will not succeed. Normal migration will succeed regardless of the load. We can also see that the effect of the background load was minimal.

VII. Conclusions and Future Work

We measured the energy consumption of high-energy physics analysis workload on two types of hardware. Results indicate that energy efficiency of a cluster can be improved by using heterogeneous hardware: it lowers the idle consumption, but does not reduce peak performance. Energy-efficient servers are good for lower load and storing, or parking, idle virtual machines, while high-efficient servers are needed to serve the peak loads and they can even be turned off when there is less load. Yet, high-efficient servers cannot be completely replaced by energy-efficient servers as the increased number of servers would consume more floor space. Migration is the key to make the management of virtual machines possible. Since the use of migration does not impose too much overhead, it is a suitable tool to manage virtual machines in large clusters and optimize the load between heterogeneous hardware.

In our tests, the choice of hardware was not optimal since the test workload was optimized for Xeon hardware. Having an up to date two processor Xeon server might have given better performance than our AMD processor. Moreover, the energy-efficient server proved to be almost too efficient for hosting idle virtual servers. In the case of the energy-efficient server, the ability to host the total memory space of the virtual machines is more important than the computing power. The memory requirement could even be optimized automatically by using the virtual machine memory size management that allows to over commit physical memory dynamically. One conclusion we made, is that when choosing hardware for a computing cluster, one could optimize both investment and running cost by purchasing task specific hardware, since there is a significant price difference between energy-efficient and high-efficiency hardware.

Our future work will focus on building an automatic system for optimizing the placement of virtual servers. Our aim is to use, e.g., CPU load values of the virtual server as an indicator whether do the migration or not in the following way: if the virtual server is running on the energy-efficient server and its short time load value goes over the threshold, it will be moved to the powerful server. Since migration is still a relatively expensive operation, it should not be performed too frequently. Thus, the virtual server is only moved back to the energy efficient server, if its long time average load value goes back under the given threshold.

REFERENCES

[1] B. Schäppi, F. Bellosa, B. Przywara, T. Bogner, S. Weeren, and A. Anglade, "Energy efficient servers in europe," Austrian Energy Agency, Tech. Rep. October, 2007.

[2] E. STAR, "Report to congress on server and data center energy efficiency," U.S. Environmental Protection Agency ENERGY STAR Program, Tech. Rep., 2007.

[3] L. A. Barroso and U. Hölzle, "The case for energy-proportional computing," Computer, vol. 40, 2007, pp. 33–37.

[4] W. Vogels, "Beyond server consolidation," Queue, vol. 6, January 2008, pp. 20–26.

[5] H. Chung-Hsing and S. Poole, "Revisiting server energy proportionality," in Parallel Processing (ICPP), 2013 42nd International Conference on, Oct 2013, pp. 834–840.

[6] J. Koomey, "Growth in data center electricity use 2005 to 2010," Oakland, CA: Analytics Press, 2011.

[7] F. Fabozzi, C. Jones, B. Hegner, and L. Lista, "Physics analysis tools for the cms experiment at lhc," Nuclear Science, IEEE Transactions on, vol. 55, 2008, pp. 3539–3543.

[8] J. Kommeri, T. Niemi, and O. Helin, "Energy efficiency of server virtualization," International Journal On Advances in Intelligent Systems, v 5 n 3&4, 2012.

[9] A. Verma, P. Ahuja, and A. Neogi, "Power-aware dynamic placement of hpc applications," in Proceedings of the 22nd annual international conference on Supercomputing, ser. ICS '08. New York, NY, USA: ACM, 2008, pp. 175–184.

[10] B. Li, J. Li, J. Huai, T. Wo, Q. Li, and L. Zhong, "Enacloud: An energy-saving application live placement approach for cloud computing environments," in Cloud Computing, 2009. CLOUD '09. IEEE International Conference on, 2009, pp. 17–24.

[11] A. Beloglazov and R. Buyya, "Energy efficient allocation of virtual machines in cloud data centers," in Cluster, Cloud and Grid Computing (CCGrid), 2010 10th IEEE/ACM International Conference on, 2010, pp. 577–578.

[12] T. Hirofuchi, H. Nakada, S. Itoh, and S. Sekiguchi, "Reactive cloud: Consolidating virtual machines with postcopy live migration," Information and Media Technologies, vol. 7, no. 2, 2012, pp. 614–626.

[13] R. Nathuji and K. Schwan, "Virtualpower: coordinated power management in virtualized enterprise systems," SIGOPS Oper. Syst. Rev., vol. 41, no. 6, Oct. 2007, pp. 265–278. [Online]. Available: http://doi.acm.org/10.1145/1323293.1294287

[14] H. Liu, C.-Z. Xu, H. Jin, J. Gong, and X. Liao, "Performance and energy modeling for live migration of virtual machines," in Proceedings of the 20th International Symposium on High Performance Distributed Computing, ser. HPDC '11. New York, NY, USA: ACM, 2011, pp. 171–182. [Online]. Available: http://doi.acm.org/10.1145/1996130.1996154

[15] C. Isci, J. Liu, B. Abali, J. O. Kephart, and J. Kouloheris, "Improving server utilization using fast virtual machine migration," IBM J. Res. Dev., vol. 55, no. 6, Nov. 2011, pp. 365–376.

[16] S. Srikantaiah, A. Kansal, and F. Zhao, "Energy aware consolidation for cloud computing," in Proceedings of the 2008 conference on Power aware computing and systems, ser. HotPower'08. Berkeley, CA, USA: USENIX Association, 2008, pp. 10–10. [Online]. Available: http://dl.acm.org/citation.cfm?id=1855610.1855620

[17] F. Hermenier, X. Lorca, J.-M. Menaud, G. Muller, and J. Lawall, "Entropy: a consolidation manager for clusters," in Proceedings of the 2009 ACM SIGPLAN/SIGOPS international conference on Virtual execution environments, ser. VEE '09. New York, NY, USA: ACM, 2009, pp. 41–50. [Online]. Available: http://doi.acm.org/10.1145/1508293.1508300

On the Use of Remote GPUs and Low-Power Processors for the Acceleration of Scientific Applications

A. Castelló*, J. Duato*, R. Mayo†, A. J. Peña‡, E. S. Quintana-Ortí†, V. Roca†, F. Silla*
*Universitat Politècnica de València, València, Spain.
Emails: {adcasgi,jduato,fsilla}@gap.upv.es
†Depto. de Ingeniería y Ciencia de Computadores, Universitat Jaume I, Castellón, Spain.
Emails: {mayo,quintana,vroca}@uji.es
‡Mathematics and Computer Science Division, Argonne National Laboratory
Argonne (IL), USA. Email: apenya@anl.gov

***Abstract*—Many current high-performance clusters include one or more GPUs per node in order to dramatically reduce application execution time, but the utilization of these accelerators is usually far below 100%. In this context, remote GPU virtualization can help to reduce acquisition costs as well as the overall energy consumption.**

In this paper, we investigate the potential overhead and bottlenecks of several "heterogeneous" scenarios consisting of client GPU-less nodes running CUDA applications and remote GPU-equipped server nodes providing access to NVIDIA hardware accelerators. The experimental evaluation is performed using three general-purpose multicore processors (Intel Xeon, Intel Atom and ARM Cortex A9), two graphics accelerators (NVIDIA GeForce GTX480 and NVIDIA Quadro M1000), and two relevant scientific applications (CUDASW++ and LAMMPS) arising in bioinformatics and molecular dynamics simulations.

***Index Terms*—High Performance Computing; Graphic Processing Units (GPUs); CUDA; Virtualization; Scientific Computing; Energy-Aware Computing;**

I. Introduction

In the quest for the enormous benefits that Exascale applications promise [1]–[4], the Top500 ranking [5] and its greener counterpart, the Green500 list [6], show an impressive 6× improvement in the performance-power ratio of large-scale high performance computing (HPC) facilities over the last five years. Furthermore, a trend clearly visible in these two lists is the adoption of hardware accelerators to attain unprecedented levels of raw performance with reasonable energy costs, which hints that future Exaflop systems will most likely leverage some sort of specialized hardware.

Many supercomputers in the first positions of these two lists currently accommodate mainstream x86 based processors along with top-of-the-line accelerator technologies. The alternative proposed by the Mont-Blanc project [7] investigates how to aggregate a large number of low-power components (specifically ARM processors and accelerators with small numbers of cores) to build a Petascale general-purpose HPC cluster. In any of these cases though, it is unlikely that the accelerators that integrate the system are used 100% of the time. Therefore, for practical purposes, a cluster with one or more accelerators per node surely leads to a waste of energy and money, due to the underutilization of these devices. In contrast to that configuration, a cluster where only a few of the nodes are equipped with hardware accelerators is a more cost-effective approach in terms of energy usage, maintenance and acquisition costs.

In order to render such a reduced amount of accelerators accessible from any node in the cluster, we have heavily invested in the development of `rCUDA`[8]–[11]with seamless access to an NVIDIA GPU residing in a remote node. Although based on a simple remote procedure call (RPC) mechanism, as of today `rCUDA` is the only CUDA 5.0-compatible solution. Furthermore, compared with many of the other remote GPU virtualization solutions [12]–[17], `rCUDA` is not only compatible with CUDA 5.0, but it is also publicly available and it supports several interconnects, including the last FDR InfiniBand fabric.

In this paper we analyze the possibilities of leveraging `rCUDA` in a "heterogeneous" environment, with clients and servers running in nodes with very different capabilities. In particular, this paper makes the following original contributions:

- We provide an experimental analysis that identifies the potential and, to some extent, the overhead sources that affect the performance of distinct CPU-GPU configurations. In particular, our hardware setup includes a variety of scenarios, with the `rCUDA` client and server running, respectively, on nodes equipped with three different types of general-purpose multicore processors and two types of GPUs
- For the evaluation we select two complex CUDA-enabled applications: CUDASW++ [18] and LAMMPS [19].
- Finally, our experimental analysis focuses on the execution time, but also considers average power dissipation and energy consumption.

The rest of the paper is structured as follows. In Section II, we provide a brief overview of the `rCUDA` software solution.

 ISBN: 978-1-61208-332-2

In Section III, we describe the applications and setup employed for the experimental study that follows next, in Section IV. We close the paper with a few concluding remarks and a discussion of future work in Section V.

II. Overview of rCUDA

rCUDA is structured following a client-server distributed architecture, where the client middleware runs in the same cluster node as the application demanding GPU acceleration services, while the server middleware runs in the cluster node where the physical GPU resides.

In rCUDA, the client middleware offers the exact same interface as the regular NVIDIA CUDA Runtime API so that the application is not aware that it is interacting with rCUDA instead of a real GPU. To support a concurrent scenario where GPUs are shared and this sequence of events occurs concurrently with analogous interactions initiated by other applications, rCUDA manages independent GPU contexts for each application.

rCUDA accommodates several underlying client-server communication technologies [11] thanks to its modular layered architecture, which supports runtime-loadable network-specific communication libraries. rCUDA currently provides communication modules tailored for Ethernet- and InfiniBand-based networks and takes advantage of the increased performance of the last FDR InfiniBand fabric [11].

Furthermore, regardless of the specific communication technology, data transfers between rCUDA clients and servers are pipelined for performance, using preallocated buffers of pinned memory.

III. Experimental Setup

In this section, we describe the applications and hardware platforms involved in the experimental study.

A. Applications

CUDASW++ [18] is a bioinformatics software for Smith-Waterman protein database searches that exploits the massively parallel CUDA architecture of NVIDIA Tesla GPUs to perform sequence searches. In our study, we use release 2.0 of the package, with the following execution parameters `-query P010008.fasta -db uniprot_sprot.fasta -use_single 0`

LAMMPS [19] is a classic molecular dynamics simulator that can be used to model atoms or, more generically, as a parallel particle simulator at the atomic, mesoscopic, or continuum scale. For the tests in the next section, we used release 3Jan13 of the software, and benchmark lj included in the release. The specific parameter list employed in the experimentation with this application was `-sf cuda -v g 1 -v x 76 -v y 76 -v z 76 -v t 2000 < in.lj.cuda`

B. Systems

Our general-purpose platforms included the following three testbeds:

- KAYLA: A SECO mITX board consisting of an NVIDIA Tegra 3 ARM Cortex A9 quad-core CPU (1.4 GHz), 2 GB of DDR3 RAM, and an Intel 82574L Gigabit Ethernet controller.
- ATOM: A board with an Intel Atom quad-core CPU S1260 (2.0 GHz), 8 GB of DDR3 RAM, and an Intel I350 Gigabit Ethernet controller.
- XEON: A server with an Intel Xeon X3440 quad-core processor (2.4 GHz), 8 GB of DDR3 RAM, and an Intel 82574L Gigabit Ethernet controller. When this platform acted as an rCUDA server, three of these cores remained disabled via BIOS to reduce the power dissipation.

All three systems operated under Linux Ubuntu 12.04 with the compiler GNU gcc/g++ version 4.6.3.

On the other hand, two types of accelerator-equipped systems were involved in the experiments:

- CARMA: A SECO development kit, with an NVIDIA Tegra 3 ARM Cortex A9 quad-core CPU (1.4 GHz) plus an NVIDIA Quadro 1000 M GPU (96 CUDA cores), 2 GB of DDR3 RAM for the ARM and 2 GB of DDR5 RAM for the GPU. These two components communicate via a PCIe $\times 4$ Gen 1 link and the network controller was an Intel 82574L Gigabit Ethernet. This system is operated under Linux Ubuntu 12.04 with the compiler GNU gcc/g++ version 4.6.3.
- FERMI: An NVIDIA GeForce GTX480 "Fermi" GPU (448 cores), with 1,280 MB of DDR3/GDDR5 RAM. This graphics card was connected to either XEON, through a PCIe $\times 16$ Gen 1 link, or to KAYLA, via a slower PCIe $\times 4$ Gen 1 link.

The CUDA Tool Kit 5.0 was employed for both accelerators.

As these systems offered us a large variety of combinations to evaluate (specifically, 4 clients $\times$ 3 servers), we made a preliminary selection based on some initial tests and considerations:

- A few initial experiments determined that, to attain relevant acceleration factors for these two applications with respect to a parallel execution using a multicore GPU-less platform, the 448-core FERMI GPU had to be involved. We therefore discarded those configurations with CARMA as a server, due to its poweless 96-core GPU.
- When acting as a client, there is no difference between CARMA and KAYLA, as the two systems include the same type of general-purpose processor. From the points of view of power and energy, when the GPU in the CARMA system is not used, it contributes little to these two factors. Therefore, we only considered the former system for the client side in our scenarios.

Table I illustrates the different hardware configurations selected for the evaluation.

Finally, the node interconnect was a CISCO SLM2009 Gigabit Ethernet switch. We note here that only the XEON-

TABLE I: DIFFERENT SCENARIOS INVOLVED IN THE EXPERIMENTAL EVALUATION.

Scenario	Client	Server	Configuration
A	CARMA	KAYLA+FERMI	Low-power ARM-based client; low-power server
B	ATOM	KAYLA+FERMI	Low-power Atom-based client; low-power server
C	XEON	KAYLA+FERMI	Power-hungry client; low-power server
D	CARMA	XEON+FERMI	Low-power ARM-based client; power-hungry server
E	ATOM	XEON+FERMI	Low-power Atom-based client; power-hungry server
F	XEON	XEON+FERMI	Power-hungry client; power-hungry server

based system can be connected to a faster InfiniBand switch. Thus, for the comparison, we restrict the study to use the Gigabit Ethernet network.

C. Power and time measurement

All power data was collected using a WATTSUP?PRO wattmeter, connected to the line from the electric socket to the power supply unit (PSU), which reports instantaneous power with an accuracy of $\pm 1.5\%$ at a rate of 1 sample/s. The measures were recorded in a separate server so that the sampling process did not interfere with the accuracy of the results. In the tests, we initially ran the application under evaluation for an initial warm-up period (about 60 s); then, the execution is repeated 5 times or until enough power samples were available, the slowest and fastest repetitions are discarded, and the result was averaged and multiplied by the corresponding run time in order to obtain the energy consumption.

IV. EXPERIMENTAL EVALUATION

In this section, we present the results obtained from the execution of the two applications chosen in this study. In order to serve as a reference, we first evaluate the performance of the codes with a local GPU (the traditional scenario), and next compare these data with the performance achieved when accessing remote GPUs.

A. Acceleration via a local GPU

We open this section with an initial evaluation of the performance of the two applications, CUDASW++ and LAMMPS, when accelerated on a system equipped with a local GPU. Therefore, the only platforms that can be included in this evaluation are CARMA, KAYLA+FERMI, and XEON+FERMI. We summarize the main results from this evaluation in Table II.

CUDASW++: From the point of view of run time, the clear winner for this application is XEON+FERMI, with an execution time of 4.81 s, which is 4.79× and 7.35× faster than those observed for KAYLA+FERMI and CARMA, respectively. On the other hand, from the perspective of power (important, e.g., for a power capped environment), CARMA exhibits a much lower average power draw, roughly by factor around 5× with respect to the other two alternatives. Nevertheless, when execution time and average power are combined into a single figure-of-merit such as energy, the high-performance XEON+FERMI is again clearly superior to the other two counterparts, by factors of 1.36× and 4.28×.

LAMMPS: The behavior of the CUDA-accelerated version of this application is quite different. Now the execution times for KAYLA+FERMI and XEON+FERMI are much closer (217.94 s and 157.91 s, respectively), and they both outperform CARMA by a wide margin (1,208.29 s). This is indicative that this application makes a much more intensive use of the GPU than CUDASW++. However, for this particular application, the average power is also significantly increased for the former two platforms so that, interestingly, the total energy usage of the three systems is very similar.

To close this initial analysis, a direct comparison between the performances of KAYLA+FERMI and XEON+FERMI for the two applications in Table II illustrates the cost of the reduced PCIe bandwidth for the former platform, as both systems leverage the same type of high-performance GPU (GeForce GTX480) and the bulk of the computation is offloaded to the accelerator. Additionally, the comparison between KAYLA+FERMI and CARMA shows the negative impact caused by the reduced number of cores in the GPU featured by the latter (96 cores vs 448 in the GTX480). That results caused partially our decision of discarding this system for the server side.

B. Acceleration via a remote GPU

In the following, we experimentally analyze the potential and overheads of the different hardware scenarios listed in Table I. In all cases the application runs in a single `rCUDA` client/node that accesses the GPU connected to a single remote `rCUDA` server/node using our software middleware. No changes were made to the original GPU-accelerated CUDASW++ and LAMMPS codes other than those already necessary to install and run these packages in the platforms with the local GPU. The only change was during the compilation phase, where the applications were linked to the current release of `rCUDA`, which replaces the usual CUDA library.

Note that possible sources of overhead (bottlenecks) are the `rCUDA` middleware; the interface to the Gigabit Ethernet interconnect in the client and/or the server; (the bandwidth of) the Gigabit Ethernet; and (the bandwidth of) the PCIe interface in the server. However, the cost of this last factor was already exposed in the previous study for the local GPU case.

TABLE II: PERFORMANCE OF THE SELECTED APPLICATIONS WHEN RUNNING ON A SINGLE PLATFORM ACCESSING THE LOCAL GPU.

System	CUDASW++			LAMMPS		
	Time (s)	Avg. power (W)	Energy (J)	Time (s)	Avg. power (W)	Energy (J)
CARMA	35.40	24.88	880.92	1,208.29	34.14	43,663.34
KAYLA+FERMI	23.07	120.30	2,775.48	217.94	203.31	44,309.17
XEON+FERMI	4.81	134.56	647.27	157.91	268.43	42,388.21

TABLE III: PERFORMANCE, AVERAGE POWER AND CONSUMPTION OF THE TWO SELECTED APPLICATIONS WHEN RUNNING ON AN rCUDA CLIENT ACCESSING THE REMOTE GPU IN THE rCUDA SERVER.

Application	Scenario	System		Time (s)	Avg. power (W)			Energy (J)		
		Client	Server		Client	Server	Total	Client	Server	Total
CUDASW++	A	CARMA	KAYLA+FERMI	73.87	11.11	123.82	134.93	820.65	9,146.69	9,967.35
	B	ATOM	KAYLA+FERMI	61.85	44.23	120.62	164.85	2,735.98	7,460.27	10,196.26
	C	XEON	KAYLA+FERMI	61.54	58.68	117.55	176.23	3,610.90	7,234.02	10,844.92
	D	CARMA	XEON+FERMI	13.72	11.10	138.45	149.56	152.26	1,898.96	2,051.22
	E	ATOM	XEON+FERMI	11.38	42.65	143.09	185.74	485.11	1,627.65	2,112.77
	F	XEON	XEON+FERMI	6.33	58.68	138.51	197.18	371.26	876.38	1,247.65
LAMMPS	A	CARMA	KAYLA+FERMI	2,124.78	10.97	134.28	145.25	23,305.79	285,314.39	308,620.19
	D	CARMA	XEON+FERMI	482.59	12.45	191.86	204.31	6,006.23	92,589.15	98,595.38
	E	ATOM	XEON+FERMI	268.92	44.29	226.48	270.77	11,910.81	60,903.28	72,814.09
	F	XEON	XEON+FERMI	236.91	68.38	251.83	320.21	16,199.09	59,661.11	75,860.21

CUDASW++: The first aspect that is manifest in Table III and the left-hand side plots of Figures 1 and 2 is the much higher execution time of the scenarios that use the KAYLA-based accelerator as a server, even though the GPU is the same as in their XEON-based counterparts (a "Fermi" GPU). Clearly, the source of this much inferior performance must be the interface to the Gigabit interconnect for KAYLA, as the results in Table II already revealed that the slow PCIe of KAYLA, while constraining the performance to a certain extent, does not justify the large gap between the configurations that leverage this type of equipment and those based on the XEON server. Actually, in the previous section, we saw a 4.79× speed up between the KAYLA and XEON platforms, whereas in Table III speed up between both platforms, when using XEON-based client increases up to 9.72×. To confirm that the interface to the Gigabit interconnect is the cause of the higher execution time, we carried out a separate test using iperf [20]. The results from these independent experiments, collected in Table IV, show that the network interface for KAYLA delivers much less than the 1 Gbps bandwidth that could be expected from a Gigabit Ethernet, and confirm this as the origin of the bottleneck. On the other hand, from the point of view of average power, the configurations with a KAYLA-based server do not show significant advantages over the XEON-based ones (considering the server only, approximately between 14 and 22 W, or 10 to 15%), which in combination with the time differences explain why the configurations that involve the latter type of server lead to a much more energy-efficient solution.

Let us focus now on the behavior of CUDASW++ when running on the three configurations with the XEON-based server, which should help to identify the source of overheads in the client. The much shorter execution time of the configuration when the XEON acts as a client (6.33 s) compared with the CARMA-client and ATOM-client (13.72 s and 11.38 s, respectively), and the smaller differences in the combined average power from the two systems (149.56 W for CARMA, 185.74 W for ATOM, and 197.18 W for XEON), render the superiority of the XEON-client solution, which is almost twice more efficient in terms of energy than the two alternative configurations. In this case, as all other factors remain the same, the source of the overhead for the CARMA-based client must be the low bandwidth of the interface to the Gigabit Ethernet in the client (see Table IV). On the other hand, for the ATOM-based client, we place responsibility for the low performance in the client CPU itself.

LAMMPS: Although the results for the previous case already identified the serious bottleneck that the interface to the Gigabit Ethernet poses for the KAYLA+FERMI server, for illustrative purposes only, we still evaluate one such type of configuration. Again, this scenario exhibits a much worse performance which, combined with the limited improvement of the average power, render a much higher energy costs for the KAYLA-based server; see Table III and the right-hand side

 ISBN: 978-1-61208-332-2

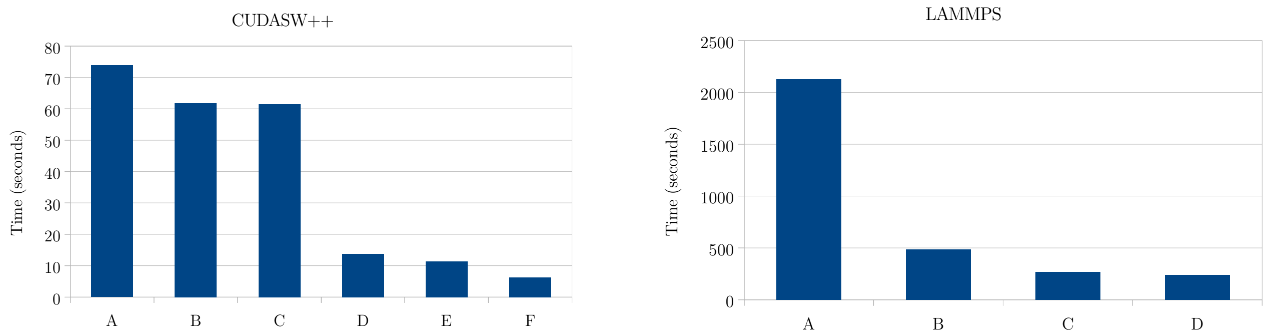

Fig. 1: Performance of CUDASW++ (rigth) and LAMMPS (left) when running on an `rCUDA` client accessing the remote GPU in the `rCUDA` server. See Table I for a description of the scenarios.

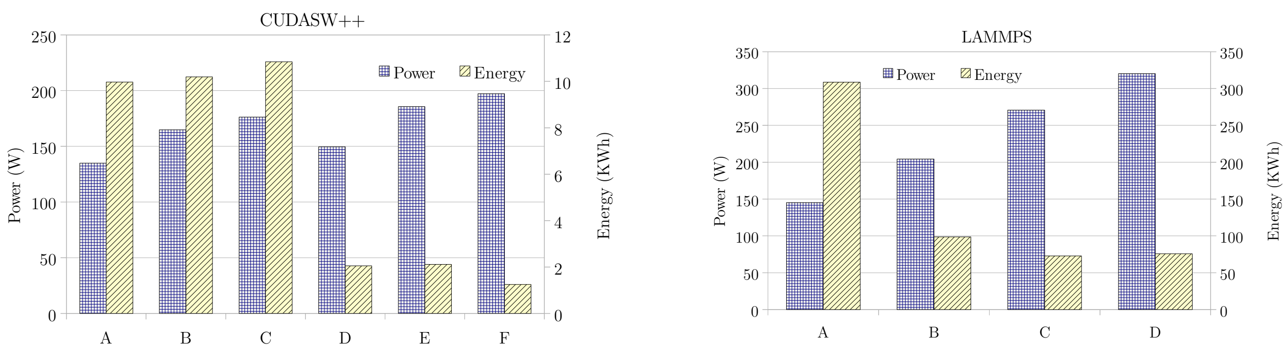

Fig. 2: Average power and energy consumption of CUDASW++ (rigth) and LAMMPS (left) when running on an `rCUDA` client accessing the remote GPU in the `rCUDA` server. See Table I for a description of the scenarios.

plots in Figures 1 and 2.

From the perspective of execution time, using the XEON equipment in both the client and server sides offers the best solution (236.91 s, compared to 482.59 s and 268.82 s in the configurations with CARMA and ATOM as clients, respectively). However, in exchange for this increase of the execution time, the ATOM-server client offers a slight but non-negligible advantage in terms of energy efficiency over the XEON-based one, of about 4.13%.

C. Overhead with respect to the local GPU

Before we perform this final analysis, it is important to realize that, from the point of view of time, the execution of an application that interacts with a remote GPU can never outperform the same application running on a platform equipped with a local GPU (provided both systems feature same type of GPU and PCIe connection). Actually, the best we can expect is that the remote access yields a low overhead, so that the execution times between the two cases (local and remote GPU) are close. From the perspective of energy, the single-client single-server configuration also works in favor of the local GPU, as we are now comparing configurations with one system (local GPU) versus two (`rCUDA` client and GPU-equipped `rCUDA` server). Nevertheless, the use of the `rCUDA` middleware allows to build a cluster with a reduced number of GPUs, even adapted for a specific purpose, which can be expected to decrease the total energy (as well as acquisition) costs with respect to a configuration where all cluster nodes are equipped with at least one GPU, that will likely remain idle a significant fraction of the time. In summary, the purpose of `rCUDA` is to reduce the energy costs for the cluster as a whole while introducing a low overhead in terms of run time.

For simplicity, let us consider only the best configurations in terms of performance, which correspond to the XEON-based client and server. Comparing this with the XEON system with a local GPU, we observe that the overhead introduced by the remote access for these two applications is rather visible, of 24.01% for CUDASW++ and 33.34% for LAMMPS. However, notice that this overhead is mainly due to the use of a slow interconnect, and the use of an FDR InfiniBand network turns them mostly negligible [11]. On the other hand, the potential benefits that a system with a moderate number of GPUs offers can be hinted by comparing the differences in the average power of the platforms when acting as clients or servers. In the same line, the power dissipated by an idle GPU can provide a more accurate estimation. In particular, the data in Table V indicate that the power dissipated by an idle FERMI GPU is between 36.4 W (compare the idle power rate of the XEON with and without the GPU) and 31.2 W (compare the figures for KAYLA –that has been included only for this

TABLE IV: BANDWIDTH BETWEEN TWO NODES, USING THE `iperf` TEST EXECUTED DURING 10 s.

Client	Bandwidth (Mbs)	
	KAYLA server	XEON server
CARMA	392	895
ATOM	439	940
XEON	435	943

TABLE V: POWER DISSIPATED BY THE EQUIPMENT WHILE IDLE.

Client Systems		Server Systems	
System	Power (W)	System	Power (W)
CARMA	10.4	KAYLA	41.3
ATOM	42.5	KAYLA+FERMI	72.5
XEON	54.4	XEON+FERMI	90.8

comparison– and KAYLA+FERMI), which is quite significant when contrasted to the power costs of a GPU-less node.

V. REMARKS AND FUTURE WORK

The tendencies captured by the Top500 and Green500 list portray a landscape for HPC consisting of heterogeneous facilities, with nodes of different types, specifically tailored for certain classes of applications. In this future, we envision that hardware accelerators, either in the form of data-parallel GPUs or more general-purpose architectures like the Intel Xeon Phi, will play a relevant role to build the first Exascale system. Furthermore, independently of whether the Exaflop barrier is reached using a "fat"-node approach (i.e., a large number of highly-multithreaded nodes) or a "thin"-node alternative (i.e., a huge number of low-core nodes), we believe that not all nodes/processes/threads in these systems will have direct access to an accelerator, and therefore some sort of remote access needs to be granted.

The main contribution of this paper is an experimental evaluation of the possibilities that state-of-the-art technology offers in today's HPC facilities, as well as low-power alternatives offer for the acceleration of scientific applications using remote graphics processors. In particular, we have assessed the potential of distinct hardware configurations, including three general-purpose multicore processors (Intel Xeon, Intel Atom and ARM Cortex A9) for the client side, and two types of graphics accelerators (NVIDIA GeForce GTX480 and NVIDIA Quadro M1000) for the server side. Our experiments with two key scientific applications in bioinformatics and molecular dynamics simulations, CUDASW++ and LAMMPS respectively, reveal an important bottleneck in the access to the network for the ARM-based client/server, with significant negative consequences on both the execution time and energy consumption. While the performance of the Xeon-based configuration is much better, the overhead when compared to the access to a local GPU is quite relevant, and clearly asks for the use of a faster interconnect.

In the future we plan to investigate some of the bottlenecks detected in this work, while monitoring new technology that may appear in principle solving the problem in the access to the network. We also plan to act on the applications to improve their performance (as described in the previous paragraph).

ACKNOWLEDGMENT

The researchers at UPV were supported by the the Generalitat Valenciana under Grant PROMETEOII/2013/009 of the PROMETEO program phase II. Researchers at UJI were supported by MINECO, by FEDER funds under Grant TIN2011-23283, and by the Fundacion Caixa-Castell Bancaixa (Grant P11B2013-21). This work was partially supported by the U. S. Department of Energy, Office of Science, under Contract No. DE-AC02-06CH11357. The authors are also grateful for the generous support provided by Mellanox Technologies.

REFERENCES

[1] S. Ashby *et al*, "The opportunities and challenges of Exascale computing," Summary Report of the Advanced Scientific Computing Advisory Committee (ASCAC) Subcommittee, November 2010.
[2] K. Bergman *et al*, "Exascale computing study: Technology challenges in achieving exascale systems," DARPA IPTO ExaScale Computing Study, 2008.
[3] J. Dongarra *et al*, "The international ExaScale software project roadmap," *Int. J. of High Performance Computing & Applications*, vol. 25, no. 1, pp. 3–60, 2011.
[4] M. Duranton *et al*, "The HiPEAC vision for advanced computing in horizon 2020," pp. 1–48, 2013.
[5] "The top500 list," Nov.,2013, available at http://www.top500.org.
[6] "The Green500 list," Nov., 2013, available at http://www.green500.org.
[7] "The Mont Blanc project," http://montblanc-project.eu, Feb.,2014.
[8] J. Duato, F. D. Igual, R. Mayo, A. J. Peña, E. S. Quintana-Ortí, and F. Silla, "An efficient implementation of GPU virtualization in high performance clusters," *Euro-Par 2009, Parallel Processing – Workshops*, vol. 6043, pp. 385–394, 2010.
[9] J. Duato, A. J. Peña, F. Silla, R. Mayo, and E. S. Quintana-Ortí, "Performance of CUDA virtualized remote GPUs in high performance clusters," in *Proceedings of the 2011 International Conference on Parallel Processing (ICPP 2011)*, Sep. 2011, pp. 365–374.
[10] J. Duato, A. J. Peña, F. Silla, J. C. Fernández, R. Mayo, and E. S. Quintana-Ortí, "Enabling CUDA acceleration within virtual machines using rCUDA," in *Proceedings of the 2011 International Conference on High Performance Computins (HiPC 2011)*, Dec. 2011, pp. 1–10.
[11] C. Reaño, R. Mayo, E. S. Quintana-Ortí, F. Silla, J. Duato, and A. J. Peña, "Influence of Infiniband FDR on the performance of remote GPU virtualization," in *IEEE Cluster 2013*, 2013, pp. 1–8.
[12] M. Oikawa *et al.*, "DS-CUDA: a middleware to use many GPUs in the cloud environment," in *SC*, 2012, pp. 1207–1214.
[13] G. Giunta *et al.*, "A GPGPU transparent virtualization component for high performance computing clouds," in *Euro-Par*, 2010, pp. 379–391.
[14] T.-Y. Liang *et al.*, "GridCuda: a grid-enabled CUDA programming toolkit," in *WAINA*, 2011, pp. 141–146.
[15] V. Gupta *et al.*, "GViM: GPU-accelerated virtual machines," in *HPCVirt*, 2009, pp. 17–24.
[16] Zillians, Inc. (Feb.,2014) V-GPU: GPU virtualization.
[17] L. Shi, H. Chen, and J. Sun, "vCUDA: GPU accelerated high performance computing in virtual machines," in *IEEE International Symposium on Parallel & Distributed Processing (IPDPS'09)*, 2009, pp. 1–11.
[18] Y. Liu, A. Wirawan, and B. Schmidt, "CUDASW++ 3.0: accelerating Smith-Waterman protein database search by coupling CPU and GPU SIMD instructions," *BMC Bioinformatics*, vol. 14, no. 1, p. 117, 2013.
[19] S. Plimpton, "Fast parallel algorithms for short-range molecular dynamics," *Journal of Computational Physics*, vol. 117, pp. 1–19, 1995.
[20] "TCP and UDP bandwidth performance measurement tool," Feb.,2014, `https://code.google.com/p/iperf/`.

Analysis of Novel Random Neural Network Controller for Residential Building Temperature Control

Abbas Javed, Hadi Larijani, Ali Ahmadinia, Rohinton Emmanuel

School of Engineering and Built Environment
Glasgow Caledonian University
Glasgow, United Kingdom
Email: abbas.javed@gcu.ac.uk, H.Larijani@gcu.ac.uk, Ali.Ahmadinia@gcu.ac.uk
Rohinton.Emmanuel@gcu.ac.uk

***Abstract*—Random neural networks (RNN) have strong generalisation capabilities and are easy to implement on hardware as compared to Artificial Neural Networks (ANN). In this paper, a novel RNN controller is proposed to maintain a comfortable indoor environment in a single zone residential building fitted with radiators for heating. This controller is capable of maintaining a comfortable indoor environment on the basis of a predicted mean vote (PMV)-based set point. The implemented RNN controller is compared with ANN controller for energy consumption, indoor room temperature, and minimum square error. Results show that for same training data and learning algorithm parameters, RNN converges faster and it consumes less energy, results in better comfortable room temperature as compared to ANN controller.**

Keywords-random neural network; artificial neural network; building simulation; residential heating system; energy efficient controller

I. INTRODUCTION

Buildings currently consume 40% of the total energy in most developed countries. The International Energy Agency (IEA) has set a target to reduce energy consumption in buildings by improving energy efficiency. This will result in estimated energy saving of 1509 million tonnes of oil equivalent (Mtoe) by 2050 [1]. The energy efficiency policy of IEA will not only save energy but it will also reduce carbon dioxide (CO_2) emission from the building sector. This will result in possible mitigation of 12.6 gigatonnes (Gt) of CO_2 emissions by 2050. According to Intelligent Energy Executive Agency (IEAE), in European households, 68% of energy consumption is for space heating, 14% for water heating and 13% for electric appliances [2].
The number of homes in UK has increased from 18 million to 26 million during the period of 1970-2011. The energy used in homes is equal to quarter of total energy used and CO_2 emission in UK [3]. UK's carbon emission reduction targets therefore will be impossible to achieve without reducing CO_2 emission in households. While many fabric-based and equipment-based interventions are needed, controllers for residential heating systems could play a useful part without compromising the occupants comfort.
According to [4], 95% of radiators were controlled using thermostatic radiator valves (TRVs). It was found that majority of TRVs failed to reduce the heating output once the room temperature is greater than the setpoint, as a result energy was wasted. The survey further revealed that 32% of TRVs were positioned at "Max" and more than 65% of TRVs were set for greater flowrates than required. If TRVs were kept within 2-3 settings (max 5), about 12.4% reduction of heat consumption could be gained compared with the situation in which the TRVs were kept fully open [5].
In non-domestic buildings for controlling the Heating Ventilation and Air Conditioning (HVAC), the control techniques are categorized in two parts, i.e., Local control and Supervisory control [6]. For local control, ON/OFF and proportional-integral derivative control schemes are normally used. The control settings of these local controllers might be optimal and energy efficient for certain subsystem however they may not be energy efficient for overall system as these control schemes are unable to maintain indoor comfort of the building by taking in to the account the ever changing indoor and outdoor environmental variables.
Supervisory control techniques are used for maintaining comfortable indoor environment by considering indoor and outdoor environment variables. Supervisory control techniques can be implemented by using physical model based techniques and black box techniques. Physical model-based techniques require physical model of the building to predict energy/cost of the concerned system which is computationally expensive and requires lot of memory. Black box techniques are normally implemented by ANN and RNN models. The ANN models are developed on empirical model of the system and are capable to mathematically relate the input and output variables of the system. ANN models are computationally less expensive than physical model based techniques but requires extensive training data to achieve accuracy.
The main contributions of this paper are:
1. Novel variable set point RNN and ANN controllers have been developed for optimisation of energy consumption by residential heating systems without compromising the thermal comfort of the occupants. The gradient descent algorithm is used to train the RNN and ANN controllers for predicting the optimised inflow of hot water in to the radiator. Variable set point estimated by PMV thermal comfort model, indoor environmental variables, and meteorological data have been used

by RNN and ANN controllers to predict the flow rate from the TRVs for optimised energy consumption and comfortable thermal environment.
2. The energy consumption by residential heating system controlled by RNN and ANN controller is compared by simulating the single zone building model in Matlab/Simulink for 100 days. The single zone building model is developed by using International Building Physics Toolbox (IBPT).
3. The training algorithm for RNN and ANN is critically analysed in terms of no. of iterations and minimum square error attained. The percentage of periods when air temperature overshoots the specified range of room temperature set point is calculated for ANN and RNN controller.
The rest of this paper is organized as follows. Brief introduction to RNN and ANN and learning algorithm for RNN are given in Section II. The implemented building model is described in Section III followed by description of intelligent heating control system in Section IV. The experimental results are provided in Section V followed by the discussions and conclusions in Section VI.

II. Related Work

A. Artificial Neural Network

The ANNs have been used in different applications for BEMS such as modeling the thermal dynamics of building space, estimation of heating loads of buildings, control of HVAC, prediction of energy consumption in buildings, and solar radiation predictions for non-domestic buildings [7]-[13]. The above mentioned ANN techniques are difficult to implement on hardware as they are computationally expensive. Residential water heating systems and radiant floor heating systems were effectively controlled by ANN based predictive control methods. In [14], the authors proposed an ANN based predictive control model. The results of their work showed that the performance of proposed predictive control is better than that of the current two-position ON/OFF control. In [15], [16], the authors showed that ANN based predictive control strategies are better than conventional control techniques in terms of energy saving, building thermal control with reduced overheating and overcooling. In [17], the authors developed an ANN based adaptive and predictive control method that ensures more comfortable thermal conditions than typical thermostat systems in terms of increased comfort periods of air temperature, humidity, PMV [18] - the most commonly used indoor thermal comfort index in buildings, and reduced over and undershoots. The ANN based models takes into account not only indoor air temperature but also PMV as a control variable in order to reduce overshoot and undershoot of the temperature which resulted in energy conservation. In [19],the author proposed the PMV index based variable set point control scheme and compared the results with fixed set point control scheme.

B. Random neural network

Gelenbe [20], [21] proposed the new class of ANN as RNN in which signals are either +1 or -1 due to which it is an excellent modeling tool. RNN can give more detailed system state description because the potential of neuron is represented by integer rather than binary value [22]. RNN is easy to implement on hardware as its neurons can be represented by simple counters [23], [24].

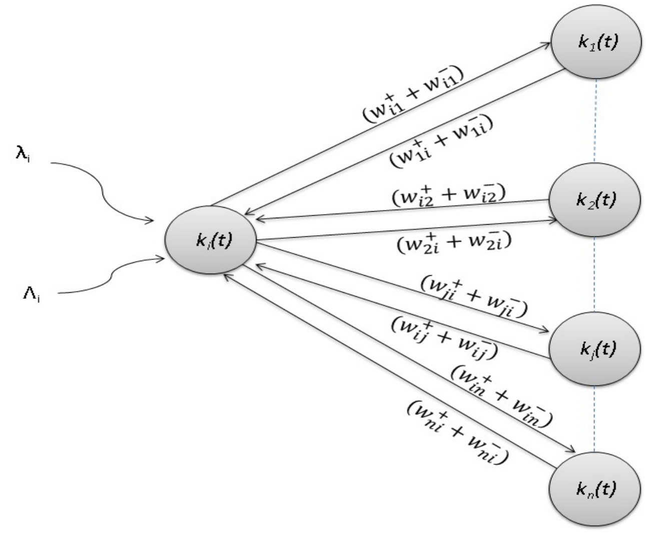

Figure 1. Random Neural Network

Applications of RNN have been reported for modeling, pattern recognition, image processing, classification, and communication systems [22], [25]. However, no such application has been reported so far in implementing control scheme for HVAC in residential/commercial buildings.
In RNN shown in Figure 1, signal travels in the form of impulse between the neurons. If the receiving signal has positive potential (+1) it represents excitation, and if the potential of the input signal is negative (-1) it represents inhibition to the receiving neuron. Each neuron *i* in the random neural network has a state $k_i(t)$ which represents the potential at time *t*. This potential $k_i(t)$ is represented by non-negative integer. If $k_i(t) > 0$ then neuron *i* is in excited state and if $k_i(t) = 0$ then neuron *i* is in idle state. When neuron *i* is in excited state, it transmits impulse according to the poisson rate r_i. The transmitted signal can reach neuron *j* as excitation signal with probability $p^+(i,j)$ or as inhibitory signal with probability $p^-(i,j)$, or can leave the network with probability $d(i)$ such that

$$d(i) + \sum_{j=1}^{N} \left[p^+(i,j) + p^-(i,j)\right] = 1 \forall i \tag{1}$$

$$w^+(i,j) = r_i p^+(i,j) \geqslant 0 \tag{2}$$

$$w^-(i,j) = r_i p^-(i,j) \geqslant 0 \tag{3}$$

combining (1)-(3)

$$r(i) = (1 - d(i))^{-1} \sum_{j=1}^{N} \left[w^+(i,j) + w^-(i,j)\right] \tag{4}$$

The firing rate between the neuron is represented by $r(i) = \sum_{j=1}^{N} [w^+(i,j) + w^-(i,j)]$. As *'w'* matrices are the product of firing rate and probabilities, therefore these matrices always hold non-negative values. External positive or negative signal can also reach neuron *i* at poisson rate Λ_i and λ_i respectively. When positive signal is received at neuron *i* its potential $k_i(t)$ will increase to +1. If neuron *i* is in excitation state and it receives negative signal the potential of neuron *i* will decrease

TABLE I
DESCRIPTION OF RNN SYMBOLS

RNN Symbols	Description
q_i	Probability neuron *i* excited at time *t*
$p^+(i,j)$	Probability neuron *j* receives positive signal from neuron *i*
$p^-(i,j)$	Probability neuron *j* receives negative signal from neuron
r_i	Firing rate of neuron *i*
Λ_i	Arrival rate of external positive signals
λ_i	Arrival rate of external negative signals
$d(i)$	Probability a signal from neuron departs from the network
$k_i(t)$	Potential of neuron *i* at time *t*

to zero. Arrival of negative signal will have no effect on neuron *i* if its potential is already 0. The description of symbols used is given in Table 1.
Consider the vector $\mathbf{K}(\mathbf{t})$= $(k_1(t),.....k_n(t))$ where $k_i(t)$ is the potential of neuron *i* and n is the total number of neurons in the network. Let $\mathbf{K}$ is continuous time Markov process. The stationary distribution of $\mathbf{K}$ is represented as:

$$\lim_{t\to\infty} Pr(K(t))) = (k_1(t).......k_n(t)) = \prod_{i=1}^{n}(1-q_i)q_i^{ni} \quad (5)$$

For each node *i*

$$q_i = \frac{G_i^+}{r_i + G_i^-} \quad (6)$$

where

$$G_i^+ = \Lambda_i + \sum_{j=1}^{N} q_j w^+(j,i) \quad (7)$$

$$G_i^- = \Lambda_i - \sum_{j=1}^{N} q_j w^-(j,i) \quad (8)$$

For three layer network, q_i for each layer is calculated as

$$q_{i\varepsilon I} = \frac{\Lambda_i}{r_i + \lambda_i} \quad where\ Iis\ input\ layer \quad (9)$$

$$q_{i\varepsilon H} = \frac{\sum_{i\varepsilon I} q_i w^+(i,h)}{r_h + \sum_{i\varepsilon I} q_i w^-(i,h)} \quad where\ His\ hidden\ layer \quad (10)$$

$$q_{i\varepsilon O} = \frac{\sum_{i\varepsilon H} q_h w^+(h,o)}{r_h + \sum_{i\varepsilon I} q_h w^-(h,o)} \quad where\ Ois\ Output\ layer \quad (11)$$

In this paper, we propose a novel variable set-point RNN controller for maintaining comfortable indoor environment in single zone residential building by controlling the motorized TRVs mounted on radiator. RNN controller uses room temperature, error (difference between current room temperature and variable setpoint), and outside temperature as Inputs to predict flow rate m' $(\frac{m^3}{hr})$ from the TRVs for optimized energy consumption while maintaining comfortable thermal environment.

C. Gradient Descent learning algorithm for RNN

Suppose we have data set F composed of M input-output pairs (x^m, y^m) where $m = 1, 2, ...M$ and $x^m = [\Lambda^m \lambda^m]$ such that x^m are pairs of excitation and inhibition signal ow rates entering each neuron from outside of the network. Output $y^m \varepsilon[0..1]$ where $y^m = f(x^m)$. The goal of the learning algorithms is to find parameters for RNN such that difference between q_i^m and y_i^m is minimum. Similarly gradient descent algorithm developed by Gelenbe [26] adjusts the parameters in order to minimize the cost function E_m

$$E_m = \frac{1}{2}\sum_{i=1}^{n} a_i(q_j^m - y_j^m)^2, \quad a_i \geqslant 0 \quad (12)$$

The rule of updating the weights by using m^{th} input-output data pair for connection between neuron *e* and *f* is

$$w_{(e,f)}^{+t} = w_{(e,f)}^{+(t-1)} - \eta \sum_{i=1}^{n} a_i(q_j^m - y_j^m)\left[\frac{\partial q_i}{\partial w_{(e,f)}^+}\right]^{t-1}$$

$$w_{(e,f)}^{-t} = w_{(e,f)}^{-(t-1)} - \eta \sum_{i=1}^{n} a_i(q_j^m - y_j^m)\left[\frac{\partial q_i}{\partial w_{(e,f)}^-}\right]^{t-1} \quad (13)$$

where

$$\left[\frac{\partial q_i}{\partial w_{(e,f)}^+}\right] = \gamma_{e,f}^+ q_e[I-W]^{-1}$$

$$\left[\frac{\partial q_i}{\partial w_{(e,f)}^-}\right] = \gamma_{e,f}^- q_e[I-W]^{-1} \quad (14)$$

$$\gamma_{e,f;i}^+ = \begin{cases} \frac{-1}{r_i+G_i^-} & if\ e=i, f\neq i \\ \frac{1}{r_i-G_i^-} & if\ e\neq i, f=i \\ 0 & else \end{cases}$$

$$\gamma_{e,f;i}^- = \begin{cases} \frac{(-1+q_i)}{r_i+G_i^-} & if\ e=i, f=i \\ \frac{-1}{r_i+G_i^-} & if\ e=i, f\neq i \\ \frac{-q_i}{r_i+G_i^-} & if\ e\neq i, f=i \end{cases}$$

$$W(i,j) = \frac{w_{(i,j)}^+ - w_{(i,j)}^- q_j}{D_j} \quad i,j = 1......, N \quad (15)$$

Steps for gradient descent learning algorithm are as followings:
$a)$ Initialize $w_{(e,f)}^+$ and $w_((e,f))^-$ ∀ e, f and choose suitable value for learning rate
$b)$ For all input output pairs, initialize Λ_{im}, λ_{im} according to X_{im}.
$c)$ Solve (6)-(8) by using current weight values
$d)$ Calculate W, $\gamma_{(e,f)}^+, \gamma_{(e,f)}^- \forall e, f$
$e)$ Calculate $\frac{\partial q_i}{\partial w_{(e,f)}^+}$ and $\frac{\partial q_i}{\partial w_{(e,f)}^-}$ by solving (14)
$f)$ Update the weights from (13)
Weights are product of firing rate and probability and can never have negative values. After solving (13) negative weights can either be set to zero or repeat the iteration with smaller value of η. Repeat the procedure (b)-(f) until convergence or maximum number of iterations.

III. IMPLEMENTATION OF MODEL

In this work, a single zone building made of three layered walls/roof, fitted with Intelligent Controllers for heating/cooling system management is modeled in Matlab/Simulink using IBPT.

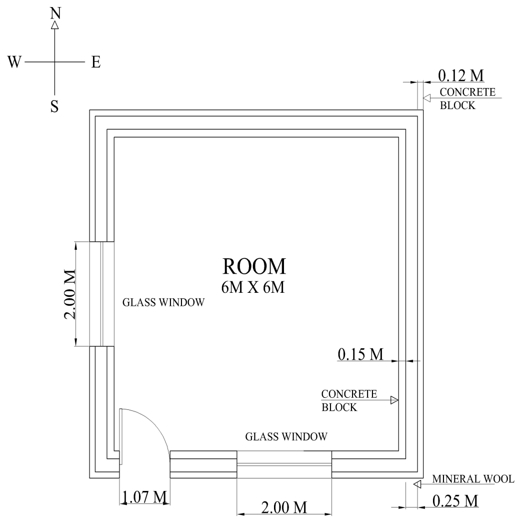

Figure 2. Layout of Walls

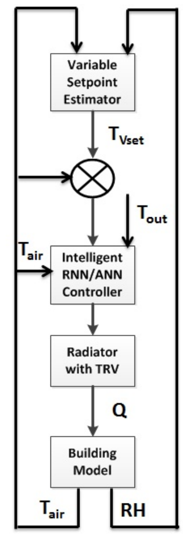

Figure 3. Control Scheme

A. Target Building

A single zone residential house is considered as the target building. The size of the house is 6 x6 x2.7 m^3. The house has a window in south and west wall with flat ceiling and three layered walls. Concrete blocks are used for exterior/interior material on each wall while mineral wool is used for insulation. Thicknesses of the materials are 0.15m, 0.25m and 0.12m respectively shown in Figure 3. Area of the uncoated double-glazed south & west windows is 4 m^2 and U-Value of window is 2.44 and resistance of air gap is 0.23.

In the simulation heat gains from lights, occupants and equipment are ignored. External parameters that affect the building environment such as ambient temperature, dewpoint temperature, global irradiation, diffuse horizontal irradiation, direct normal irradiation, longwave sky radiation, global luminance, diffuse horizontal illuminance, direct normal illuminance, wind direction, and wind speed are included in the weather data file. The building model is affected by meteorological parameters and heat emitted from radiator inside the room. The room is fitted with the radiator and TRV is mounted on inlet pipe of the radiator for controlling the flow rate of hot water entering the radiator. Heat emission from radiator is controlled by changing the flow rate m' of hot water.

B. Radiator Model

The heat transfer from radiator to surrounding is represented by (16). At low mass flow rate the radiator exhaust temperature is nearly equal to room temperature and heat emission from the radiator is linear function of mass flow rate.

$$q = m' c_p (T_{su} - T_{en}) \tag{16}$$

Where,
$q =$ heat flux W
$m' =$ fluid mass flowrate $\frac{Kg}{s}$
$c_p =$ specific heat capacity of fluid $\frac{J}{Kg-K}$
$T_{su} =$ radiator supply temperature °C.
$T_{en} =$ environment temperature °C

C. Thermostatic Radiator Valve Model

The TRV used in this model is 15mm valve whose flow rate Kv varies between 0 to 0.56 ($\frac{m^3}{hr}$) at differential pressure of 0.6 bar. The flowrate through the valve is represented by (17).

$$Kv = \frac{Q}{\sqrt{\triangle P}} \tag{17}$$

$Q =$ the flow rate in $\frac{m^3}{hr}$
$\triangle P =$ differential pressure of 0.6 bar.
The relationship between the position of the thermostatic radiator valve, flow rate and differential pressure is implemented in Simulink. The kv -value as a function of the TRV position is included in a lookup table.

IV. INTELLIGENT HEATING CONTROL SYSTEM

The main goal of the controller design is to reduce energy consumption while maintaining acceptable indoor thermal comfort for occupants. In Figure 4, the block diagram of intelligent heating system controller is shown. The variable set point estimator estimates the variable set point TVset for the building temperature by using PMV index as proposed in [19]. The intelligent controller takes the difference between the T_{Vset} and room air temperature (T_{air}), room air temperature (T_{air}), and outside temperature (T_{out}) as inputs and controls the heat emission from the radiator by changing the flowrate (m') through TRV mounted on the radiator. Two types of intelligent controllers are investigated in this study, i.e., ANN and RNN controller.

A. Variable Set Point Estimator

The Institute for Environmental Research at the Kansas State University conducted the research study under the contract of ASHRAE and defined PMV in terms of easily measured parameters [19].

$$PMV = at + bp_v - c \tag{18}$$

Where a, b, c are constants defined in Kansas State university research, p_v is vapour pressure and t is temperature.

 ISBN: 978-1-61208-332-2

By setting the required PMV (in this work between -0.3 and +0.3) and using the constants (a=0.220, b=0.233, c=5.673) heating setpoint for room temperature varies between room temperature varies between 22.76 ^oC to 23.77 ^oC and cooling setpoint varies between 26.59 ^oC to 27.58 ^oC.

B. Training Data

The training dataset have been generated by simulating the single zone building for 30 days in Matlab/Simulink using IBPT toolbox. During this period the outside temperature of the building varies between -20 ^oC and 10.2 ^oC. The building has radiator fitted with motorized TRV and simple ON/OFF controller to control the TRV for maintaining the required flow rate ($m^{'}$) of hot water entering the radiator. The training data is recorded after every 30 seconds and both ANN and RNN controllers are trained with this data.

C. Artificial Neural Network Controller

Feed forward neural network (FFNN) consists of simple neuron like processing unit and is organised in layers. In this network all neurons in *i*th layer are connected to all neuron in (i-1)th layer. In FFNN, the learning rule tries to adjust the weights and biases of the network in order to move the network output closer to the target. The output of each neuron in the hidden layer and output layer is the result of non linear transfer function f represented by (19).

$$y_i = f(\sum_{j=1}^{m} w_{ij}x_j + b_i) \quad (19)$$

Where x is the input presented to the network, w are the weights of the network, b is the constant term which is referred as bias, and y is the output predicted by the network.
The intelligent ANN controller is three layered neural network model and it has three neurons in input layer, eight neurons in hidden layer and one neuron in the output layer. The ANN controller is trained by using gradient descent algorithm with learning rate η = 0.01. It took 20000 iterations by gradient descent algorithm to converge with minimum square error (MSE) of 2.3e-04.

D. Random Neural Network Controller

Similar to ANN controller the proposed RNN controller is three layer random neural network model with three neurons in the input layer, eight neurons in the hidden layer and one neuron in the output layer. When random neural network is trained by using gradient descent algorithm it took only 500 iterations to achieve the MSE of 1.6288e-06 at learning rate η = 0.01.

V. RESULTS AND DISCUSSION

The single zone building is simulated for 100 days for testing the performance of RNN and ANN controller. During this period of 100 days the outside temperature varies between -20.89 ^oC to 14 ^oC. The indoor air temperature of the building with RNN and ANN controller is shown in Figure 5. The heat supplied by the RNN controller is shown in Figure 6. Similarly, Figure 7 represents the heat supplied by the ANN controller. The variable heating set point for room temperature varies

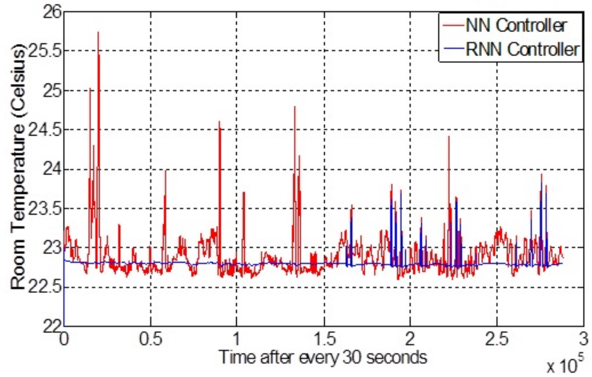

Figure 4. Indoor Air Temperature during testing

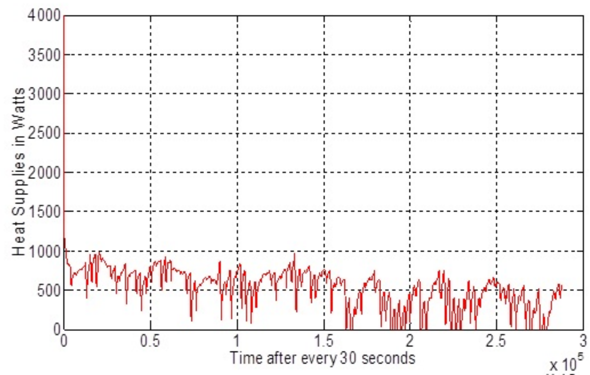

Figure 5. Heat Supplied by the RNN controller

between 22.66 ^oC to 23.00 ^oC in order to maintain the PMV of at least -0.3. As shown from Figure 5, ANN controller failed to keep the temperature within the specified range of variable set point while RNN controller achieved more accurate results than ANN controller.
In Figure 7, it is shown that ANN controller caused frequent oscillation between maximum and minimum flow rate as a result heat supplied to the rooms oscillate between 0 and 4000 watts. The RNN controller kept the stable flow rate due to which the heat suplied to the room didnt oscillate between minimum and maximum. In Table 2, the comparison of RNN and ANN in terms of MSE, no. of iterations, energy consumption, maximum overshoot and percentage of overshoot periods is given. The percentage of overshoot periods is percentage of instances when air temperature T_{air} exceeds the specified range of room temperature set point during 100 days simulation. From Table 2, it is shown that percentage of overshoot periods is only 4.27% for RNN controller while for ANN it is 45.96%. Similarly the energy consumption by heating system with RNN controller is 1282.4 MWh while with ANN controller energy consumption is 1292.5 MWh.

VI. CONCLUSION

In this paper, two variable set point intelligent heating system controllers are developed and their performances are compared for energy efficiency, and accuracy for maintaining the comfortable room temperature. To compare the performance of RNN and ANN controller, both controllers were trained

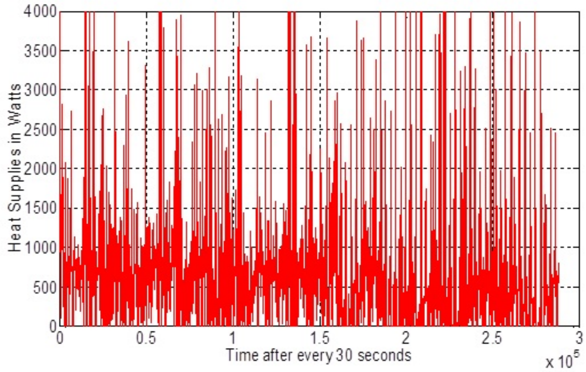

Figure 6. Heat Supplied by the ANN controller

TABLE II
COMPARISON OF RNN AND ANN CONTROLLER

	RNN	ANN
Learning Algorithm	Gradient Descent	Gradient Descent
No. of Iterations	500	20000
MSE	1.8266e-06	2.3 e-04
Energy Consumption 100 days (MWh)	1282.4	1292.5
Max Overshoot (^oC)	0.81	2.73
Percentage of overshoot periods	4.27%	45.90%

with same dataset and same training algorithm i.e., gradient descent algorithm. During training RNN showed impressive generalization capabilities and gradient descent algorithm for RNN converges in 500 iterations while gradient descent algorithm of ANN took 20000 iterations to converge. The RNN controller outperformed the ANN controller in testing phase where both controllers were tested for unknown data set. The heating system with RNN controller consumes 10 MWh less energy than with ANN controller. The RNN controller stopped the flowrate of hot water to the radiator by sensing the increase in outside temperature at correct time as a result percentage of overshoot periods is less compared to the ANN controller. The PMV index based variable set point control scheme ensured the comfortable indoor environment by suggesting the variable set points for maintaining PMV index of 0.3. The performance of RNN controller can further be improved by training the RNN with Levenberg Marquardt algorithm [27].

REFERENCES

[1] O. Publishing and I. E. Agency, Energy Technology Perspectives 2010: Scenarios and Strategies to 2050. Organisation for Economic Co-operation and Development, 2010.

[2] A. I. Dounis and C. Caraiscos, “Advanced control systems engineering for energy and comfort management in a building environment: A review,” Renewable and Sustainable Energy Reviews, vol. 13, no. 6, 2009, pp. 1246–1261.

[3] U. Parliament, Climate Change Act 2008. London: The Stationery Office Limited, 2008.

[4] Z. Liao, M. Swainson, and A. Dexter, “On the control of heating systems in the uk,” Building and Environment, vol. 40, no. 3, 2005, pp. 343–351.

[5] B. Xu, L. Fu, and H. Di, “Dynamic simulation of space heating systems with radiators controlled by trvs in buildings,” Energy and buildings, vol. 40, no. 9, 2008, pp. 1755–1764.

[6] S. Wang and Z. Ma, “Supervisory and optimal control of building hvac systems: A review,” HVAC&R Research, vol. 14, no. 1, 2008, pp. 3–32.

[7] S. A. Kalogirou, “Applications of artificial neural networks in energy systems,” Energy Conversion and Management, vol. 40, no. 10, 1999, pp. 1073–1087.

[8] M. Gouda, S. Danaher, and C. Underwood, “Application of an artificial neural network for modelling the thermal dynamics of a building’s space and its heating system,” Mathematical and Computer Modelling of Dynamical Systems, vol. 8, no. 3, 2002, pp. 333–344.

[9] S. A. Kalogirou and M. Bojic, “Artificial neural networks for the prediction of the energy consumption of a passive solar building,” Energy, vol. 25, no. 5, 2000, pp. 479–491.

[10] S. A. Kalogirou, “Artificial neural networks in renewable energy systems applications: a review,” Renewable and sustainable energy reviews, vol. 5, no. 4, 2001, pp. 373–401.

[11] A. A. Argiriou, I. Bellas-Velidis, and C. A. Balaras, “Development of a neural network heating controller for solar buildings,” Neural Networks, vol. 13, no. 7, 2000, pp. 811–820.

[12] S. A. Kalogirou, “Applications of artificial neural-networks for energy systems,” Applied Energy, vol. 67, no. 1, 2000, pp. 17–35.

[13] A. E. Ben-Nakhi and M. A. Mahmoud, “Energy conservation in buildings through efficient a/c control using neural networks,” Applied Energy, vol. 73, no. 1, 2002, pp. 5–23.

[14] J.-Y. Lee, M.-S. Yeo, and K.-W. Kim, “Predictive control of the radiant floor heating system in apartment buildings,” Journal of Asian Architecture and Building Engineering, vol. 1, no. 1, 2002, pp. 105–112.

[15] M. Gouda, S. Danaher, and C. Underwood, “Quasi-adaptive fuzzy heating control of solar buildings,” Building and Environment, vol. 41, no. 12, 2006, pp. 1881–1891.

[16] A. Ruano, E. Crispim, E. Conceição, and M. Lucio, “Prediction of building’s temperature using neural networks models,” Energy and Buildings, vol. 38, no. 6, 2006, pp. 682–694.

[17] J. W. Moon and J.-J. Kim, “Ann-based thermal control models for residential buildings,” Building and Environment, vol. 45, no. 7, 2010, pp. 1612–1625.

[18] P. Fanger, “Thermal comfort: analysis and applications in environmental engineering,” 1972.

[19] J. A. Orosa, “A new modelling methodology to control hvac systems,” Expert Systems with Applications, vol. 38, no. 4, 2011, pp. 4505–4513.

[20] E. Gelenbe, “Random neural networks with negative and positive signals and product form solution,” Neural computation, vol. 1, no. 4, 1989, pp. 502–510.

[21] ——, “Stability of the random neural network model,” Neural computation, vol. 2, no. 2, 1990, pp. 239–247.

[22] S. Timotheou, “The random neural network: a survey,” The computer journal, vol. 53, no. 3, 2010, pp. 251–267.

[23] C. Çerkez, I. Aybay, and U. Halici, “A digital neuron realization for the random neural network model,” in Neural Networks, 1997., International Conference on, vol. 2. IEEE, 1997, pp. 1000–1004.

[24] H. Abdelbaki, E. Gelenbe, and S. E. El-Khamy, “Analog hardware implementation of the random neural network model,” in Neural Networks, 2000. IJCNN 2000, Proceedings of the IEEE-INNS-ENNS International Joint Conference on, vol. 4. IEEE, 2000, pp. 197–201.

[25] H. Bakırcıoğlu and T. Koçak, “Survey of random neural network applications,” European journal of operational research, vol. 126, no. 2, 2000, pp. 319–330.

[26] E. Gelenbe, “Learning in the recurrent random neural network,” Neural Computation, vol. 5, no. 1, 1993, pp. 154–164.

[27] S. Basterrech, S. Mohammed, G. Rubino, and M. Soliman, “Levenbergmarquardt training algorithms for random neural networks,” The Computer Journal, vol. 54, no. 1, 2011, pp. 125–135.

Flow-based Routing Schemes for Minimizing Network Energy Consumption using OpenFlow

Shota Oda, Daiki Nobayashi, Yutaka Fukuda and Takeshi Ikenaga
Kyushu Institute of Technology
Fukuoka, Japan
Email: {n349405s@tobata.isc, nova@ecs, fukuda@isc, ike@ecs}.kyutech.ac.jp

***Abstract*—With the growing amount of traffic on the Internet, network devices are consuming increasing energy as their transmission capacity increases. Therefore, the reduction of network power consumption is an important issue. Software-defined networking is receiving significant attention as a means to manage increased network traffic efficiently. We consider energy efficient routing schemes with traffic aggregation for each flow that maintain acceptable performance using OpenFlow. The proposed schemes aggregate traffic to build a network using the minimum active devices necessary. Through simulation evaluation, we show that the proposed schemes can reduce power consumption, in comparison with the conventional routing protocol.**

Keywords-Routing; Energy-Efficiency; Green network; OpenFlow; QoS.

I. INTRODUCTION

Power saving technology for network devices has become an important issue due to growing demand for power conservation. In addition, Software-Defined Networking (SDN) is receiving significant attention as a means to manage increasing network traffic efficiently. OpenFlow, which enables flexible and centralized network management, plays a central role in SDN enabled routing for each flow and allows dynamic control of network routing and metrics.

Many researchers have proposed network power saving schemes [1][2][3], such as ElasticTree [4]. Some of these schemes control traffic aggregation with using network devices, effectively building a power saving sub-network. These schemes can reduce unnecessary power consumption by restructuring the network relative to traffic activity and requirements. However, these schemes have some inherent issues. For example, increasing wait time for optimization of all network flows and reduced responsiveness of new flows.

In this study, we consider a dynamic and energy efficient routing scheme that applies traffic aggregation for each flow while maintaining acceptable performance using OpenFlow. The proposed schemes can build a power saving sub-network by changing only the operating mode of nodes and links as required without restructuring the entire network. Simulation results indicate that the proposed schemes can reduce network power consumption, in comparison with Open Shortest Path First (OSPF) [5].

In Section II, we describe related works in this area. Section III shows the proposed scheme. Section IV shows the simulation model, and its results will be discussed in Section V. Finally, we describe our conclusion in Section VI.

TABLE I. POWER CONSUMPTION SUMMARY FOR ENTERPRISE SWITCHES

Configuration	Rack Switch(W)	Tier-2 switch(W)
Chassis	146	54
Linecard	0 (include in chassis)	39
10Mbps/port	0.12	0.42
100Mbps/port	0.18	0.48
1Gbps/port	0.87	0.90

II. RELATED STUDY

Existing routing protocols, such as OSPF or IS-IS [6] provide shortest path. Also, many researchers have proposed TE (Traffic Engineering) which uses effective network resource for achieving QoS. Many of TE approaches try to spread the load among multiple paths, so they do not consider the power saving. Therefore, we require the new routing scheme for achieving power saving without deteriorating communication quality.

Many researchers have proposed power saving schemes in a network area. In this section, we first show the power saving effect on the network. Then, power saving schemes for node and link are described. Finally, we show the power saving with traffic engineering.

A. Power saving effect

Mahadevan et al. measured the power consumption of network devices in their data center [7]. They measured power usage of each linecard and chassis, and of different transmission rate at a port. The obtained measurement is summarized in Table I. From this table, the power consumption of network devices such as switch or router is about 100 W, and of link is about 1.0 W. Therefore, we can expect power saving if the number of active nodes and links are dynamically adapted to the arriving traffic.

B. Node and Link state control for power saving

Gupta et al. proposed three policies for power saving by putting network components to sleep mode [8]. In the first policy, the device sets a sleep timer, and only wakes up when the timer expires. All packets arriving during the sleep period are lost. This approach can achieve stable power saving

 ISBN: 978-1-61208-332-2

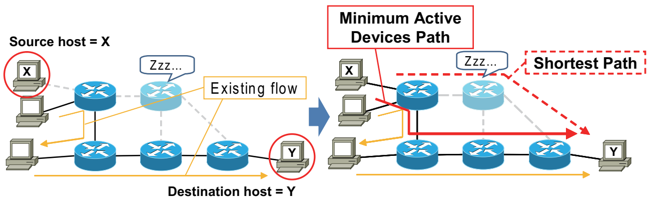

Figure 1. Overview of proposed schemes

performance, but cause performance degradation because of packet loss. The second is Hardware Assisted Sleep (HAS) policy in which the interface wakes up by an incoming packet; however, the packet is lost. The third is the Hardware Assisted Buffered Sleep (HABS) policy, to buffer the incoming packets while the interface is in the sleep mode. HABS consumes the power in the sleep state to buffer the incoming packets, but avoid the packet loss.

Uncoordinated sleeping informs its neighbor before network devices put into the sleep mode [9]. In this scheme, in order to avoid the packet loss, the neighbor hosts send a wake up packet to the hosts with sleep state. Nedevschi et al. proposed the power management schemes to reduce the energy consumption in the network [10]. Putting network elements to sleep mode during idle times and rate adaptation according to the arrival traffic are used. Their work focuses on edge routers in the Internet.

C. Traffic control for power saving

Heller et al. proposed the ElasticTree, which is the network wide power manager [4]. ElasticTree can dynamically adjust the number of active nodes and links. This scheme increases the more computational cost by growing the size of the network. Therefore, throughput is lower because of decreasing the processing speed of devices.

Vasic and Kostic proposed the Energy-Aware Traffic engineering (EATe), which reduces the power consumption while maintaining same traffic rate [11]. This scheme distributes traffic while periodically exchange link utilization between nodes. Therefore, we think that power saving performance is relatively low.

From the above, we propose a scheme that can achieve higher energy efficiency without deteriorating the throughput.

III. PROPOSED SCHEMES

In this section, we propose two schemes that consider not shortest paths but dynamic flow aggregation to increase sleep state in nodes and links. Overview of both proposed schemes is shown in Fig. 1. OpenFlow is used in the proposed schemes to manage network operations, allowing centralized and dynamic control of network routing and metrics.

A. Routing for Minimization of Active Devices

We refer to this scheme as Routing for Minimization of Active Devices (RMAD). This scheme selects a path to maximize the number of active state links from several shortest path for minimization of active links in the whole of middle or low load network. In the RMAD scheme, we use not only the shortest hop count but also number of active links as a routing metric. The number of active links is changed dynamically according to flow arrival using OpenFlow routing management. The procedure of RMAD scheme is shown in Fig. 2 and is as follows:

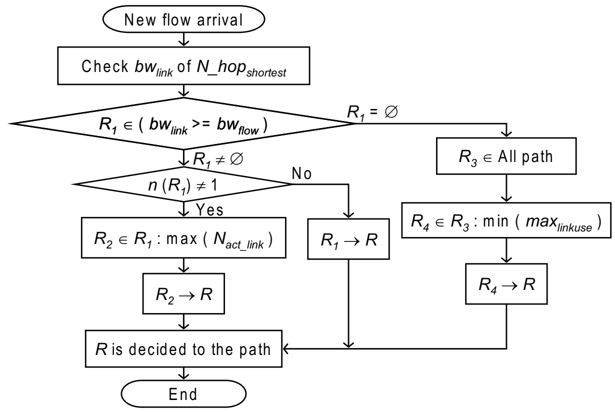

Figure 2. Procedure of RMAD scheme

(1) We obtain the shortest paths to the destination. Then, the paths of which residual bandwidth (bw_{link}) is larger than reguired bandwidth of new flow (bw_{flow}) are denoted as candidates R_1.
(2) If $R_1 \neq \phi$, we select the path that maximizes the number of active links (N_{act_link}) from R_1. If the chosen path includes the sleep state link, the link is set to active state.
(3) If $R_1 = \phi$, we select the path that minimizes the maximum link utilization ($max_{linkuse}$) on the path.

RMAD can aggregate new flows to the links of active state, so that the number of active links in the network is reduced. In addition, communication performance degradation is avoided by using shortest path and the path which minimizes the maximum link utilization on the path even if we cannot find the R_1.

B. Routing for Minimization of Active Devices plus

As previously mentioned, RMAD considers only active links to reduce power consumption. However, we can expect greater power savings by decreasing the number of active nodes using dynamic routing controlled by OpenFlow because node power consumption is greater than link power consumption [7]. We refer to this scheme as "RMAD+". The procedure of RMAD+ scheme is shown in Fig. 3 and is as follows:

(1) We obtain paths with hop counts that do not exceed the shortest path ($N\_hop_{shortest}$) plus δ. Next, the paths with residual bandwidth (bw_{link}) larger than the reguired bandwidth of new flow (bw_{flow}) are denoted as candidates R_1.
(2) If $R_1 \neq \phi$, we select the path R_2 that minimizes the number of active nodes in the entire network ($N\_all_{act\_node}$) from R_1.
(3) We select the path R_3 that minimizes the number of active links in the entire network ($N\_all_{act\_link}$) from R_2.

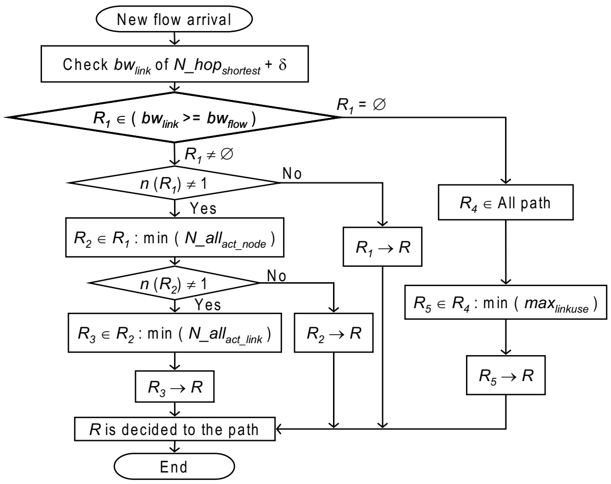

Figure 3. Procedure of RMAD+ scheme

(4) If $R_1 = \phi$, we select the path that minimizes the maximum link utilization ($max_{linkuse}$) on the path.

RMAD can aggregate new flows to active links, thus reducing the number of active links in the network. Moreover, RMAD+ considers the number of active nodes and concentrates flows on active paths. Therefore, we can expect greater power saving compared to RMAD. In addition, communication performance degradation is avoided by using the path that minimizes the maximum link utilization on the path even if we cannot find a satisfactory R_1. Furthermore, RMAD+ avoids paths with high link utilization, thus communication performance does not degrade significantly.

IV. SIMULATION MODEL

We evaluate the proposed schemes in terms of both data transfer and power saving performance. Simulation experiments were performed by using the network simulator, *ns-2*.

In our simulation, we used a multi rooted fat-tree topology, which is commonly used in many data centers. The multi rooted fat-tree topology consists of three-layer (i.e., Edge, Aggregation, and Core layers) on top of the end hosts. A pod of the fat-tree topology is a set of Edge and Aggregation routers. In our simulation, we use a 4-pod fat-tree topology with 16 end hosts. The average communication period follows the *exponential distribution*, and its average was set to α. In the data transmission, end hosts transmit data only during the *ON* time, and the *OFF* time is determined at the end of each data transmission according to the exponential distribution. Thus, the hosts are silent in the *OFF* time and start data transmission again after this interval. We defined the each *ON* time as a flow, and end hosts send Constant Bit Rate (CBR) packets at transmission rates of 500 Kb/s as a flow in each *ON* time. The network load ρ is defined by the following equation:

$$\rho = \frac{\text{Sum of } \alpha}{\text{Simulation Time}}$$

In this simulation, we vary the load factor ρ from 0.1 to 0.9.

In the simulation, each end host selected the destination randomly or specifically. In the first case, the destination was selected uniformly from the end hosts. In contrast, 70% of the flows in the second case were distributed among four end hosts in a single pod. Figure 4 shows an overview of case of concentrating flows. We conduct our simulation experiments for 20 seconds. The parameters used in our simulation are summarized in Table II.

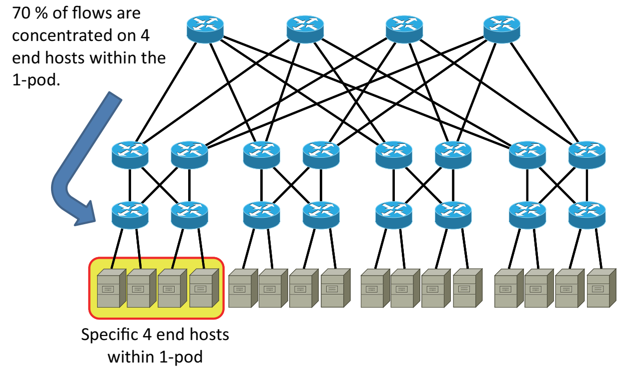

Figure 4. Simulation model of concentrating flows

TABLE II. SIMULATION PARAMETER

Simulation Time	20 s
Link Bandwidth	5.0 Mb/s
Link Delay	1.0 ms
Packet Size	1500 Byte
Transmission Rate	500 Kb/s
Average Time of Flow, α	1.0 s
Number of Flow Concurrent	1 / end hosts
Network load, ρ	0.1 $\sim$ 0.9

V. SIMULATION RESULTS

We compared the performance of RMAD and RMAD+ against OSPF in terms of both power efficiency and communication performance. To quantify the performance of the proposed schemes, we used Node Sleep Ratio, Average Transmission Time, and Packet Loss Ratio as performance measures. Node sleep ratio defined as the percentage of device sleep time indicates power saving efficiency, and the average transmission time and packet loss ratio indicate communication performance.

Figure 5 shows the node sleep ratio for the case in which flow destination was selected from the end hosts with uniform probability. Figure 6 shows the node sleep ratio for the case in which the flow destinations were concentrated on specific end hosts. In both figures, RMAD+ (+δ) shows the number of hops to allow.

For both simulated traffic patterns, we observed that the node sleep ratio of the RMAD and RMAD+ schemes was higher than OSPF under any load. In addition, the node sleep ratio of RMAD+ was higher than that of RMAD under low load when the flow destination was selected from the end hosts with uniform probability. Moreover, the node sleep ratio of RMAD+ was always higher than that of RMAD when the flow destinations ware concentrated on specific end hosts.

We also researched the average transmission time and the ratio of packet loss of each schemes in order to show the effect on communication performance. Figure 7 shows the average

 ISBN: 978-1-61208-332-2

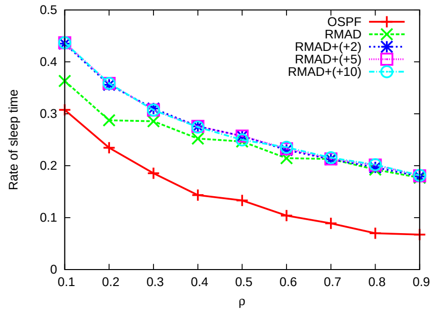

Figure 5. Node sleep ratio of uniform probability model

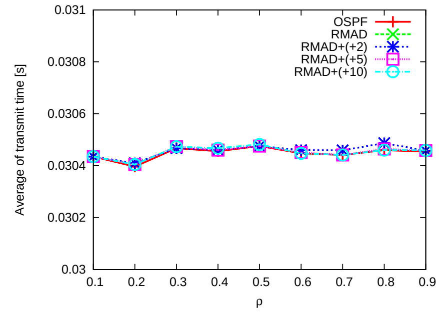

Figure 7. Average transmission time of uniform probability model

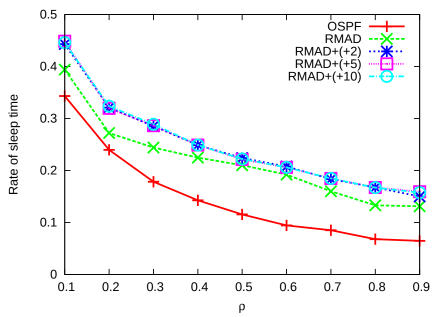

Figure 6. Node sleep ratio of concentrated flow model

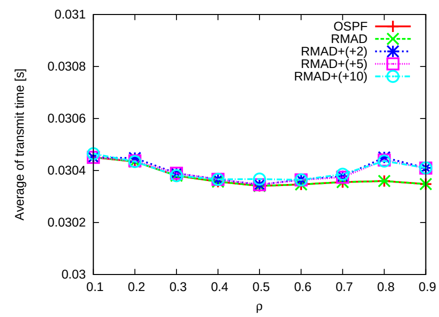

Figure 8. Average transmission time of concentrated flow model

transmission time for the case in which flow destination was selected from the end hosts with uniform probability. Figure 8 shows the average transmission time for the case in which the flow destinations were concentrated on specific end hosts. We observed that RMAD+'s average transmission time was longer than that of OSPF and RMAD. Furthermore, we did not observe any packet loss in the proposed schemes. If flow destinations are selected from specific end hosts, RMAD+ uses a concentrated path by considering power consumption rather than the shortest path. As a result, the node sleep ratio of RMAD+ was higher than that of OSPF and RMAD even under high load. However, the difference in average transmission time between the proposed schemes and OSPF is negligible. From these results, we conclude that the proposed schemes can achieve high energy efficiency without deteriorating communication quality.

VI. Conclusion and Future Work

We have proposed two energy efficient routing schemes that aggregate traffic on active network paths to decrease the number of active network devices while maintaining acceptable performance. By placing network devices without traffic into a sleep mode, power consumption in the network is reduced assuming that information management and path selection is performed by OpenFlow. The performance of the proposed schemes, RMAD and RMAD+ was compared against OSPF in terms of energy efficiency and communication quality. The simulation results indicate that the proposed schemes, on comparison with OSPF, can reduce power consumption without deteriorating communication quality. As a future work, we will compare proposed schemes with the existing energy efficient routing schemes. Also, we will evaluate the performance of the proposed schemes in various network environment.

Acknowledgement

This work was supported in part by the Japan Society for the Promotion of Science, Grand-in-Aid for Scientific Research (B) (No. 23300028)

References

[1] R. Bolla, R. Bruschi, F. Davoli, and F. Cucchietti, "Energy Efficiency in the Future Internet: A Survey of Existing Approaches and Trends in Energy-Aware Fixed Network Infrastructures," IEEE COMMUNICATIONS SURVEYS & TUTORIALS, vol. 13, May. 2011, pp. 223-244.

[2] L. Chiaraviglio, M. Mellia, and F. Neri, "Minimizing ISP Network Energy Cost: Formulation and Solutions," IEEE/ACM TRANSACTIONS ON NETWORKING, vol. 20, no. 2, Apr. 2012.

[3] N. Wang, C. Michael, and K. Hon Ho, "Disruption-Free Green Traffic Engineering with NotVia Fast Reroute," IEEE COMMUNICATIONS LETTERS, vol. 15, no. 10, Oct. 2011.

[4] B. Heller, S. Seetharaman, P. Mahadevan, Y. Yiakoumis, P. Sharma, S. Banerjee, and N. McKeown, "ElasticTree: Saving Energy in Data Center Networks," Proc. USENIX NSDI'10, Apr. 2010.

[5] J. Moy, "OSPF Version 2," IETF RFC 2328, Apr. 1998.

[6] R. Callon, "Use of OSI IS-IS for Routing in TCP/IP and Dual Environments," IETF RFC 1195, Dec. 1990.

[7] P. Mahadevan, P. Sharma, S. Banerjee and P. Ranganathan, "Energy Aware Network Operations," Proc. 28th IEEE international conference on Computer Communications Workshops, Apr. 2009, pp. 25-30, INFOCOM'09.

[8] M. Gupta, S. Grover, and S. Singh, "A Feasibility Study for Power Management in LAN switches," Proc. 12th IEEE International Conference on Network Protocols, Oct. 2004, pp. 361-371.

[9] M. Gupta and S. Singh, "Greening of the Internet," Proc. ACM SICCOMM '03, Aug. 2003, pp. 19-26.

[10] S. Nedevschi, L. Popa, G. Iannaccone, S. Ratnasamy, and D. Wetherall, "Reducing Network Energy Consumption via Sleeping and Rate-Adaptation," Proc. USENIX NSDI'08, Apr. 2008.

[11] N. Vasi'c and D. Kosti'c, "Energy-Aware Traffic Engineering," Proc. 1st International Conference on Energy-Efficient Computing and Networking, Apr. 2010, pp. 169-178.

Keeping High Availability of Connected End-point Devices in Smart Grid

Hiroshi Isozaki[1, 2], Jun Kanai[1]
[1] Corporate R&D Center, Toshiba Corporation,
Kawasaki, Kanagawa, Japan
{hiroshi.isozaki, jun.kanai}@toshiba.co.jp
[2] Graduate School of Media and Governance,
Keio University,
Fujisawa, Kanagawa, Japan

Shunsuke Sasaki, and Shintarou Sano
Center for Semiconductor Research & Development,
Toshiba Corporation,
Kawasaki, Kanagawa, Japan
{shunsuke.sasaki, shintarou.sano}@toshiba.co.jp

***Abstract*— Security is an important aspect to achieve Smart Grid success in terms of commercial deployment. Particularly, availability gets the highest priority in Smart Grids. For end-point devices, such as smart meters or concentrators, this must be true since they must always be working. We present LiSTEE™ Recovery, an architecture for a fault tolerant system for end-point devices to monitor the status of the operating system and to recover even if they stop working due to unexpected behavior or cyber attack including zero-day attack. LiSTEE™ Recovery provides further functions to prevent illegitimate memory modification and to notify a head-end system once a security incident occurs. We demonstrate a full implementation of LiSTEE™ Recovery on a TrustZone capable ARM based processor. Our experiment shows that the performance degradation is small enough to be ignored. Furthermore, we observed that the cost of production and maintenance can be minimized.**

Keywords-Smart Grid, Smart Meter, Concentrator, Security, High Availability, TrustZone

I. Introduction

In Smart Grids, requirements about supporting various protocols and functions to network connected end-point devices, such as smart meters or concentrators, make their systems more complicated. Because a large quantity of source code is necessary to implement a complicated system in general, the risk of including vulnerability in the system increases. Moreover, since the devices are connected to home networks, the risk of devices of being attacked is high compared with legacy devices connected to managed network only. In fact, it is reported that smart meters from a variety of vendors were found to improperly handle malformed requests which could be exploited to cause buffer overflow vulnerability; allowing an attacker to cause a system to become unstable or freeze [1]. To keep devices secure under this situation, many security protocols and algorithms have been proposed to securely distribute a shared key between devices and head-end systems or to store privacy data in devices in a secure manner [2][3]. However, confidentiality and integrity are not enough to solve the security problem in Smart Grids. It is strongly desired for the devices to keep high availability since they must always be working to provide demand response service or to use consumption data for payment [4]. As only one vulnerability may cause the system to go down, it is very difficult to keep high availability in a complicated system. Furthermore, unlike interactive devices, such as PC or smart phone, it is difficult to expect that end users reset and restart devices once they freeze or hang since end users cannot recognize the status of the devices and cannot determine the device should be rebooted or not. Thus, how to keep the availability of the devices is a significant challenge in Smart Grids.

To address these problems, we propose LiSTEE™ Recovery, an architecture for fault tolerant systems which automatically recovers from error status. To achieve this goal, LiSTEE™ Recovery isolates a surveillance process observing the state of the system and recovery process which reboots the system when it detects the system freezes. In the LiSTEE™ Recovery, surveillance and recovery processes run in an isolated secure environment while general purpose processes, including operating system, such as network or storage access run in a non-secure environment with hardware access control performed with respect to memory. Hence, a memory area where surveillance and recovery processes are arranged cannot be accessed by general purpose processes. As a result, even if the operating system is attacked and crashes, it becomes possible to prevent from interference in the surveillance and recovery processes.

The remainder of this paper is organized as follows. In Section-II, problems are defined. Section-III indicates background information. Section-IV and V propose framework and implementation of LiSTEE™ Recovery. The evaluation is shown in the Section-VI; and the conclusion and future work are in at the end.

II. Problem Definition

In a legacy system, surveillance and recovery processes and their execution environment are monolithically configured. In other words, the reliability of surveillance and recovery processes depends on the reliability of their execution environment. In order to keep reliability high, a system needs to be implemented without vulnerability. In order to detect and eliminate vulnerability in source code, various testing methods have been proposed [5][6]. However, since end-point devices will be deployed without maintenance over a long period of time within Smart Grids, there is a large risk such devices continue operating without vulnerabilities being fixed even if those devices had no vulnerabilities at the time of shipping. Attackers may exploit the vulnerability, such as buffer overflow or malformed network input, in order to and cause the device to crash. To make matters worse, attackers are in a somewhat advantageous position in launching a large attack since the number of device vendors is limited and the software installed in the devices is uniform. Furthermore, attackers can reverse-engineer code without administrators noticing in order to find a vulnerability since, unlike a server application, devices are located at the user side. Therefore, when attackers find one vulnerability in a single device, they can exploit it on a lot of devices. Considering the above situation, the following problems are to be solved in order to keep high availability under a legacy system.

A. *Difficult to Keep a High Level of Surveillance Continuity*

In order to implement a complicated application program or a minor network protocol on the end point device, Linux will be used as a software execution environment. In Linux, the surveillance and recovery processes can be implemented as a user task executed on the operating system or as an interrupt handler in the operating system. When a surveillance target process is implemented as a user task running on the operating system then support functions in the operating system, such as the "cron" service in Linux, can be used to detect a failure of the user task and to automatically restart the target process. When the surveillance process is implemented as an interrupt handler in the operating system then sophisticated implementation is necessary compared to an application program; it is automatically and periodically called by a timer interrupt as long as the operating system works. Another legacy approach is implementing a monitoring and detecting mechanism in the operating system. For example, in order to find buffer overflow attacks, a protection element monitors system call frequencies, and if the frequencies are different from normal behavior, it can detect the attack [7]. However, the fundamental problem of a legacy approach is that there is no way to restart the process if the operating system itself crashes for some reason. Furthermore, the protection mechanism itself could be a target of the attack, as the result the protection mechanism could be invalidated. Thus, there is a large risk where devices in a Smart Grid breakdown and the attack may be able to cause a blackout to vast areas in the worst case. In order to prevent devices breaking down, they are required to provide a robust method to recover the system from failure in order to keep high level of availability. Still there are some existing hardware devices supporting a watchdog timer function which detects the status of the operating system and automatically reboots the system. Since not all devices support the function and it is difficult to implement complicated functions as described below inside it, a new approach is desired. To clarify the conditions, only a software failure including an attack is assumed in this paper. A physical fault, such as a hardware failure or loss of power, or a hardware attack, such as physically destroying devices or cutting cables are out of scope in this paper.

B. *Difficult for an Administrator to Detect when Incident Occurs*

End-point devices are connected with a head-end system through the network to provide a demand response service. When the devices detect an error status, such as a surveillance target process being stopped for an unknown reason, it is desirable for these devices to send a report to the head-end system so that an administrator can realize the situation and use the report to investigate the reason for the failure. However, for the same reason as described above, there is no way for devices to send a message to the head-end system if the operating system crashes. Even in such a case, it is desirable for devices to provide a method to send a message to acknowledge the error situation to the system administrator. Besides notification of the error situation to the system administrator, software update function is also desirable. However, since many existing hardware devices have already supported a secure firmware update function and its method highly depends on each device, it is out of scope in this paper.

In addition to the problem described above, the following business problem needs to be considered when introducing a new architecture to the market.

C. *Development and Production Cost*

Cost is an important aspect in evaluating the proposed security architecture. When implementing an end-point device, if the new security architecture requires a complete rebuild of software, the architecture will never be introduced to the market. Thus, it is desirable to reuse existing software asset, such as libraries, middleware and applications as much as possible to minimize the development cost including verification cost. Specifically in Smart Grids, the verification cost is large since reliability is strongly required. Besides the development cost, we need to consider the cost per device. One approach to solve the problems described above is to utilize the dedicated hardware security chip. However, since the chip is sometimes very expensive, it causes a rise of production cost per device. Therefore, it is also desirable to use widely available existing commodity hardware to minimize production cost.

III. BACKGROUND (TRUSTZONE)

In this section, we provide background information on the hardware technologies leveraged by LiSTEE™ Recovery.

TrustZone is a hardware security function supported by a part of ARM processor [8][9]. General ARM processor defines two modes, user mode and privileged mode. In privileged mode, execution of all instructions and access to all memory regions are allowed while in user mode availability of instructions and accessibility of memory regions are restricted. In general system, operating system is executed in privileged mode while application programs are executed in user mode. In addition to the two modes, a TrustZone enabled ARM processor supports two worlds which are independent of the modes. One is secure world for the security process and the other is non-secure world for everything else. Fig. 1 shows the relationship between worlds and modes conceptually. The processor is executed by selectively switching the worlds if needed. For example, it is assumed that key calculation process is executed in secure world while all other general processes, such as

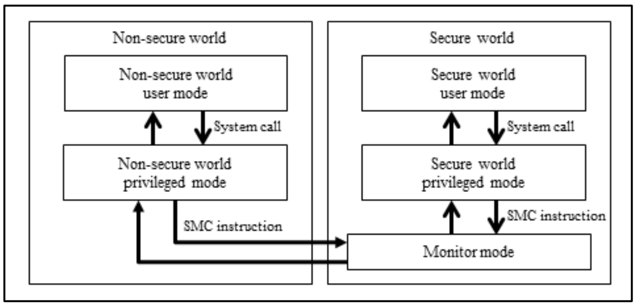

Figure 1. Mode and world in ARM.

storage access or network access are executed in non-secure world. Moreover, using TrustZone enable hardware, it is possible to make a system where a process running in secure world can access all system resources while a process running in non-secure world can access a part of system resources only. For example, when used in combination with the TrustZone Address Space Controller (TZASC), access to memory can be restricted for a process running in non-secure world even if the process runs in privileged mode to install an access control policy on TZASC.

TrustZone provides a dedicated instruction, the Secure Monitor Call (SMC) instruction, to transit between the worlds. As soon as the SMC instruction is called, the processor switches into monitor mode. A software program running in monitor mode saves a context of the program running in the current world on the memory and restores a context of the program running in the previous world, then changes the world, and finally executes the program running in the previous world. Besides the SMC instruction, hardware exceptions can be configured to cause the processor to switch into monitor mode.

IV. FRAMEWORK OF LISTEE^{TM} RECOVERY

LiSTEE^{TM} Recovery provides a method for an end-point device to automatically recover from an error status. It also provides a high level of memory protection mechanism. Hence, the recovery process is securely executed without interference. Fig. 2 shows the entire architecture of LiSTEE^{TM} Recovery. LiSTEE^{TM} Recovery consists of three components, Normal OS, LiSTEE^{TM} Tracker Application (LiSTEE^{TM} TA), and LiSTEE^{TM} Monitor.

- Normal OS: Operating system which executes eneral purpose processes, such as storage access or network communication. It is executed in non-secure world. All applications implementing smart meter functions or concentrator functions run on this operating system.
- LiSTEE^{TM} Tracker Application (LiSTEE^{TM} TA): Surveillance and recovery processes executed in secure world. LiSTEE^{TM} TA includes three modules: Watcher module, Recovery module, and Notification module. The Watcher module is an entry point of LiSTEE^{TM} TA. It is executed periodically by a timer interrupt through LiSTEE^{TM} Monitor. Whenever it is called, it investigates the status of Normal OS. If it detects Normal OS is not working, it calls Recovery module to reboot the system. Otherwise, it calls SMC instruction to switch to Normal OS. Moreover, the Notification module is called before Recovery module reboots the system. It sends a message to notify that the system is about to reboot to the head-end system through network.
- LiSTEE^{TM} Monitor: LiSTEE^{TM} Monitor is a program running in the monitor mode. It initializes configurations of TrustZone related hardware when booting the system. It also provide context switching function between worlds in hardware interrupt handler and SMC handler. Regarding the configuration of the hardware, it sets configuration register of timer interrupt so that hardware interrupt handler of LiSTEE^{TM} Monitor is called when timer interrupt is generated. Moreover, LiSTEE^{TM} Monitor manages the access control policy and installs the policy on TZASC. Regarding context switching, the SMC handler in LiSTEE^{TM} Monitor is executed when the SMC instruction is called and it transits from secure world to non-secure world. In contrast to the SMC handler, timer interrupt triggers transit from non-secure world to secure world based on the initializing configuration.

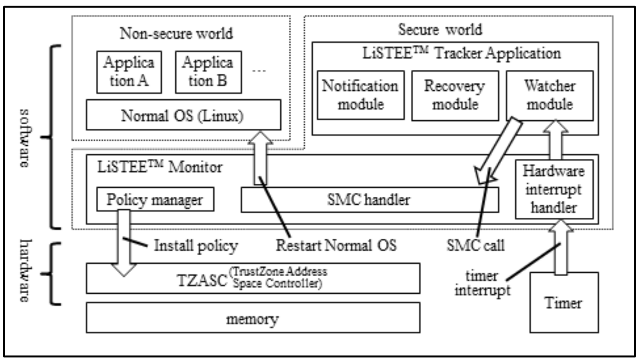

Figure 2. System Architecture of LiSTEE^{TM} Recovery.

The primary feature of LiSTEE^{TM} Recovery is to provide a method for the end-point device to detect the status of Normal OS and to recover it even if Normal OS crashes or stops working. Furthermore, it provides two additional functions. One is to enhance the security protection for LiSTEE^{TM} Monitor, LiSTEE^{TM} TA and Normal OS against attacks. The other is sending a message to the head-end system when an incident occurs. The details of these functions are described below.

A. *Periodical Surveillance and Recovery*

When booting the system, LiSTEE^{TM} Monitor is executed after executing initial program. Then, LiSTEE^{TM} Monitor loads and executes LiSTEE^{TM} TA and Normal OS respectively. While executing Normal OS, whenever the timer interrupt occurs, the processor jumps to the hardware interrupt handler in LiSTEE^{TM} Monitor. The hardware interrupt handler context switches from non-secure world to secure world and calls LiSTEE^{TM} TA. Specifically LiSTEE^{TM} Monitor saves a context of Normal OS to

Copyright (c) IARIA, 2014. ISBN: 978-1-61208-332-2

memory and restores a context of LiSTEE™ TA, then changes the world and finally calls the Watcher module of LiSTEE™ TA. The Watcher module checks the status of Normal OS. If it judges that Normal OS is not working, the Watcher module calls the Recovery module which reboots the system. Otherwise it calls SMC instruction. Then, SMC handler in the LiSTEE™ Monitor is executed. It first context switches from LiSTEE™ TA to Normal OS, and restarts Normal OS at the point just before the timer interrupt occurred. While executing LiSTEE™ Monitor and LiSTEE™ TA, the execution of Normal OS is suspending. That is, Normal OS continues to be processed as if nothing was executed during the execution of LiSTEE™ TA. Fig. 3 shows the flowchart of periodical surveillance and recovery process.

There are many ways for the Watcher module to determine whether Normal OS is working or not. One of the methods is to check the data area of Normal OS. In general, when an operating system is working, there must be a certain data area which is updated regularly. By checking this data area, it is possible for the Watcher module to judge whether Normal OS is working or not.

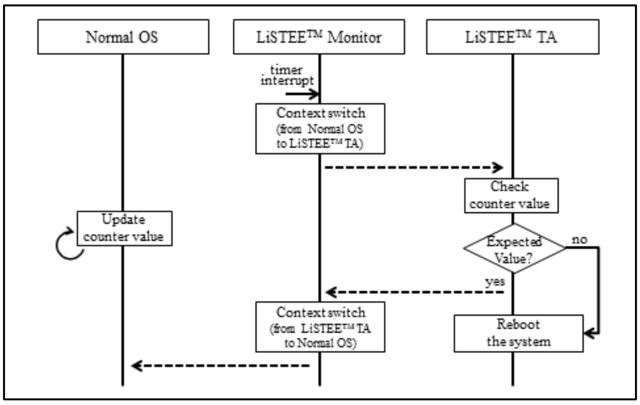

Figure 3. Flowchart of periodical surveillance and recovery.

B. Memory Protection

LiSTEE™ Recovery provides two memory protection mechanisms. Fig. 4 shows how these memory protection mechanisms work. One is protection for the kernel area of Normal OS. To realize this protection, LiSTEE™ Monitor

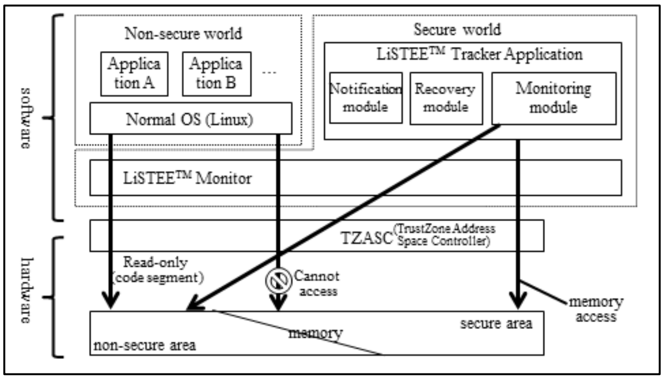

Figure 4. Memory protection mechanism.

provides virtual read-only memory. In virtual read-only memory, Normal OS running in non-secure world cannot overwrite the content on the memory while LiSTEE™ TA and LiSTEE™ Monitor running in secure world can access it using ordinary SRAM as the memory which is, of course, physically writable memory. In general, when a program is loaded into memory, a data region (data segment) and a code region (code segment) are assigned. In the initial state before booting the system, all regions are allowed to be accessed from Non-secure world by default. In order to allow the boot loader to write the code segment into the memory, LiSTEE™ Monitor leaves the memory region as is until the code segment is loaded. Just after executing the kernel of Normal OS, LiSTEE™ Monitor sets the memory region as read-only for kernel code segment of Normal OS. As the result, even Normal OS is prohibited from overwriting its own code segment.

Another mechanism is protection for LiSTEE™ Monitor and LiSTEE™ TA. To realize this protection, LiSTEE™ Monitor installs an access control policy where Normal OS cannot access the memory area allocated to LiSTEE™ Monitor and LiSTEE™ TA while LiSTEE™ TA and LiSTEE™ Monitor can access all areas when booting the system. For this policy, LiSTEE™ Monitor and LiSTEE™ TA can be protected from illegitimate falsification by Normal OS, even if Normal OS is attacked and under control of an attacker.

C. Message Notification

LiSTEE™ Recovery provides a function to notify the head-end system that Normal OS has stopped working and is rebooting the system by sending a message through network. The Notification module has the role of sending a message. Although Normal OS has network connectivity function, such as TCP/IP stack, LiSTEE™ TA cannot use the function since it is not working when sending a message. Thus, LiSTEE™ TA supports network connectivity function including network application, network protocol stack and network driver.

V. PROTOTYPE IMPLEMENTATION

We used ARM C/C++ Compliler 5.01 as a compiler to build LiSTEE™ Monitor and LiSTEE™ TA. We used gcc 4.4.1 to build Linux 3.6.1 as Normal OS. We chose Motherboard Express uATX with CoreTile Express A9x4 processor which supports TrustZone as an execution environment.

Regarding memory map, from 0x48000000 through 0x4A000000 is assigned for SRAM, and from 0x60000000 through 0xE0000000 is assigned for DRAM. Tab. I shows memory map with access control policy of the memory. In Tab. I, Normal OS (code) indicates Linux kernel code. Normal OS (data) includes Linux data, application code and application data. For clarification, Read/Write access is applied from Non-secure world for the area not described in Tab. I.

In order for the LiSTEE™ Monitor to install an access control policy on TZASC, the start address and the size of each memory region are predefined. After the boot loader loads Linux at the predefined value, LiSTEE™ Monitor installs the access control policy on TZASC. As shown in Tab. 1 the access to the memory regions allocated to LiSTEE™ Monitor, LiSTEE™ TA and code segment of Normal OS is restricted for the Normal OS running in non-secure world while the access to the region allocated to the data segment of Normal OS and shared memory is not. For clarification, LiSTEE™ Monitor and LiSTEE™ TA run in secure world can access all regions. Furthermore, since LiSTEE™ Monitor sets the configuration registers of TZASC to prohibit Normal OS from accessing them, Normal OS cannot change this configuration.

TABLE I. MEMORY MAP

Data	Start Address	Size	Access Policy (Non-sec world)
Vector table + Initialization code + LiSTEE™ Monitor + LiSTEE™ TA	0x48000000	0x01B00000	Cannot Access
Normal OS (code)	0x60000000	0x002FE000	Read Only
Normal OS (data)	0x602FE000	0x3EF02000	Read/Write
Shared memory	0x9F200000	0x00C00000	Read/Write

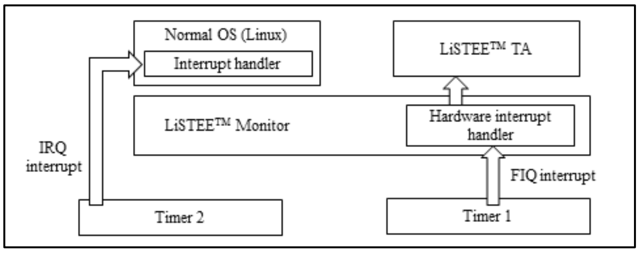

Figure 5. Assignment of timer interrupt.

Fig. 5 shows the assignment of the timer interrupt. We allocated a timer interrupt caused by timer (timer 1) to Fast Interrupt Request (FIQ) and timer interval was set to 1 second. The FIQ interrupt is handled by hardware interrupt handler in LiSTEE™ Monitor, and then it calls LiSTEE™ TA, as the result LiSTEE™ TA is periodically called. We used another timer (timer 2) and allocated to Interrupt Request (IRQ), and timer interval was set to 4 milliseconds. The IRQ interrupt is handled by interrupt handler in Linux. Since Linux assumes timer interrupt is allocated to IRQ, it is not necessary to modify the Linux source code to adopt LiSTEE™ Monitor. Tab. II shows a configuration of hardware interrupt. We configured Secure Configuration Register (SCR) and Current Program Status Register (CPSR) so that FIQ handler of LiSTEE™ Monitor is called when FIQ handlers occurs while IRQ handler in Linux is called when IRQ interrupt occurs during executing Linux. In the same manner as the configuration register of TZASC, we set FIQ and IRQ configuration registers to prohibit Normal OS from accessing them. Thus, Normal OS cannot change the configurations. For example, Normal OS cannot mask FIQ to stop timer 1, or cannot change the timer interval of timer 1.

In order to determine whether Linux is working or not, we made a small application program which runs on Linux and communicates with LiSTEE™ TA. Shared memory is used to exchange data between LiSTEE™ TA and Normal OS. The application program writes a counter value into the shared memory periodically. Then LiSTEE™ TA reads the counter value from the shared memory. When Normal OS is crashed, the application program cannot update the counter value. If the counter value is not updated in certain amount of time or the counter value is not an expected value, LiSTEE™ TA determines that Normal OS is not working.

TABLE II. RELASHONSHIP BETWEEN WORLD AND INTERRUPT

World when interrupt occurs	Interrupt	Jumps to
Non-secure world	FIQ	Hardware interrupt handler (FIQ handler) in LiSTEE™ Monitor
	IRQ	IRQ handler in Normal OS (Linux)
Secure world	FIQ	Pending FIQ
	IRQ	Pending IRQ

When LiSTEE™ TA determines that Linux is not working, it sends the head-end sever a message. In order to send a message to the head-end system when LiSTEE™ TA detects that Linux is not working, we ported a network driver and UDP/IP stack to LiSTEE™ TA. We defined a proprietary protocol and data format over UDP to notify the head-end system that LiSTEE™ TA starts reboot the system.

VI. EVALUATION

In this section, we describe the result of the evaluation in terms of security to verify the problems of the legacy system defined in Section-II can be solved. Performance and cost analysis of LiSTEE™ Recovery is also described below.

A. Security Analysys

1) Surveillance and Recovery: LiSTEE™ Recovery can recover from a failure to reboot the system even if Normal OS crashes. The reason for the crash could be a software bug or a cyber attack including zero-day attack caused by unknown vulnerabilities. In either case, since hardware timer interrupt continues working regardless of the state of Normal OS, LiSTEE™ TA is always periodically called and can detect a failure of Normal OS. At the next level, it is desirable to detect the failure as soon as possible. Detection time depends on how frequently LiSTEE™ TA checks the status of Normal OS. Since the execution time of LiSTEE™ TA and context switching by LiSTEE™ Monitor is very short, LiSTEE™ Recovery can detect the crash of Normal OS very quickly. Some attackers may continue to attack just after rebooting the system. One possible approach to a countermeasure for the attack is to let LiSTEE™ TA have the minimum function like the "safe mode", we have not implemented that though.

2) Attack Prevention: The proposed system provides two levels of attack prevention mechanism. The first level is to prevent Normal OS from illegitimate modification. When an attacker gains full control of Normal OS to misuse the vulnerability, the attacker may overwrite the code segment of Normal OS to directly overwrite the memory. In fact, many vulnerabilities (e.g., CVE-2013-4342, CVE-2013-1969, and CVE-2008-1673) allowing a remote attacker to execute arbitrary code are reported [10]. In the case of Linux, for example, once arbitrary code is executed with an administrator privilege by an attacker, it is possible for the attacker to overwrite an arbitrary area of code segment through /dev/mem, resulting in system crash or misbehavior. Overwriting the code segment in memory is difficult in general though it is relatively easy in the case of end-point devices since hardware configuration is fixed. As the result, system may go down. However, since LiSTEE™ Monitor sets the access control of the memory region for the code segment of Normal OS as read-only, and its configuration can be changed only from secure world, it is impossible for Normal OS to overwrite the code segment of Normal OS. An advantage is the protection does not cause any side effects. Since data segment is used to store the state of the program, Normal OS updates the content of data segment frequently during its execution. In contrast to the data segment, since code segment is used to store program code, it is not expected to update its content after booting the system. Particularly because devices, such as smart meters or concentrators are not expected to change their function after being deployed, the dynamic update function is not required. Thus, this protection mechanism can protect Normal OS from illegitimate modification without side effects. The second level is to protect LiSTEE™ Monitor and LiSTEE™ TA from illegitimate modification and suspension. Since the first level of protection is effective on code segment of Normal OS only, an attack which overwrites a data segement cannot be prevented. Thus, there are still possibilities that control of Normal OS is gained by an attacker. Even in such cases, thanks to TZASC, since Normal OS is prohibited from overwriting the content of memory where LiSTEE™ TA and LiSTEE™ Monitor are allocated, illegitimate modification is prevented. Moreover, since the interrupt configuration register is accessible only from secure world, there is no way for Normal OS to stop the timer interrupt.

3) System Reliability: In a legacy system, one single bug could affect the entire system to cause a critical failiure. Considering a defensive viewpoint, the entire system including operating system must be bug free to archieve high availability ideally. However, it is not practical to build a system without bugs in a complicated system. In fact, Linux 3.6.1 consists of over 15 millions of lines of code and a lot of new bugs causing critical crash are reported frequently (e.g., CVE-2013-4563, CVE-2013-4387, and CVE-2012-2127) even though it is carefully reviewed by many professionals [10]. Thus, the smaller the critical component that has to be robust within a system, the better. In the case of LiSTEE™ Recovery, the critical components corresponds to LiSTEE™ TA and LiSTEE™ Monitor. In contrast to Linux, the code size of LiSTEE™ Monitor and LiSTEE™ TA is relatively small. In fact the volume of source code for LiSTEE™ Monitor is about 700 lines and its code and data size are 2.1 KB and 1.6 KB respectively. Similarly the volume of source code of LiSTEE™ TA is about 41200 lines and its code and data size are 1.09 MB. Compared to the volume of source code of Linux, the risk where LiSTEE™ Monitor and LiSTEE™ TA includes bugs is small.

4) Response to Failure: The Notification module in LiSTEE™ TA sends a message to the head-end server just before rebooting the system. The message, which tells which the particular devices are about to reboot, is sometimes useful information for administartors. For example, if messages are sent by devices having a particular software version number, the reboot could be caused by an attack which aims at the vulnerability specific to the software. If messages are sent by devices located in one particular network, the reboot could be caused by a network worm distributed in the specific network. Although LiSTEE™ Recovery cannot prevent an attack in advance, the notification feature can help the administrator to investigate the reason of the failure during or after the incident. The attackers try to block sending the message to circumvent the notification. However, Normal OS cannot interfere with Notification module in sending a message to the head-end server since Notification module is executed inside LiSTEE™ TA. Moreover, since LiSTEE™ TA is processed in an isolated environment from Normal OS, security processes, such as encrypting a message, are easy to implement in LiSTEE™ TA. In the next step, it is possible to include a firmware update feature to implement functions receiving data from the head-end system and writing the data into the file system to extend the function of Notification module. In combination with “safe mode” described above, this function is effective against a countinuous attack which occurs just after the system recovers.

B. Performance Analysis

As well as the implementation environment, we used Motherboard Express uATX which contains ARM Cortex-A9x4 processor running at 400MHz as an experimental environment. The size of level 1 instruction cache, level 1 data cache, and level 2 cache are 32 KB, 32 KB, and 512 KB respectively. It contains 1 GB DRAM as the main memory and we assigned the same memory map described in Section-V.

First, we measured the execution time of LiSTEE™ TA during execution of Normal OS. Precisely, the time period from the beginning of hardware interrupt handler in LiSTEE™ Monitor through to the execution of the SMC instruction. Without calling the Notification module, the average time is 1.7 microseconds over 10,000 trials. However, if the Notification module is called, the average time is 4.1 milliseconds over 10,000 trials. Note that the Notification module is called when rebooting the system which is rarely occurs. Thus, this performance overhead is no problem.

Next, we measured the performance degradation of Normal OS. Since the execution of Normal OS is suspended during execution of LiSTEE™ TA, the performance of Normal OS degrades in any case. The total time of Normal OS suspension time depends on the frequency where LiSTEE™ TA is called. There is a tradeoff between the performance degradation of Normal OS and the delay for detecting the crash of Normal OS. When the frequency is increased, the performance degradation of Normal OS is also increased. On the other hand, when the frequency is decreased, the delay for detecting the crash of Normal OS becomes larger. To measure the performance degradation, we used dhrystone as a benchmark program [11].

Fig. 6 shows the result of the experiment. The bar graph shows dhrystone score and the line graph shows the performance degradation. The performance is better as the score value is high. Each bar shows timer interval which LiSTEE™ TA is called and its value is default (never called), 5 seconds, 1 second, 0.2 seconds and 0.04 seconds respectively. When timer interval was set to 5 seconds, the performance degradation was suppressed within 0.001 %. Even if the interval was set to 0.04 seconds, the performance degradation was less than 0.2 %. The result shows that although there is a trade-off between performance degradation of Normal OS and detection rate logically, the performance degradation can be ignored in practical even if the frequency where LiSTEE™ TA is called is increased.

C. Cost Analysis

1) Development Cost: LiSTEE™ Recovery does not require any modification to Linux to run it as Normal OS on LiSTEE Monitor. Thus, in terms of application developer's cost, since developpers can reuse all existing programs including libraries, middlewares, and applications running on Linux, no additional develping cost is necessary.

2) Production Cost: LiSTEE™ Recovery is software based technology and no additional hardware except TrustZone capable ARM processor and address space controller is required. Today TrustZone capable processors are widely available. In fact, all ARM Cortex A series processors support TrustZone. Therefore, the additional cost is mitigated. As the result, developing cost per device can be minimized.

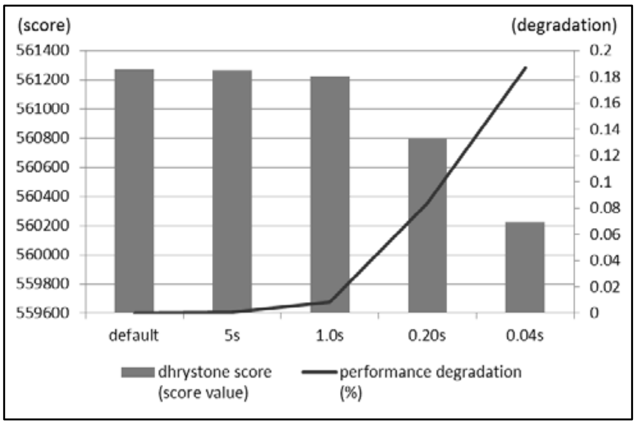

Figure 6. Result of performance degradation.

3) Maintainance Cost: It is assumed that tremendous number of devices are deployed in the field in Smart Grid. Specifically in the case of cyber attack, since many devices could be a target of the attack and the attack could be done in a very short period of time through network, it is not practical for field service engineers to physically visit each site and reboot them in terms of both cost and time. The auto recovery feature of LiSTEE™ Recovery mitigates this problem. Moreover, the report is sent to the head-end server once the device reboots. This function contributes in reducing the cost of trouble shooting. Thus, LiSTEE™ Recovery provides opportunity to reduce maintainance cost compared with legacy systems.

VII. RELATED WORK

To recover from an operating system failure, various approaches have been proposed.

The simplest approach is including the recovery mechanism inside operating system. One method is to use NMI as a watchdog timer [12]. Non-maskable Interrupt (NMI) is a processor interrupt that cannot be ignored. When NMI is generated, NMI handler implemented inside operating system is called regardless of the status of the operating system. Thus, NMI can be used as surveillance and recovery process to implement NMI handler so that it detects whether operating system hangs or not. Although NMI is easy to use for watchdog timer as it has already been implemented in Linux, it is vulnerable because NMI handler could be invalidated to overwrite the code segment of the operating system. Furthermore, it is not anticipated to implement a rich application in an interrupt handler, such as network communication function or data encryption function, it is difficult to realize the notification function.

Another approach to recover from the failure is to check the status of the operating system from outside using virtualization technology. It is easy to realize resource isolation environment by utilizing virtualization technology. Karfinkel developed trusted virtual machine monitor (TVMM), on which general-purpose platform and special-purpose platform executing security sensitive process run separately and concurrently [13]. The libvirt project develops a virtualization abstraction layer including a virtual hardware watchdog device [14]. To cooperate with the watchdog daemon installed in guest OS, a virtual machine monitor can notice that the daemon is no longer working when periodically trying to communicate with it. Although virtualization technology is widely deployed in PC-based systems, it is difficult to implement it in embedded devices as less hardware devices support it. Moreover, since the volume of source code for virtual machine monitor (VMM) tends to become large, the risk where VMM includes bugs also becomes large. To overcome the restriction, Kanda developed SPUMONE, which a light weight virtual machine monitor designed for working on embedded processors [15]. It provides a function to reboot the guest OS. However, SPUMONE does not provide memory protection mechanism between virtual machine monitor and the guest OS (Normal OS). Thus, it is vulnerable to an attack on the virtual machine monitor from guest OS.

VIII. CONCLUSION AND FUTURE WORK

In this paper, LiSTEE™ Recovery works effectively to resist critical bugs or attacks including zero-day causing system crash in order to keep availability of end-point devices. The performance evaluation is presented to show that the degradation of the existing system is small enough. Considering the deployment in the market, we show that the development cost and production cost can be minimized. Moreover, it can save maintenance cost.

Future work includes the resistance to sophisticated attacks. One possible attack is that an attacker illegitimately modifies the shared memory area to fake as if Normal OS works correctly while almost all Normal OS functions actually stop. As the result, LiSTEE™ TA misunderstands that Normal OS works correctly. One approach to solve this attack is to implement LiSTEE™ TA so that it itself checks the status of Normal OS without the support of an application program running on Normal OS. For example, whenever Normal OS is running, it must update a certain data area, such as page tables or process tables. Therefore, LiSTEE™ TA can determine Normal OS is working or crashed to monitor the data area. An advantage of LiSTEE™ is that it is impossible for Normal OS to reverse-engineer and to tamper an algorithm of LiSTEE™ TA because of memory protection mechanism. Thus, an attacker cannot know how to compromise Normal OS producing misleading information. We have not implemented this though. Another possible attack is damaging file system locating Normal OS. Network boot can be a solution where LiSTEE™ TA downloads a small rescue program from the head-end system when it fails booting.

REFERENCES

[1] C4 Security. “The Dark Side of the Smart Grid”. [Online]. Available: http://www.c4-security.com/The%20Dark%20Side%20of%20the%20Smart%20Grid%20-%20Smart%20Meters%20%28in%29Security.pdf [Accessed 19 Feb 2014]

[2] D. Forsberg, Y. Ohba, B. Patil, H. Tschofenig, and A. Yegin, “Protocol for carrying authentication for network access”, IETF RFC 5191 [Online]. Available: http://tools.ietf.org/html/rfc5191 [Accessed 19 Feb 2014]

[3] F. Zhao, Y. Hanatani, Y. Komano, B. Smyth, S. Ito, and T. Kamibayashi, “Secure authenticated key exchange with revocation for smart grid”, The third IEEE PES Conference on Innovative Smart Grid Technologies (ISGT 2012), IEEE Power & Energy Society (PES), Jan 2012, pp.1-8.

[4] H. Khurana, M. Hadley, N. Lu, and D. A. Frincke, "Smart-grid security issues", IEEE Security & Privacy, vol.8, Issue.1, Jan-Feb 2010, pp.81-85.

[5] K. Li, "Towards Security Vulnerability Detection by Source Code Model Checking", Software Testing, Verification, and Validation Workshops (ICSTW), 2010 Third International Conference on, Apr 2010, pp.381-387.

[6] M. Rinard, C. Cadar, D. Dumitran, D. M. Roy, and T. Leu, "A Dynamic Technique for Eliminating Buffer Overflow Vulnerabilities (and Other Memory Errors)", Computer Security Applications Conference, 2004. 20th Annual, Dec 2004, pp.82-90.

[7] S. M. Varghese and K. P. Jacob, "Anomaly Detection Using System Call Sequence Sets", Journal of Software, Vol.2 No.6, Dec 2007, pp.14-21.

[8] ARM. “ARM Security Technology” [Online]. Available http://infocenter.arm.com/help/topic/com.arm.doc.prd29-genc-009492c/PRD29-GENC-009492C_trustzone_security_whitepaper.pdf [Accessed 19 Feb 2014]

[9] T. Alves and D. Felton, “TrustZone: Integrated Hardware and Software Security”, Information Quarterly, Vol.3, No.4, 2004, pp.18-24.

[10] MITRE. “Common vulnerabilities and exposures” [Online]. Available http://cve.mitre.org [Accessed 19 Feb 2014]

[11] ARM. “Dhrystone Benchmarking for ARM Cortex Processors” [Online]. Available http://infocenter.arm.com/help/topic/com.arm.doc.dai0273a/DAI0273A_dhrystone_benchmarking.pdf [Accessed 19 Feb 2014]

[12] A. Kleen, “Machine check handling on linux”, Technical report, SUSE Labs, Aug 2004 [Online]. Available http://halobates.de/mce.pdf [Accessed 19 Feb 2014]

[13] T. Garnkel, B. Pfa, J. Chow, M. Rosenblum, and D. Boneh, “Terra: A virtual machine-based platform for trusted computing”, In Proceedings of the Symposium on Operating System Principles, Oct 2003, pp.193-206.

[14] "libvirt - the virtualization API." [Online]. Available: http://libvirt.org [Accessed 19 Feb 2014]

[15] W. Kanda, Y. Yumura, Y. Kinebuchi, K. Makijima, and T. Nakajima, "SPUMONE: Lightweight CPU Virtualization Layer for Embedded Systems", In Proc. of Embedded and Ubiquitous Computing, Dec 2008, pp.144-151.

Concerning the Sustainability of Smart Grids

A critical analysis of the sustainability of current Smart Grid models and on indicators of Smart Grid sustainability assessment

Antonio Caló, Jean-Nicolas Louis, Eva Pongrácz
NorTech Oulu, Thule Institute
University of Oulu
Oulu, Finland
e-mail: antonio.calo@oulu.fi

***Abstract* — There is currently no generally accepted definition of Smart Grids, especially on what is expected to make them smart and how. On the other hand, the general assumption about Smart Grids implies the development of power networks toward more efficient, greener and sustainable systems. In this article, we present the intrinsic limitations of this view, showing how any aim for a future sustainable energy system will require analyses considering the multidimensional structure of the problem, including the integrated dynamics of different energy sources and carriers. The aim of this paper is to present the case for the development and conceptualization of a smarter and more comprehensive system known as Smart Energy Networks. We also put forward a set of indicators that we believe could aid in assessing the sustainability of smart grids and follow progress toward the development of a more sustainable energy regime.**

Keywords - Smart Energy Networks; Smart Grids; Sustainable Development; Sustainability Assessment; Energy Services.

I. INTRODUCTION

There is currently no accepted definition of Smart Grids (SGs). An argument can be made that, in many respects, this also reflects a lack of generally accepted understanding of what SGs should be, do or achieve. Any discussion on SGs, though, includes two commonly accepted premises: they represent the (on going) evolution of the power network and their development is driven by our – ever more pressing – needs for a more efficient, environmentally performing and sustainable energy system.

In recent years, the quest for energy efficiency and environmental performance has led to extending the boundaries of SGs and a broader concept of energy network has emerged [1],[2]. The expression Smart Energy Network (SEN) indicates a concept intended to go beyond the SG idea, where multiple energy carrier and their synergies can be considered. SENs are therefore defined by broader boundaries and they can be expected to provide a better systemic tool for the development of a more sustainable energy regime.

At the same time, analyses concerning the sustainability of SEN focus primarily on the environmental performance, and even more limited to center on the reduction of greenhouse gases (GHGs) emissions. The general expectation of lower GHG emissions is based on an increased use of Renewable Energy Sources (RES) and improved efficiency through owing to, and supported by, the technological development of the last few decades especially in the fields of Information and Communication Technology (ICT). On the other hand, to the best of our knowledge, no in-depth consideration concerning the overall system performance of SEN has been introduced or discussed, leading to, in our opinion, a number of significant misapprehensions concerning the expected results in terms of sustainability.

In this paper, we intend to address these issues, presenting how the concept of SG alone is intrinsically insufficient and a broader concept, such as the SEN is required to properly address the sustainability of the energy system. Furthermore, we intend to discuss the idea of sustainability within the framework of SGs and SENs. Far beyond the mere notion of environmental performance, the concept of sustainability in itself is a subject of intense research work; in this paper we intend to provide the readers with a set of elements to be considered in the quest for a more sustainable (and smarter) energy regime.

This manuscript, beside the introduction, develops on 3 sections. Section II introduces the concept of sustainability applied to the generally understood development of SGs, underlying, from different point of view, the limiting factors the concept of SG carries along in comparison with the more comprehensive concept of SENs. In Section III we introduce the paradigm shift required to overcome the presented limitations, suggesting what we believe to be the two key factors: addressing the multilayered and multidimensional character of sustainable energy system, and overcoming of the intrinsically limited nature of a quantitative understanding of the energy system. In Section IV we present a summary of our critical analysis of the generally accepted sustainable development of smart grids and our final conclusions.

II. ADDRESSING SUSTAINABILITY

Addressing the concept of sustainability, in our work we refer to the widely accepted definition of a process that "… meets the needs of the present without compromising the ability of future generations to meet their needs" [3]. Although often associated with environmental performance, and consequent separated references to economic or

 ISBN: 978-1-61208-332-2

technological challenges and social acceptance/involvement, the concept of sustainability is based on what is referred to as the triple bottom line (TBL) (Figure 1). The TBL view indicates that companies ought to manage also their social and environmental capital in as much as they would manage their economic bottom line.

In this section, we discuss the currently generally accepted view of SG development from each of the mentioned sustainability dimensions.

A. *The Environmental Case*

The advantages and expectations concerning the environmental performance of SGs can, in most cases, be summarized in terms of significant reduction in GHG emissions through a wider use of RESs and improvements in terms of energy efficiency. We believe that, placed in these terms, both the goal and the means represent indeed fundamentally important factors for the development of a better energy regime, and yet they are an incomplete answer to the quest for a sustainable energy regime.

The reduction of GHGs is essential for combating global warming. On the other hand, it should be kept in mind that it is not the only element of concern in terms of environmental performance. Factors such as the competing uses of resources such as biomass and water, direct and indirect land-use impacts usually falls outside the spectrum of elements describable by using exclusively GHGs reduction based criteria.

Figure 1. The three pillars of sustainability [4].

Concerning the means through which we try to achieve our goals, it should be evident that the use of an increasing share of RES does not automatically imply an improved system performance toward sustainability. When it comes to evaluating the environmental performance on the use of RES, a good starting point is the fundamental limiting factors usually introduced in distinguishing renewable and non-renewable energy sources. Non-renewable energy sources are limited by the available stock: the progressive exhaustion of the stocks define the energy source availability usually in terms of financial viability and profitability. Renewable energy sources are defined by the flow: our limitations in tipping in these sources come, consequently, to be described in terms of limited capacity in “catching” these flows. If, in some cases, limitations are physical or technological, a limitation often overseen is defined by the limited resources we have at our disposal to bring to bear the means for “catching” the energy flow in the first place. Issues related to energy efficiency fall also under these limitations. In addition, the production and deployment of smart technologies massively rely on ICTs using a number of materials whose availability is given for granted. Scarcity of key component materials is a potential barrier to both large-scale deployment and reductions in technology cost. This is because many SG technologies (ICT and renewable energy) resort to critical metals. A raw material is labeled ‘critical’ when the risks of supply shortage and their impacts on the economy are higher compared with most of the other raw materials [5]. Critical materials important for low carbon energy generation and electronics share the same insufficiencies in their life cycles; low recycling rates and a high degree of dissipative losses, with a major share of the material being lost into other material flows [6]. Due to the rapid growth of SG technologies, end-of-life (EOL) management will become increasingly important. The role of recycling is negligible at the moment, but shall play an important role when recent and future installations will reach their EOL.

An important conclusion concerning the analysis of environmental implications for the development of a sustainable energy regime is that, also in the case of RES, we need to deal with stock limited resources. Since the development of a SG based system is expected to rely on these technologies, it is important to take these elements into considerations.

B. *The Economic Case*

In this paper we do not intend to produce an in depth discussion concerning the economy and economic mechanisms behind the development and deployment of SGs. It is, on the other hand, our intention to present few key points that we believe are generally missing and which are essential for sustainability.

Concerning the economic aspects of sustainability, this element is often understood in terms of financial profitability. Although an essential aspect of any system development, financial considerations should not entirely replace the economic ones (economic being here understood as the science that study the production, distribution and consumption of goods and services). Already the concept of sustainable use of non-renewable resources comes with some important economic consequences (e.g., access and use rate of mined resources) and, as mentioned in the previous section, it represents the determining factor for stock defined resources. On the other hand, even accepting the intrinsic limitations of a predominantly finance based approach to the problem, major concerns regarding possible contradictions in the development of a sustainable energy system arise. Simply put, the key issue resides in the maintenance of the core business model: as long as the customers pay for energy

 ISBN: 978-1-61208-332-2

units (kilowatt-hours), the more are the kilowatt-hours sold, the better are the revenues supporting the growing energy business. The argument then can be made that there is little actual incentive to systematically reduce consumption or, in contradiction with a popular and generally supported understanding, energy efficiency does not – cannot – automatically translate in cheaper bills for the end-users at the end of the value chain.

C. *The Social Case*

The social dimension of SG development is often related, on one hand, with user response and interaction with a smarter power network, and, on the other hand, with concerns related to privacy and security with respect to a technology potentially very intrusive. We can argue that this is also a valid, important, and yet incomplete assessment of the social dimension related to SGs. A large use of RESs implies the distributed deployment of power generating units, access to resources often not necessarily accepted by the involved parties. Furthermore, privacy and security issues are often addressed with "end-of-pipe" solutions such as data protection, encryption, etc., while the efficient and effective system operation often relies on virtually unrestricted possibilities of access and control. A number of solutions concerning these issues have been presented, and they generally rely on the concentrated accumulation of data. Setting aside, for a moment, pure technological considerations and focusing on more technical aspects of this approach, a question should be asked on the direction we pursue for the design of the SGs system. The new system will be born from the synergic integration of the energy and the ICT networks, relying on, and supported by, an ever more decentralized power production and communication. Under the new paradigm, it should be questioned if physically decentralized power production system still relying on a centralized data management system can be still technically referred to as Smart Grid.

D. *Sustainability Indicators*

It is expected that long-term profitability of energy provision should go hand-in-hand with social justice and protecting the environment [7]. Indicators are needed to assess the environmental, economic and social sustainability of smart grids, to monitor trends in conditions over time, or to provide an early warning signal of change. It is widely recognized that some socio-economic indicators are related to environmental indicators (e.g., resource conservation) and that public acceptance depends on environmental impacts [8]. This manuscript recommends using a set of indicators under four categories, based on Dale *et al.* [9] and Global Reporting Initiative [7] to assess the sustainability of smart grids:

- Energy security and profitability
  - Security of supply
  - Net present value (NPV)
  - Energy price volatility
  - Trade volume
- Social acceptability
  - Energy quality
  - Data privacy and security
  - Transparency and compliance
  - Public opinion, stakeholder participation
- Social well-being
  - Employment
  - Impact on customer well-being
  - Access to energy
  - Diverse services
- Resource conservation
  - Depletion of non-renewable resources
  - Energy return on investment (EROI)
  - Land-use changes
  - (CO_2) emissions

This selection of indicators of SG sustainability is based on the availability of information about socioeconomic conditions for each category, on other efforts to identify sets of indicators, and on established criteria for selecting indicators. This set of 16 indicators is not as detailed or comprehensive as other approaches but, we argue, that it is more practical to apply. The indicators were selected based on being practical, unambiguous, resistant to bias, sensitive to changes, related to those changes, predictive, estimable with known variability, and sufficient when considered collectively [9].

III. PARADIGM SHIFT

The expression "paradigm shift", a ubiquitous phrase often used to describe the radical changes SGs are expected to bring about on century old power system, finds its origins in the book *The Structure of Scientific Revolutions* by Thomas S. Kuhn [10]. In his book, Kuhn, a historian and philosopher of science, argued that science does not progress linearly but though a series of anomalies that challenge the existing theories and, after some time, generate a revolutionary change in thoughts, a paradigm shift. An often cited example of this process is the Copernican revolution vs. the Ptolemy's view of the universe. The comparison between the Copernican and the SG revolutions is perhaps useful to underline the deep and radical system level re-thinking of the energy system required for a sustainability assessment. It is indeed the case that innovative solutions might appear more like an addition of system appendices, the equivalent of *epicycles* and *deferents* in the Ptolemaic system, rather than a radical, concept level, change of perspective.

Essentially, there are two key factors that need to be addressed: the intrinsic limitations of addressing the power system alone and the fundamentally quantitative understanding of the energy system.

A. *A more comprehensive approach*

As mentioned earlier, the expression Smart Energy Network refers to an information-based distributed system including a multilayered and multidimensional interaction

 ISBN: 978-1-61208-332-2

of different energy carriers and vectors. The synergies among different forms of energy (end their network flows) allow for a more comprehensive theoretical framework upon which we could base our analyses.

A sustainability analysis focusing exclusively on SGs does not allow a proper and complete assessment of the entire energy system value chain. In particular, essential aspects of the energy production are amiss, as they do not directly contribute to the power system. Consequently, a *conditio sine qua non* is the consideration of a larger, more comprehensive system that would include the synergic flow of resources and energy. At the same time, broadening the system boundaries generally reduces the risk of externalities introduced in the previous section and calls for a general re-evaluation of system sustainability in the face of the new paradigm.

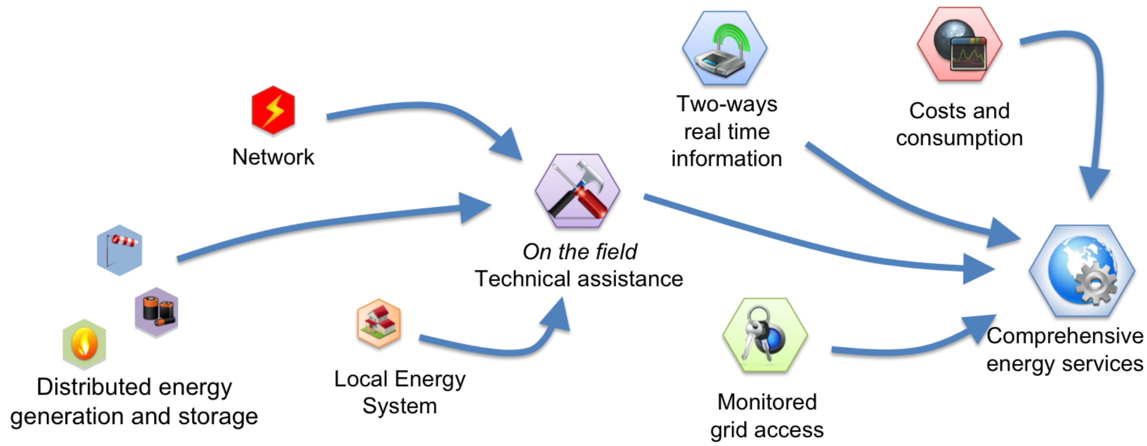

Figure 2. Indicative representation of the different elements required for an effective and comprehensive energy service to play a role in a smart grid based energy system.

B. *Energy Services*

In order to sustain a radical transformation of the future energy network, it is required for an important share of the future energy business to move from the production and commercialization of energy to the offering of comprehensive energy services (Figure 2). Customers' services would include, among others, monitored grid access, real-time information on costs and consumption, possibility to customize the service's profile and the possibility to sell small-scale renewable energy production in an open, transparent and easily accessible manner. These services rely primarily, but not exclusively, on efficient and effective data sharing and information extraction processes used by different operators to maintain a real-time overview of the network and feed back necessary, and valuable, information to end-users (e.g., energy consumption and price forecast). The other element energy services are to rely on is customized "on the field" technical assistance. This can potentially support the development of a more sustainable energy regime on two fronts. On one hand, end-users' living environments are expected to develop and grow in efficiency and complexity, requiring a number of dedicated and customized technical solutions and support services often currently not available on the market. This would potentially support the development of the system sustainability, especially for what concern the social dimension we described in section 2C. On the other hand, these services can support the development of an energy quality based systems rather than a material production ones. The key idea is to prevent moving the focus of the sustainability problem from the production of energy to the production of the materials required to produce energy, therefore avoiding the shortcomings we discussed in section 2B.

It is, notwithstanding, important for the new energy system to remove the commercialization of energy on a quantitative base as the primary economic driver. Based on our previous consideration, we can claim that any scenario built outside this premise is destined to present inevitable limitations, and any claim of sustainability can be put forward only by limiting the analyses to some specific element of the energy system (i.e., Smart Grids). Consequently any consideration of (un-)sustainability becomes, de facto, an externality.

Energy services are currently not widely available, or shared among a limited number of market players in a non-user-friendly way. Furthermore, each element of the network needs monitoring and assistance, they need to be combined at the local level and they need to be properly included in the common grids where their activity can be properly monitored in order to maintain the efficiency and the stability of the common networks. All these elements, even if, in most of the cases, are technically available, are not provided in a coordinated and viable fashion to the end-users. Other elements, especially for what concern the inclusion of the models considered in this work in the common grids, are currently not available. Apart from a number of administrative and political aspects, there are a number of technical obstacles that need to be overcome.

IV. CONCLUSIONS

In this paper, we presented a critical analysis of the sustainability of smart grid development. We pointed out that, in terms of assessing the sustainability of smart grids, attention is often limited to environmental considerations only (such as CO2 emissions and the share of renewables), often ignoring entirely the economic and social dimensions of sustainability. More specifically, we showed that the concept of sustainability, in its sometimes underestimated complexity, comes, for what concern the future development of SGs, with two key elements of consideration: the first one is that the concept of SG is, in itself intrinsically inadequate for a proper sustainability evaluation as it represent a far too limited portion and, therefore, an incomplete representation of the energy system. The second point is that the economic driver of the future energy regime will have to oversee a transition from an energy based system to an energy service based system. Finally, we foresee that the social well-being impact of smart grids will be a key characteristic in European smart grid development efforts, with social acceptability providing a "make or break" feature, especially in terms of stakeholder participation, transparency and compliance.

ACKNOWLEDGMENT

A. Caló and E. Pongrácz thank the funding provided for the SMARCTIC project by the Finnish Funding Agency for Technology and Innovation (Tekes). J.-N. Louis thanks the funding of the Thule Institute Research Programme.

REFERENCES

[1] D. S. Chai, "Simulation of congeneration within the concept of smart energy networks", M.A.SC. thesis, University of Waterloo, Mechanical Engineering, 2012.

[2] D. S. Chai, J. Z. Wen and J. Nathwani, "Simulation of congeneration within the concept of smart energy networks", Energy. Convers. Manage., vol. 75, 2013, pp. 453–465.

[3] World Commission on Environment and development, "Report of the World Commission on Environment and Development: Our Common Future", United Nations, 1987.

[4] W. M. Adams, The Future of Sustainability: Re-thinking Environment and Development in the Twenty-first Century, Report of the IUCN Renowned Thinkers Meeting, University of Cambridge, 2006. URL: http://cmsdata.iucn.org/ downloads/iucn_future_of_sustanability.pdf.

[5] European Commission, DG Enterprise and Industry, "Critical raw materials for the EU: Report of the Ad-hoc Working Group on defining critical raw materials", study of the European Commission, Enterprise and Industry's study, 2013.

[6] T. Zimmermann and S. Gößling-Reisemann, "Critical materials and dissipative losses: A screening study", Science Total Environ., 461–462, 2013, pp. 774–780.

[7] Global Reporting Initiative G4, "Sustainability reporting guidelines", 2013. https://www.globalreporting.org/ resourcelibrary/GRIG4-Part1-Reporting-Principles-and-Standard-Disclosures.pdf.

[8] S. L. Collins et al. An integrated conceptual framework for long-term social–ecological research, Front. Ecol. Environ., 9 (2011), pp. 351–357:

[9] V. H. Dale et al. Ecological Indicators 26 (2013) 87–102, http://dx.doi.org/10.1016/j.ecolind.2012.10.014.

[10] T. S. Kuhn, "The structure of Scientific Revolutions", University of Chicago press, 1962.

Energy Management in a Smart Grid-integrated Hydrogen-based Storage

Electric Grid balancing when integrating large-capacity Renewable Energy sources

Diego Arnone, Alessandro Rossi, Massimo Bertoncini
Engineering I. I. S.p.A.
Palermo and Roma, Italy
diego.arnone@eng.it, alessandro.rossi@eng.it,
massimo.bertoncini@eng.it

Diana Moneta
Ricerca sul Sistema Energetico S.p.A.
Milano, Italy
diana.moneta@rse-web.it

Rosario Proietto
Onyx Technology S.r.l.
Roma, Italy
proietto.rosario@gmail.com

Giuseppe Tondi
ENEL Distribuzione S.p.A.
Roma, Italy
giuseppe.tondi@enel.com

Carlos García-Santiago
Tecnalia
Miñano, Spain
carlosalberto.garcia@tecnalia.com

***Abstract*—The integration of plants for the distributed production of electric energy from renewable (especially solar and wind) sources into the electric grid has become an increasingly complex task over the last years because the amount of “green” energy nowadays injected into the grid is sometimes comparable to the energy supplied by traditional thermal plants. This introduces a strong variability, a partial unpredictability and a significant dependency from the location, three of the main features that characterize the energy produced by photovoltaic plants and wind turbines. This paper introduces a high level description of an innovative system that integrates high-capacity hydrogen-based storage into the grid with the aim of contributing to the grid balancing and to the improvement of the power quality. A concrete demonstrator is being designed and will be connected to a properly selected primary substation where most of the power lines connected to the medium voltage bus bar are active because they are connected to many photovoltaic plants and wind turbines.**

Keywords-energy management system; hydrogen-based storage; integration of renewable energy sources; grid balancing.

I. INTRODUCTION

Thanks to a growing environmental awareness as well as to many government directives and incentives, the distributed production of energy from renewable sources is rapidly expanding. The amount of “green” energy nowadays injected into the grid is in some places comparable to the energy supplied by traditional thermal plants. In order to obtain an efficient Renewable Energy Sources (RES) injection into the grid in terms of stability and power quality, both transmission and distribution grids need a certain level of “smartness”, because RE is characterized by non-controllable variability, partial unpredictability and locational dependency [1]. So, the integration of photovoltaic plants and wind farms into the grid entails the need to provide proper solutions to a wide range of problems not entirely new to the network but still more critical. The huge amount of RE can negatively affect the efficiency of classic fossil fuel-based generators and reduce their lifetime because it forces them to operate in not optimal and continuously changing conditions to balance load and overall supply.

Today, Distribution System Operators (DSOs) must compensate RE output fluctuation, grid faults, conventional generation outages, load variation with more flexibility consisting in conventional generation flexibility (e.g., “unit commitment” techniques), demand/response, energy storage (e.g., batteries, hydrogen production and storage, water pumping, etc.) and grid-friendly RE generation. In this context, the INGRID FP7 European co-funded project [2][3][4] is studying several solutions, which this work, still in progress, is part of. In order to balance power supply and demand by empowering the grid flexibility, the proposed solution may either shift the electricity adsorption or modify the energy injected into the grid. These methods must be applied in the context of a strict cooperation with the distribution system operator. Our integrated approach involves combining a solid-state high-density hydrogen storage system with advanced solutions for smart distribution grids, which monitor and control a large number of RES. In the vision of the INGRID project, when a huge amount of energy produced by RES injected into the distribution grid reaches a critical level in terms of balancing, power quality, grid stability and so on, this amount of green energy (or part of it) may be transformed and moved on a different energy vector, e.g., hydrogen, syngas, Full Electric Vehicle (FEV)

 ISBN: 978-1-61208-332-2

mobility, or simply injected into the grid again, by providing proper ancillary services. The brain of the proposed solution is the Energy Management System (EMS).

The remainder of the paper is organized as follows: in Section II, an overview of the related works is provided; in Section III, the INGRID System is described; in Section IV, more details are provided about the role played by the EMS; in Section V, the EMS optimisation problem is modelled as well as the solution that is being investigated to address it; in Section VI, the conclusions are given.

II. LITERATURE REVIEW

Many research projects and initiatives are focused on proposing new approaches and solutions to balance energy supply and demand when renewable sources are integrated into the electric grid. Recently closed EU co-funded FP7 projects, like SEESGEN-ICT [5], INTEGRIS [6] or HiPerDNO [7], aim at designing and implementing new solutions for EMS, in particular through communication networks improvements. Other research works, like the recently ended MIRABEL project [8], propose an approach on a conceptual and infrastructural level that allows energy distribution companies to balance the available supply of renewable energy sources and the current demand in an ad-hoc fashion. Moreover, many research works concerning optimization strategies for EMS have been published. As reported in [9][10][11] the optimization strategy is based on a proper objective function that may be a profit function or a losses function. However, despite a considerable amount of research undertaken in ICT-based EMS, significant research challenges do exist in terms of improving sustainability, reliability and cost-efficiency of energy supply. Of course, in order to cope with these challenges, significant changes towards an ICT-empowered Smart Energy Grid are required. Moreover, the expected transition to Smart Grids requires deep technological transformations and huge financial investments. This involves enhancing and upgrading existing infrastructure, implementing new systems and improving integration throughout the ICT operating environment.

In general, very partial and fragmented solutions have been proposed so far. A comprehensive yet fully integrated contextual data model is missing, as well as a suitable intelligence in the processing of the captured information; all of these aspects will be fundamental for proactively predicting the energy production and accordingly fine tuning the energy produced and transmitted into the grid.

Current EMSs are lagging behind due to their low integration level among the different subsystems, like power grid, RE sources, storage system, etc., which prevent EMSs to effectively accomplish their task of mitigating the intermittent energy supply from RE sources through balancing activities offered by storage systems.

III. THE INGRID SYSTEM

The INGRID system will strictly collaborate with the Distribution Management System (DMS). That collaboration will allow either to adsorb or to inject active or reactive energy considering both DSO indications and results coming from each subsystem.

When needed, the INGRID system may indirectly adsorb the electricity produced by RES based plant outside INGRID system by means of an electrical connection to the grid. Moreover it may directly adsorb/supply electricity if an internal RES (iRES) based plant inside INGRID system is taken into account.

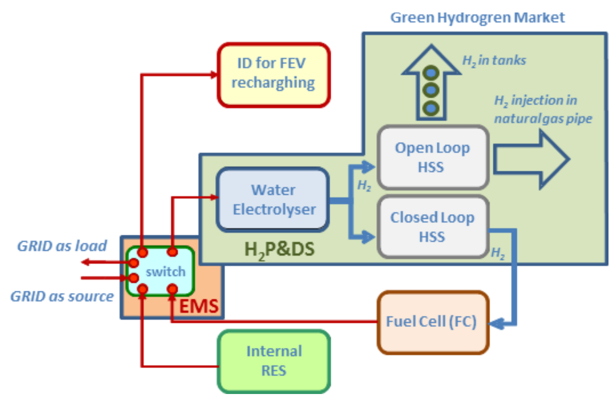

Figure 1. Graphical representation of the INGRID System.

As depicted in Figure 1, the adsorbed electricity is used to supply a Water Electrolyser (WE) to produce hydrogen which may be either stocked in an innovative solid-state storage system by exploiting a patented magnesium hydrides-based technology, or directly injected into existent methane pipeline (if feasible). The hydrogen, stored in high capacity but conventionally sized tanks, can be sold in the green H_2 market (open loop) or can be used to supply a Fuel Cell (FC) that, according to the specific strategy adopted by the System, converts it again in electricity (closed loop). This electricity is either injected into the grid or used by the charging stations of an innovative green urban mobility system [12] whose management system is named Intelligent Dispenser (ID). Grouping all the subsystems in two main categories, i.e., loads or sources, the INGRID EMS, addressed by this paper at a very preliminary stage, is the ICT-based component in charge of monitoring and controlling the power flows from sources to loads.

TABLE I. POSSIBLE ENERGY CONNECTIONS

		Loads			
		Grid	AUX	ID	H_2P&D
Sources	Grid	no	yes	Yes	Yes
	iRES	yes	yes	yes	Yes
	FC	yes	Yes	Yes	No

In Table 1, a schema of the possible electric connections inside the plant is reported, where AUXiliary equipment (AUX) is considered as well.

In order to perform a real-time demand/supply balance according to a specific strategy (e.g., economic sustainability), the EMS must exchange information about energy requests and offers, costs and prices, constraints and degrees of freedom, with all the System sub-units that will be equipped with proper management systems. For instance, the Hydrogen Production and Distribution System (H_2P&DS)

will control hydrogen production, manage its storage, accept purchase orders, supervise and manage the tanks delivery and its injection into the gas distribution network. But, it also will predict its own power adsorption in the future hours and communicate this adsorption profile to the EMS that, in its turn, will decide if and to what extent the required power adsorption profile will be followed. In general, the EMS receives (sends) from (to) each subsystem a load or generation trend intended as a graph of power vs time. The time (horizontal axis) is divided in different "time slots" where the EMS, in accordance with a certain optimisation strategy, defines the power level (vertical axis) of each subsystem, i.e., the set-point of each component, its electrical connections and widely the plant configuration.

The System that is being currently designed will be instantiated in a concrete 39 MWh energy storage facility that will be deployed and will operate in Troia (Puglia region, Italy). As a variable passive load (0 - 1200 kW), the plant will be connected to a middle voltage feeder of a primary substation (150/20 kV/kV). As a variable generator (0-90 kW), the plant will be connected to a low voltage power line.

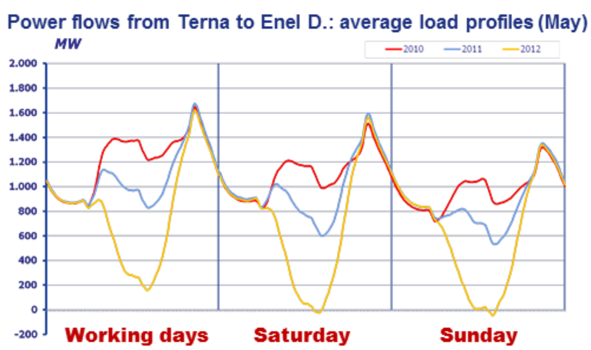

Figure 2. Exchange profile from Terna to Enel Distribuzione in Puglia region.

The primary substation has been identified among the ones having an high value of power reverse flow, calculated as the number of hours in one year in which the phenomena occurs. The curves in the Figure 2 represent the power requested by the distribution network managed by Enel Distribuzione in the Puglia region, to the transmission and transport grid, managed by Terna, the unique Transport System Operator (TSO) in Italy. The diagram shows the effect of the growing distributed generation from 2010 to 2012. The curves are lowered along years reaching a value of zero (in 2012, the local RES generation balanced the local adsorption) and even a reverse flow during the low load days occurred.

IV. THE ROLE PLAYED BY THE EMS

The EMS is the core of the INGRID system defining the energy adsorption/supply of the INGRID components. Energy Management Systems can be found in a number of different applications, since this generic denomination can be adopted whenever some kind of management must be applied to energy (e.g., mobile devices, data centres, FEVs, smart buildings, etc.).

In the Smart Grid domain, the concept of EMS-Family has been recently proposed [13]; it includes different types of EMS for generation plants, for TSOs, for DSOs, for substations, for micro-grids and final energy consumers. All these systems are coordinated by a global EMS that can be considered as the Smart Grid brain. The system described in this paper is in charge of dispatching electric power flows among the different sub modules of the overall infrastructure by significantly contributing to grid balancing, while ensuring the economic sustainability by selling hydrogen and providing FEV charging services.

The proposed EMS will be based on the concept of a closed-loop feedback control process. The feedback cycle starts with the collection and monitoring of relevant data from the interested sources that reflect a fine-grained state of the whole system. A comprehensive contextual information model will fully describe the whole status of the involved parties, i.e., the Green Energy Storage (GES), the grid or the Green Urban Mobility System, as well as factors which affect the decisions for energy dispatching, i.e., price of both energy and hydrogen, historical energy demand, etc. Structuring and reasoning about the collected data will be done during the subsequent analysis phase which will include machine learning methods and semantic contextual information extraction, business intelligence and statistical methods, consumption forecasting and distribution rules.

The analysis phase will originally combine and integrate intelligent processing technologies, e.g., basic and advanced statistics, data mining, complex event processing (CEP), predictive analytics, etc., to plan all the power flows internal to the plants and towards the grid as well as the hydrogen tanks transfers into the green hydrogen market.

Applying CEP techniques in the domain of the Smart Grid will allow to detect relevant events from the distributed and heterogeneous data sources and to analyse their impact in real time. Based on the results coming from the analysis, the EMS will take a decision about how to reach a desirable state, often taking into account opposing goals, involving also external factors, such as local policy for a risk analysis about reliability of the whole energy system.

Finally, in order to accomplish all the decision task balancing the energy flows, the EMS must monitor each of the overall system sub-modules, i.e., the WE, the Hydrogen Solid-state Storage (HSS), the FC, the ID and its scheduling system, the local RES production and related predictions as well as all the request coming from the DSO (e.g., through the DMS), the price of energy, the price of hydrogen, etc.,

The full feedback process can follow two approaches where, according to the requirements of reliability and/or to risk policies, a human intervention can be required or not, enabling the EMS system to follow either fully automatic or semi-automatic paradigm by requesting the human intervention in order to achieve the desired goal. So, human operators will be able to question, and possibly, to override the decisions. Similarly, the knowledge will be incremental, learning from past successful and unsuccessful decisions taken by the system, as well as from issued human advices.

V. The EMS optimisation problem

The optimisation problem is related to the maximisation of an objective function. In this case revenues and costs have to be taken into account and a profit function is considered, as fully described in [14], and just reported here.

$$\sum_{i=1}^{N}\{P_{WE-HSS-OL}(t_i)\cdot\Delta t\cdot K\cdot\eta\cdot p_{H_2}\cdot w_1+P_{IDdemand}(t_i)\cdot\Delta t\cdot p_{EV}(t_i)+$$
$$+\beta\cdot P_{FC}(t_i)\cdot\eta\cdot\Delta t\cdot p_{EV}(t_i)\cdot w_2+(1-\beta)P_{FC}(t_i)\cdot\eta\cdot\Delta t\cdot p_{ANC}(t_i)\cdot w_2$$
$$+[P_{RES}(t_i)\cdot\Delta t\cdot c_{grid}(t_i)-P_{IDgrid}(t_i)\cdot\Delta t\cdot c_{grid}(t_i)-(P_{WE-HSS}(t_i)\cdot\Delta t\cdot$$
$$\cdot c_{grid}(t_i)]\cdot[1-\alpha\cdot d_{\%}^{-}] \qquad (1)$$

The WE power level at a generic instant is represented by P_{WE-HSS} where $P_{WE-HSS-OL}$ is the part dedicated to the open loop. $P_{IDgrid}\Delta t$ is the energy adsorbed by the grid to supply the ID and is the part of the overall required energy ($P_{IDdemand}\Delta t$) sold at the price p_{EV} . βP_{FC} represents the FC power dedicated to the ID, while $(1-\beta)P_{FC}$ is related to the ancillary services. p_{ANC} represents the related incentive, c_{grid} the electricity price, P_{RES} the RES power. K is a conversion parameter from electrical energy to H_2 quantity, η is the round-trip efficiency, w_i are technical weights and $d_{\%}^{-}$ is the incentive granted by the DSO for following its profiles.

The first part of the function represents the possible revenues from electricity supply and hydrogen production, whereas the second one is related to the incurred costs. Some weights are used to take into account technical constrains.

Moreover, that function has to be maximized considering different technical constraints as the maximum FC power, the advised range of the WE power level, the capability of the HSS system and so on. As an example, the constraint related to the storage capability is reported below:

$$0\leq HSS_{CL-state}\,(t_i-1)+P_{WE-HSS-CL}(t_i)\cdot\Delta t\cdot K_1\cdot\eta-P_{FC}(t_i)\cdot\Delta t\cdot$$
$$\cdot K_2/\eta\leq HSS_{CL-capacity} \qquad (2)$$

where $HSS_{CL-state}$ and $HSS_{CL-capacity}$ represent the closed loop HSS charge state and its capability, respectively.

Considering the complexity and nonlinear trends, three different algorithms are being implemented to solve the optimisation problem: two heuristic algorithms, Simulated Annealing and Tabu Search, and the nonlinear programming method Generalised Reduced Gradient. By means of these algorithms, the EMS will control a subset of the input variables defining an optimal plant configuration for each considered time slot (Δt) in a wide time horizon (N*Δt).

VI. Conclusion and Future Developments

The proposed solution is a variable and controllable load/generator, whose cooperation with the DSO will support the electric grid balancing, following the adsorption or generation profiles sent periodically. The development of new incentive models to support such a flexible system, the management of the hydrogen production and sale and energy supply to electric vehicles will help to accomplish this task.

The system potentiality in real conditions will be tested by a prototype connected to a feeder of a primary substation where a power reverse flow occurs, since the renewable energy produced exceeds the adsorbed one. The proposed solution will be as much general and versatile as possible, allowing the adoption of different optimization approaches.

The results achieved by applying the proposed algorithms will be processed and released soon and will help to evaluate the economic sustainability of the proposed solution.

Acknowledgment

This work is part of the INGRID project, co-funded by the European Commission (EC) within the FP7 Framework Programme. Authors thank the members of the INGRID Consortium as well as the EC for supporting any project dissemination activities. This work reflects only the authors' views. The Commission is not liable for any use that may be made of the information contained therein.

References

[1] "Grid integration of large-capacity Renewable Energy sources and use of large-capacity Electrical Energy Storage", White Paper, IEC MSB, Oct. 2012, http://www.iec.ch/ 2014.1.13

[2] INGRID (High-capacity hydrogen-based green-energy storage solutions for grid balancing) EU FP7 project: http://www.ingridproject.eu/ 2014.1.13

[3] D. Arnone et al., "An ICT-based Energy Management System to Integrate Renewable Energy and Storage for Grid Balancing", ACM e-Energy 2013, Berkeley, USA, May 2013.

[4] A. Rossi, D. Arnone, M. Bertoncini, and D. Moneta, "Sistema di gestione energetico innovativo con stoccaggio a idrogeno per l'integrazione di fonti rinnovabili a supporto del bilanciamento della rete", AEIT 2013, Palermo, Oct. 2013.

[5] SEESGEN-ICT (Supporting Energy Efficiency in Smart GENeration grids through ICT) EU FP7 project: http://seesgen-ict.rse-web.it/ 2014.1.13

[6] INTEGRIS (INTelligent Electrical Grid Sensor communications) EU FP7 project: http://fp7integris.eu/ 2014.1.13

[7] HiPerDNO (High Performance Computing Technologies for Smart Distribution Network Operation) EU FP7 project: http://dea.brunel.ac.uk/hiperdno/ 2013.1.13

[8] MIRABEL (Micro-Request-Based Aggregation, Forecasting and Scheduling of Energy Demand, Supply and Distribution) EU FP7 project: http://www.miracle-project.eu/ 2014.1.13

[9] M. Aguado et al., "Economical assessment of a wind hydrogen energy system using WindHyGen® software", International Journal of Hydrogen Energy, vol. 34, issue 7, Apr. 2009, pp 2845-2854.

[10] R. Dufo-López, J. L. Bernal-Agustín, and J. Contreras, "Optimization of control strategies for stand-alone renewable energy systems with hydrogen storage", Renewable Energy, vol. 32, issue 7, June 2007, pp. 1102-1126.

[11] R. Carapellucci and L. Giordano, "Modeling and optimization of an energy generation island based on renewable technologies and hydrogen storage systems", International Journal of Hydrogen Energy, vol. 37, issue 3, Feb. 2012, pp. 2081-2093.

[12] F. D'Errico, A. Screnci, and M. Romeo, "A Green Urban Mobility System Solution from the EU Ingrid Project", REWAS 2013: Enabling Materials Resource Sustainability, pp. 179-191.

[13] H. Sun, B. Zhang, W. Wu, and Q. Guo, "Family of energy management system for smart grid", ISGT Europe 2012, 3rd IEEE PES International Conference and Exhibition, 14-17 Oct. 2012, pp. 1-5.

[14] "D2.4 - EMS Requirements and Algorithms Analysis", FP7 INGRID Project's report, unpublished.

Smart Grid Architectures and the Multi-Agent System Paradigm

Cagri Yilmaz, Sahin Albayrak, and Marco Lützenberger
DAI-Labor
Technische Universität Berlin
Berlin, Germany
{cagri.yilmaz|sahin.albayrak|marco.luetzenberger}@dai-labor.de

***Abstract*—Currently, coal-, gas-, oil and nuclear power plants are used to meet the ever increasing energy demand. Nevertheless, due to an increasing environmental awareness and entailing regulations, energy prices rise considerably. As for today, many companies try to facilitate the production and utilization of much cheaper energy from renewable energy sources, such as wind and solar energy. However, due to constructional deficiencies, today grid infrastructures support only a small amount of renewable energy sources. To counter this problem, sophisticated control applications were designed, deciding on when and where to procure energy and safeguarding the stability of the entire grid architecture. Over the last years we have developed many smart grid applications and—due to the distributed nature of the problem solving task—we always applied multi-agent technology for this purpose. Java Intelligent Agent Component (JIAC), version five, innately provides many features indispensably required for smart grid applications. In this paper, we present our experiences in successfully using JIAC framework for the smart grid applications.**

Keywords–Power Network;Smart Grid;Distributed Power Generation; Electric Vehicles;Critical Infrastructures Protection;Demand and Response;Evolution Strategy; Multi-agent Systems;JIAC Agent Framework

I. Introduction

The world's energy demand increases continuously as the use of electronic devices grows. Increasing energy demand is accompanied by an ever increasing environmental awareness. 25.9% of total carbon emission is due to the consumption of high carbon fuels in current power industries [1]. In addition, today grid infrastructure is not able to cope with the increasing energy requirement. Aged components of current electric grid and out-dated technologies cause significantly inefficient and unstable electric systems. Beyond that, developing technology in last decades will change the methods and technics of generation and delivery of electricity. Integration of Information Communication Technologies (ICT) systems with grid infrastructure allows sophisticated control over generation, transportation, storage and consumption of electricity power. In current grid infrastructure, local distributed generators are accommodated on the electric grid in a fit-and-forget manner. Centralized electricity generation also creates an unreliable power load distribution system. Frequently, occurring blockages and outages inhibits sustainable power flow to the consumers.

Innovative techniques in smart grid will sustain utilization of electricity more securely and efficiently. Flexible and reorganisable virtual power plants ensure better management of power flow in case of any failures in electric grid. Assistance systems are developed in order to aid operators to handle with the problems. Managing charging and discharging processes of electrical vehicles (EVs) will be another significant issue. Uncontrolled charging profiles of EVs will increase the peak demands. Load management systems are developed to shift the EV demand to the other non-peak periods. Another concept, demand and response mechanism, makes the customers get involved into the power management. Considerable amount of savings can be achieved with demand management. Further, the demand and response concept allows customers to deliver surplus electricity to the power grid. In addition to that, policies for the reduction of CO_2 emission also become vital issues for future electric grid. According to the International Energy Agency (IEA) estimates, introduction of smart grid technologies will provide 0.9 to 2.2 gigatonnes reduction in CO_2 emissions by 2050 [1]. Over the last years, we have developed many smart grid applications and, due to the distributed nature of problems that were solved by with these applications, we always used multi-agent technology for this purpose. Nevertheless, agent-technology is still a domain which is mainly confined to the realm of research [2]. However, it is our experience that agent-technology should be seriously considered by industrial players. In order to substantiate this thesis, we present selected projects, in which we used agent technology, namely the JIAC agent framework [3]. It is our experience that most problems in smart grid architectures come from the distributed nature of the smart grid. The multi-agent system approach is a mature and efficient mechanism to master these problems.

The remainder of this paper is structured as follows: Current architecture of the electric grid and related challenges are described in section II. It is our purpose to emphasise the distributed nature of many grid related problems. In Section III, we outline smart grid architectures and potential extensions. In Section IV, we present selected applications that tackle smart grid related problems by using multi-agent technology. In Section V, we wrap up with a conclusion.

II. Electric Grid

An electric grid is an interconnected grid networks composed of several hardware and software components to distribute electricity from generation utilities to the end-users. Power delivery and communication infrastructures are two main layers of the electric grid.

 ISBN: 978-1-61208-332-2

A. Current Grid Architecture and Challenges

Maintaining continuous stability in the electric grid is a great challenge since most of the consumed electricity is provided from centralized power plants. Supervisory Control and Data Acquisition (SCADA) systems are used for data transmission among electronic devices of the current electric grid. SCADA allows a utility's operators to monitor and control the conditions of the grid equipment and network processes. Computers at the control centre and Remote Terminal Units (RTU) in the field measure the data sets from hundreds to tens of thousands of data points [4]. Control model is based on slow automatic control by the control centre in order to balance load and generation. Synchronous Phasor Measurement Units (PMU) gathers data several times in each power cycle to get detailed picture of the grid dynamics for systems planning, control and post-incident analysis. Special Protection Schemes (SPS), which is also called as Remedial Action Schemes (RAS) have been developed to meet some of the wide-area control needs. An SPS involves instituting hardwired, point-to-point communication among substations [5]. In current infrastructure, these data can not be reasoned beyond the substations. Customers demand are not involved into the load management strategies. The present load management system is mostly oriented to the electricity generation of the traditional power plants. This grid structure has a limited ability to cope with fast cascading phenomena such as blockages and outages.

III. Innovative Smart Grid Concepts

Traditional electric grid has been evolving into the smart grid composed of spatially complex, intelligent and autonomous power and telecommunication networks. Renewables energy sources, plug-in electric vehicles and distributed battery systems will increase the reliability and stability of the electric grid. Intelligent power management system should be designed to manage the power load among various decentralized energy sources and consumers. In this sense, better understanding of the innovative concepts is vital to analyse and control new power system.

A. Electric Vehicles

Electric vehicles are regarded as dispatchable energy storage units. They become an integrated part of the power network with the well-known applications which are Vehicle-To-Grid (V2G) and Grid-To-Vehicle (G2V) technologies [6]. Cooperation of electric vehicles with renewable energy sources make the power network more secure and reliable. Required data regarding power quality and availability of Distributed Energy Resources (DERs) and EVs are aggregated and processed by responsible control units. There are several approaches to model charging and discharging processes of electric vehicles. The study [7] describes a decentralized online EV charging scheduling scheme as a cyber-physical system. In another approach [8], a strategy to coordinate the charging of plug-in EVs is prepared by using the non-cooperative games. In [9], a Threshold Admission with Greedy Scheduling (TAGS) algorithm is presented as a model of harging problem with admission control and charging capacities. This approach includes a reserve dispatch algorithm for compensating the intermittency of renewables. Integration of electric vehicles into power grid will also increase the peak demand. To decrease the demand in peak times, the load shifting methods are applied to the charging processes of electric vehicles.

B. Renewable Energy Sources

New smart grid concepts increase the efficiency of renewables, some of which are wind turbines, solar Photo-Voltaics (PVs) and biofuel cells. In Germany, renewable energy sources are expected to supply minimum 33% of the primary energy consumption by 2020 [10]. Centralized traditional power plants meet the power demand at the expense of high operation costs whereas renewable energy sources operate relatively at lower costs. Several studies are conducted to increase the efficiency of the grid operations of renewables. Different algorithms to predict the daily distribution of solar irradiation and hourly distribution of wind speed are compared in [11]. Performances of algorithms are compared and ranked to select the best suitable algorithm for renewable power management. Power shifting on demand side relying on generation of renewables provides high economic benefits for end-users. In addition to the cost benefits, renewable energy sources cut down high carbon emission, causing greenhouse effect.

C. Virtual Power Plants

Virtual Power Plant (VPP) refers to a diverse pool of wholesale renewable energy sources, electric vehicles and energy storage systems in supply side. According to the flexible power plant concept, numerous small-scale DER meet the power demands with sustainable power loads. VPPs are so modular that sources can be swapped in or out depending on the demands. The results of the study [12] shows that load reduction obtained with control schedules minimize the transmission congestion in the grid. In other words, flexibility of the VPP decreases the transmission congestion in distribution grid. In another study [13], an optimization algorithm is developed to manage a VPP composed of a large number of customers with thermostatically controlled appliances.

D. Micro Grids

In micro grids, distributed generation and storage units are tied together to the common feeder and this feeder is linked to the single point of common coupling, providing connection to the larger grid [14]. Micro grid concepts bring solutions to the several problems of the current electric grid. In case of any outages and blockages in utility grid, micro grids island from the utility grid and continue to serve for the power needs. In the event of any incidents, control system kills or activates the power stream according to the priority level of end users. Prioritized customers such as hospitals or telecommunication companies are fed with sustainable power load during the electricity interruption. Self-healing of micro grids increases considerably the reliability of the power grid. The control schemes required for operation of micro grid in grid-connected and islanded modes are studied in [15]. In micro grids, advanced power electronic systems control the operations and services among DER and surrounding Alternating Current (AC) electric system. The power electronic technology enables the micro grid operators to access smart meters and other intelligent field devices so that ancillary services are delivered to the electricity consumers. In this context, micro grids are regarded as autonomous sub-grids, providing its own power demands and related supportive services.

E. Demand Response Management

Customers request power load with scheduling the operations of household appliances. Home energy management system enables end-users to participate in planning and scheduling the power management. Smart appliances are the integral components of the smart grid infrastructure. These intelligent household appliances are connected to the grid via smart meters which are electronically controlled interface devices. Smart meters have capability to receive and respond to the command signals coming from management systems of micro grids and utility grids. In doing so, demand side systems schedule the energy consumption of the electronic devices based the market prices and power capacity levels of supplies. Demand response mechanism encourages the energy consumers to shift the operations of the smart appliances from on-peak hours to off-peak hours. In [16], a heuristic is developed, allowing electricity demands to be shifted according to the dynamic energy market conditions.

F. Flexible Manufacturing System

Increasing energy costs and political policies on environmental precautions influence the methods of developing production plans. Achieving production targets and energy savings are two determinants for optimization of production plans. Several energy efficiency approaches are applied for the different manufacturing scenarios. In [17], the production plan achieving targets with minimum energy consumptions is chosen among alternative production plans. Process scheduling software determines the time slots, on which manufacturing processes use up the cheapest energy. While creating production plan, the scheduling software considers the physical constraints and interdependencies of the manufacturing processes. In advance of the implementation, the configuration of selected production plan can also be modified by the human operators. Overall, smart grid encourages the procurement of cheaper and less energy in the manufacturing industry.

G. Grid Management Model for JIAC Agent Implementations

New technologies arising in smart grid are integrated to the existing grid with innovative physical and cyber components [18]. Electric grid equipped with smart devices communicating with the control center enables efficient remote grid controlling and monitoring.

Smart metering system collects the energy data from smart meters and measured data are concentrated by smart grid data concentrators. Those collected data are unified by the central system called Meter Data Management (MDM) system. Some of end users in the grid are industrial companies, which plan manufacturing processes, or homeowners, who control their electricity consumption with use of home energy assistant systems. Data sets unified in central system are processed by the corresponding subsystems aiming at providing services for particular end users. Grid service providers in the proposed model are Residential Demand Management (RDM), Self- healing and Outage Management (SOM), Flexible Manufacturing Management (FMM), Electric Vehicle Charging/Discharging Management (EV - CDM) 1. The smart autonomous grid entities are represented as JIAC agents and these agents are scattered over the telecommunication network such as Local Area Network (LAN) and General

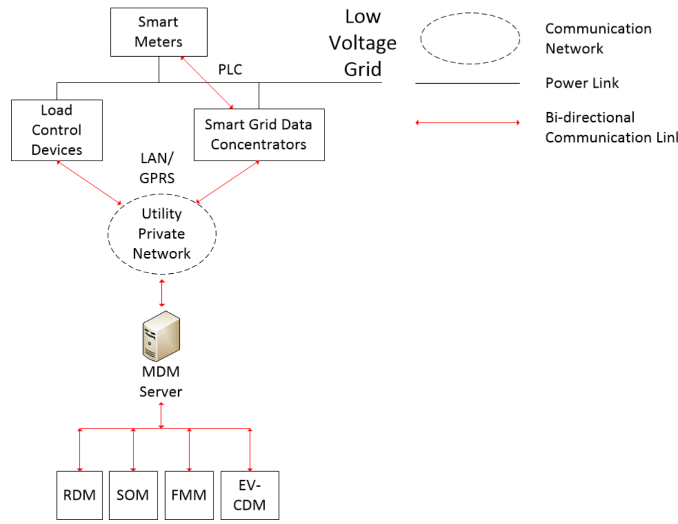

Figure 1. Structure of Grid Management Model

Packet Radio Service (GPRS). Figure 1 shows the coordination of the smart meters and the smart grid data concentrators via Power Line Communications (PLC). JIAC agents, deployed on grid control devices, grid concentrators, and software components exchange JIAC messages by complex interaction protocols. Proposed grid management models include the control schemes and algorithms which are incorporated with intelligent agents. We will present smart grid applications based on JIAC agent framework in the next section.

IV. JIAC Agents in Smart Grid Applications

In the previous section, we outlined challenges and possible (novel) areas of application that evolve from an ever increasing development of smart grid architectures. Over the last years, we developed several applications that fall under the umbrella of the one or the other category of evolving business concepts that we have outlined above.

It is our experience, that the multi-agent paradigm works very well for the smart grid domain, simply because of the distributed nature of most smart-grid-related problems. In this section, we present selected applications in more detail and respectively reference the categories (from the previous section), that our applications belong to. All applications that we present here are based on the JIAC agent framework [3]. JIAC agent framework is integrated with service-oriented paradigm. The framework provides the agent-platform, comprising agent nodes, physically distributed and it enables runtime environments for JIAC agents. JIAC agent multi-system can be defined by one or more spring configuration files. Runtime environment can be monitored and controlled by Java Management Extension Standard (JMX).

Our decision to use JIAC was based on a couple of factors. First, as most project were developed in an academic context, we required a free software framework. Furthermore, most applications were supposed to be executed in real life applications, thus, reliability and a comprehensive set of development tools and debugging capability was required as well. JIAC provides most of these features and had already been used in series of industrial applications [3]. For these reasons, we decided to use JIAC for our applications.

We proceed by presenting selected smart grid applications in more detail. In doing so, we respectively emphasize the distributed nature of the problem(s) that we tackled. The following presentations as well as additional project descriptions can be found in Lützenberger et al [19].

A. ILIas

The *Intelligente Lösungen zum Schutz vor Kaskadeneffekten in voneinander abhängigen kritischen Infrastrukturen* (ILIas) project [20] is based on the assumption that modern infrastructures are interdependent. Consider power- and telecommunications grids as examples for interdependent infrastructures. This interdependency may cause cascading effects in all involved infrastructures, whenever failures occur. ILIas aims to research and create intelligent and scalable management systems that provide prediction and reaction to cascading failure effects, so that actions to stabilize the managed infrastructure can be taken. Figure 2 presents the general architecture of the ILIas framework.

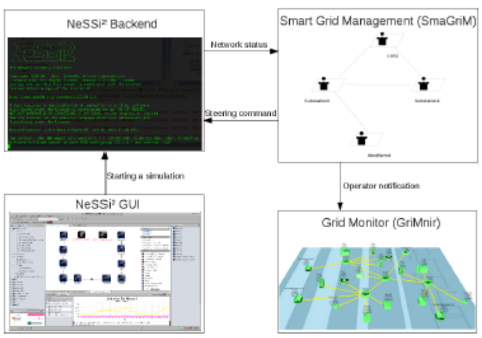

Figure 2. Overview over the ILIas Framework

An example for this is the reaction to power outages and the consequent failure of telecommunication networks in the affected areas. In ILIas, we developed an agent-based decentralized smart grid management system, which observes and controls the grid. Grid behaviour is predicted by means of simulation. Power and telecommunication networks are simulated by the agent-based simulation framework, Network Security Simulator, (NeSSi^2) [21].

NeSSi^2 is developed upon the JIAC framework, resulting in a distributed and easy-to-extend architecture. Detection algorithm plugins, traffic analysis, and automated attack generation are some of the capabilities, allowing for security research and evaluation purposes. We distinguish between an offline simulation, in which pre-defined scenarios are simulated, and an online simulation, in which the current grid state is used as a starting point for the calculation of predictions. The smart grid management is agent-based, such that JIAC agents in a Peer-to-Peer (P2P) network take over management duties for one smart grid entity. The distributed assembly allows our management system to work through partial or total power outages and to deal with sub-networks, which are either connected or independent from each other (islanded). We allowed our management agents to communicate with both physical as well as simulated smart grid entities. This design eases the testing—even the testing of large scale systems. Furthermore, we put particular emphasis on support for human operators.

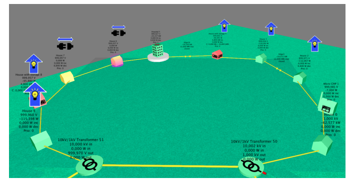

Figure 3. Smart grid live monitoring in ILIas (image detail)

Visualization is provided by our monitoring application (Figure 3), which shows the smart grid topology, live sensor data, the results of a situation analysis and action recommendations. The monitoring application can be used by human operators in order to check the current grid state and acknowledge or deny the recommended actions. Both, situation analysis and action recommendations are provided by the JIAC rule engine. Intelligent solutions developed in ILIas project provide electricity and telecommunication services to general public.

B. Gesteuertes Laden V2.0

The goal of Gesteuertes Laden V2.0 [22] was to use electric vehicles as mobile and distributed energy storages in order to increase the utilization of wind energy and to balance energy grids. The focus of our application was on the driver, such that mobility was ensured at all times. In order to account for this, we developed an application, which optimizes charging- and feeding processes of electric vehicles with respect to the driver's daily schedule. The application was implemented as a live system including real vehicles and charging stations.

The application [23] was designed as a distributed mobility and energy management system, in which relevant stakeholders (e.g., the driver, the vehicle manufacturer, the energy provider, the charging station, and the grid operator) were represented by JIAC Agents. The JIAC agents, distributed over the Gesteuertes Laden V2.0 architecture is shown in figure 4.

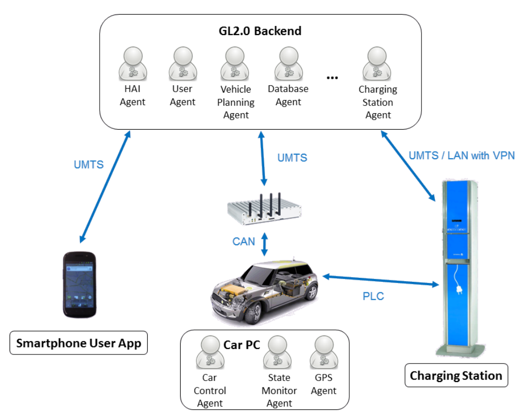

Figure 4. Architecture of Gesteuertes Laden V2.0

The system produces day schedules for the drivers. These schedules include journeys as well as charging- and feeding events [24]. In order to calculate these schedules, stakeholder-specific preferences and constraints had to be considered. Examples for such properties are: the driver's appointments, wind forecast, characteristics and current state of the electric vehicles, characteristics and availability of charging stations, and power grid constraints. In total, the system comprises twelve JIAC agents. These agents were permanently running on a back-end server and on the hardware of participating vehicles. In addition, more than 100 services were executed simultaneously, providing different functionality, from simple information services to complex planning algorithms. Beyond that, users and vehicles were represented by JIAC agents, which were deployed on the back-end server, taking the main responsibility for developing user- and vehicle schedules.

The data exchange between the electric vehicles and back-end agents was safeguarded by failover mechanisms in order to guarantee a reconnection after network-failures. Charging- and feeding events are triggered by the electric vehicle agents, which interact with the charging stations via power line communication. We also integrated external services, such as wind forecasts into our system. For this purpose, we used the *Web Service Gateway*, which is also provided by JIAC. An My Structured Query Language (MySQL) database was included as well. This connection was done with a generic database agent. We put particular emphasis on allowing users to participate in the scenario, thus, a smartphone application was developed. In order to connect this interface to the agent system, the *Human Agent Interface*, a JIAC component, was used. Finally, a monitoring component was installed to observe all agents and to notify developers in case services are down. The system was evaluated within a three-week field test [25]. To summarize, Gesteuertes Laden V2.0 brings in new approaches for integration of V2G to maximize the use of renewable energy for power grid stabilizaton.

C. The EnEffCo Project

In the Energy Efficiency Controlling (EnEffCo) project, we developed an optimization software [17, also [26]] that optimises production processes of the automotive industry in terms of energy costs. Optimization schema for production processes is illustrated in figure 5. The application uses the fact that (at least in Germany), industrial players are able to buy energy by means of short-term strategies at the day-ahead energy market. The prices at the day-ahead market are highly dynamic, thus, it is possible to minimize energy costs by shifting energy-consuming sub-processes to time slots with less expensive prices. Since contemporary production processes are highly complex and comprise many co-depending sub-processes, shifting parts of the overall process is not trivial.

We implemented our solution based on *Evolution Strategy* to produce reliable optimization results quickly. Entities were designed as JIAC agents, such that a server agent receives information about the production process in form of a bipartite graph. The graph contains architectural information, e.g., the sequential arrangement of all production steps, involved machinery and storages, but also meta information, such as the duration and the energy consumption of production steps. In addition, the server agent receives a production target, a timeframe and information on the energy price development. Based on this data, the initial production plan is generated (Figure 5). This initial plan is mutated, such that sub-processes are randomly shifted. Each server agent produces a defined number of mutations. Quality assessment is done by means of fitness function, which selects the most effective production plans. These plans are used as input for the next stage of evolution, where the production plans are mutated again. The process terminates whenever the quality of the best mutations remains steady across several generations.

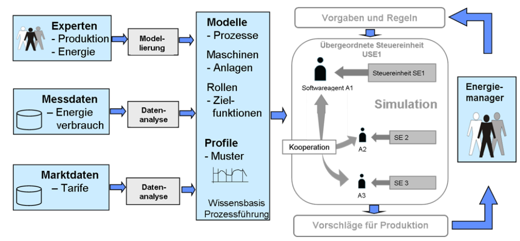

Figure 5. Energy Efficiency Controlling in the Automotive Industry

In addition to the server agent, we developed an optimization client (also a JIAC agent), which broadcasts an optimization problem by means of a custom protocol. Available server agents either accept and initiate the optimisation process or refuse as a result to different reasons. Results are sent back to the client, who compares available results and selects the best. The best production schedule is presented to the operator, who is able to configure his production line accordingly.

JIAC facilitated the development in two aspects. First, JIAC allows for the execution of several optimization agents in parallel. This capability can be used to counter the problem of stochastic optimisation algorithms to get stuck in local optima. It is our experience, that the simultaneous execution with different initial populations and random mutations increases the chance to overcome local maxima significantly. In doing so, the application's execution speed remains equal while the quality of the optimisation increases. The reason for this is that JIAC agents are truly multithreaded, capable of running on different Central Processing Units (CPUs). The second reason for the appliance of JIAC is the reliable communication of the agent framework. It is possible to distribute agents among different hosts and thus to increase the overall performance of the optimization software. As an example, complex optimisation scenarios can be supported by additional servers. These servers can be removed when this additional calculation performance is not required any more. The reason for this flexibility is the loose coupling between JIAC agents, which simplifies this process greatly and also supports maintenance issues. Erroneous or obsolete server agents can be replaced individually, without shutting down the entire application. To sum up, outcomes,which are gained from temporal shifting of energy demand support the manufacturing companies in the automative industry in terms of reducing energy consumption.

V. Conclusion

Smart grid architectures involve many autonomous systems with partial knowledge that have to communicate and cooperate in order to solve problems that each one of them for itself is not able to solve. It is difficult to provide a generic perspective on smart grid architectures, but in most applications that we developed, we designed energy consuming- and energy

 ISBN: 978-1-61208-332-2

producing entities as agents. These agents represent relevant stakeholders that are capable to make decisions. We design these agents to pursue their goal to allocate required, or to dispose produced energy, respectively. To achieve this goal, agents communicate, cooperate, or negotiate and determine their actions themselves. Especially, in ILIas and Gesteuertes Laden 2.0, we had to account for aspects of data security. Furthermore, if one develops solutions for smart grid architectures, one has to deal with real distribution. As an example, consider the Gesteuertes Laden project, in which decision making software was deployed on electric vehicles, charging stations and backend servers. The main advantage to use agent technology for the implementation of these systems is the development support. Agent-oriented approaches for smart grid problems are modelled and implemented in described projects.

REFERENCES

[1] E. Meer, "Smart grid estimated to reduce co2 emission by up to 2.2 gigatonnes by 2050." [retrieved: 01, 2014]. Available from: http://www.prweb.com/releases/2012/11/prweb10169445.html, 2013.

[2] B. Hirsch, T. Balke, and M. Lützenberger, "Assessing agent applications — r&D vs. R&d," in Multiagent Systems and Applications — Volume 1:Practice and Experience, ser. Intelligent Systems Reference Library, M. Ganzha and L. C. Jain, Eds. Springer Berlin / Heidelberg, 2013, vol. 45, pp. 1–20.

[3] M. Lützenberger, et al., "JIAC V — a MAS framework for industrial applications (extended abstract)," in Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems, Sait Paul, Minnesota, USA, M. G. T. Ito, C. Jonker and O. Shehory, Eds., May. 2013, pp. 1189–1190.

[4] K. Barnes, B. Johnson, and R. Nickelson, "Review of Supervisory Control and Data Acquisition (SCADA) Systems," Idaho National Engineering and Environmental Laboratory, Idaho, Internal Technical Report, January. 2004.

[5] C. H. Hauser, D. E. Bakken, and A. Bose, "A failure to communicate: next generation communication requirements, technologies, and architecture for the electric power grid," Power and Energy Magazine, IE, vol. 3, no. 2, March. 2005, pp. 47–55.

[6] T. Yiyun, L. Can, C. Lin, and L. Lin, "Research on vehicle-to-grid technology," in International Conference on Computer Distributed Control and Intelligent Environmental Monitoring, February. 2011, pp. 1013–1016.

[7] R. Jin, B. Wang, P. Zhang, and P. B. Luh, "Decentralized online charging scheduling for large populations of electric vehicles: a cyber-physical system approach," in International Journal of Parallel, Emergent and Distributed Systems, January. 2012, pp. 1–17.

[8] Z. Ma, D. S. Callaway, and I. A. Hiskens, "Decentralized charging control of large populations of plug-in electric vehicles," IEEE Trans. Control Syst. Technol., vol. 21, no. 1, January. 2013.

[9] S. Chen, Y. Ji, and L. Tong, "Deadline scheduling for large scale charging of electric vehicles with renewable energy," in IEEE 7th Sensor Array and Multi-channel Signal Processing Workshop, June. 2012, pp. 13–16.

[10] J. Burgermesiter, "Germany: The world's first major renewable energy economy." [retrieved: 01, 2014]. Available from: http://www.renewableenergyworld.com/rea/news/ article/2009/04/ germany-the-worlds- first-major-renewable-energy-economy, pp. 67–78, June. 2009.

[11] G. M. Shafiullah, M. T. Amanullah, D. Jarvis, M. S. Ali, and P. Wolfs, "Potential challenges: Integrating renewable energy with the smart grid," in 20th Australasian Universities Power Engineering Conference (AUPEC), 5-8 December. 2010, pp. 1–6.

[12] M. Wooldridge, "Intelligent agents: The key concepts," in Multi-Agent Systems and Applications II, ser. Lecture Notes in Computer Science, H. K. V. Mařík, O. Štěpánková and M. Luck, Eds. Springer Berlin Heidelberg, 2002, vol. 2322, pp. 3–43.

[13] N. Ruiz, I. Cobelo, and J. Oyarzabal, "A direct load control model for virtual power plant management," IEEE Transaction on Power Systems, vol. 24, no. 2, March. 2009, pp. 959–966.

[14] P. Asmus, "Microgrids, virtual power plants and our distributed energy future," The Electricity Journal, vol. 22, no. 3, April. 2010, pp. 72–82.

[15] A. Moallem, A. Bakhshai, and P. Jain, "Local smart micro grids," in 2011 IEEE 33rd International Telecommunications Energy Conference (INTELEC), October. 2011, pp. 1–8.

[16] M. Deindl, C. Block, R. Vahidov, and D. Neumann, "Load shifting agents for automated demand side management in micro energy grids," in SASO '08. Second IEEE International Conference on Self-Adaptive and Self-Organizing Systems, October. 2008, pp. 487–488.

[17] T. Küster, M. Lützenberger, and D. Freund, "Distributed optimization of energy costs in manufacturing using multi-agent system technology," in Proceedings of the 2nd International Conference on Smart Grids, Green Communications and IT Energy-aware Technologies, St. Maarten, St. Maarten, P. Lorenz and K. Nygard, Eds. IARIA, March. 2012, pp. 53–59.

[18] A. St. Leger, J. John, and D. Frederick, "Smart grid modeling approach for wide area control applications," in IEEE Power and Energy Society General Meeting, San Diego, CA. IARIA, 22-26 July. 2012, pp. 1–5.

[19] M. Lützenberger, et al., "A multi-agent approach to professional software engineering," in Engineering Multi-Agent Systems — First International Workshop, EMAS 2013, St. Paul, MN, USA, May 6–7, 2013, Revised Selected Papers, ser. Lecture Notes in Artificial Intelligence, M. Cossentino, A. E. F. Seghrouchni, and M. Winikoff, Eds., vol. 8245. Springer Berlin / Heidelberg, 2013, pp. 158–177.

[20] T. Konnerth, et al., "Integration of simulations and MAS for smart grid management systems," in Proceedings of the 3rd International Workshop on Agent Technologies for Energy Systems(ATES), Valencia, Spain, 2012.

[21] D. Grunewald, et al., "Agent-based network security simulation (demonstration)," in Proceedings of the 10th International Joint Conference on Autonomous Agents and Multiagent Systems, L. S. K. Tumer, P. Yolum and P. Stone, Eds., May. 2011, pp. 1325–1326.

[22] Vattenfall, BMW, TU Berlin, TU Chemnitz, and TU Ilmenau, "Increasing the effectiveness and efficiency of the applications wind-to-vehicle (W2V) and vehicle-to-grid (V2G) including charging infrastructure (Managed Charging V2.0)," Technische Universitätsbibliothek Hannover (TIB), 2011.

[23] J. Keiser, M. Lützenberger, and N. Masuch, "Agents cut emissions - on how a multi-agent system contributes to a more sustainable energy consumption," Procedia Computer Science, vol. 10, no. 0, August. 2012, pp. 866–873.

[24] N. Masuch, J. Keiser, M. Lützenberger, and S. Albayrak, "Wind Power-Aware Vehicle-to-Grid algorithms for sustainable EV energy management systems," in 2012 IEEE International Electric Vehicle Conference (2012 IEVC), Greenville, South Carolina, USA, Mar. 2012.

[25] M. Lützenberger, J. Keiser, N. Masuch, and S. Albayrak, "Agent based assistance for electric vehicle- an evaluation," in International Conference on Active Media Technology, Macau, e. a. R. Huang, Ed., vol. 7669. Springer Berlin / Heidelberg, December. 2012, pp. 145–154.

[26] T. Küster, M. Lützenberger, and D. Freund, "An evolutionary optimization for electric price responsive manufacturing," in Proceedings of the 9th Industrial Simulation Conference, Venice, Italy, S. Balsamo and A. Marin, Eds., Eurosis. EUROSIS-ITI, June. 2011, pp. 97–104.

 ISBN: 978-1-61208-332-2