Report from Dagstuhl Seminar 15091

Smart Buildings and Smart Grids

Edited by

Hans-Arno Jacobsen[1], Randy H. Katz[2], Hartmut Schmeck[3], and Christoph Goebel[4]

1 **TU München, DE,** arno.jacobsen@tum.de
2 **University of California – Berkeley, US,** randy@cs.berkeley.edu
3 **KIT – Karlsruher Institut für Technologie, DE,** hartmut.schmeck@kit.edu
4 **TU München, DE,** christoph.goebel@tum.de

Abstract

This report provides an overview of the program, discussions, and outcomes of Dagstuhl Seminar 15091 "Smart Buildings and Smart Grids", which took place from 22–27 February 2015 at Schloss Dagstuhl – Leibniz Center for Informatics. The main goal of the seminar was to provide a forum for leading Energy Informatics (EI) researchers to discuss their recent research on Smart Buildings and Smart Grids, to further elaborate EI research agenda and methods, and to kick-start new research projects with industry. The report contains abstracts of talks that were held by the participants and the outcomes of several discussion sessions on the focal topics of the seminar (e.g., information technology driven developments in building and power system management, as well as cross-cutting topics, such as computer networks, data management, and system design.

Seminar February 22–27, 2015 – http://www.dagstuhl.de/15091
1998 ACM Subject Classification D.0 [Software] Software, G.0 [Mathematics of Computing] General, H.0 [Information Systems] General, J.2 [Computer Applications] Physical Sciences and Engineering – Engineering
Keywords and phrases Energy Informatics, Smart Grids, Smart Buildings, Cyber-Physical Systems
Digital Object Identifier 10.4230/DagRep.5.2.128

1 Executive Summary

Hans-Arno Jacobsen
Randy H. Katz
Hartmut Schmeck
Christoph Goebel

Motivation

Motivated by the increasing importance of producing and consuming energy more sustainably, a new and highly dynamic research community within computer science has evolved: Energy Informatics (EI). Researchers active in the EI field investigate information age solutions for monitoring and controlling large cyber-physical infrastructures with a focus on the following goals: (i) an overall reduction of the energy consumption of these infrastructures, and (ii) the integration of distributed renewable energy sources into these infrastructures. This seminar focused on two use cases of existing cyber-physical systems, buildings and power grids. These use cases were chosen due to their relevance in terms of energy footprint. The seminar has three major goals: (i) to provide a forum for leading EI researchers to discuss their recent research on Smart Buildings and Smart Grids, (ii) to further elaborate EI research agenda and methods, and (iii) to kick-start new research projects with industry.

Smart Buildings. Modern buildings already incorporate increasingly sophisticated Building Management Systems (BMS) that integrate building control with improved sensors and better data collection and presentation capabilities. However, these systems currently only enable simple, decoupled control of building services like lighting, ventilation, heating and cooling. Their architectures and Application Programming Interfaces (APIs) are not standardized, and often proprietary: only the BMS vendor can add functionality. This slows the pace of innovation, buildings remain rigid in the functions and services they provide, and their quality and effectiveness remain difficult to quantify. Contemporary BMS attempt to achieve global service levels based on local control instead of meeting individual occupant requirements based on global control. Standardized building management APIs and scalable middleware solutions that enable reliable communication between building sensors, users, control systems, and machinery could accelerate energy innovation in the building sector.

Smart Grids. Contemporary electricity grids and markets were designed for a scenario in which large and mostly fossil-fueled power plants are dispatched to meet an almost inflexible demand. Achieving sustainable energy supply, however, requires moving towards a scenario where the variable power supplied by distributed renewable resources like wind and solar has to be absorbed by supply-following loads and energy storage whenever it is available. Thus, instead of dispatching a relatively small number of large generators, the large-scale integration of new types of generators and loads into electric grids requires new types of information systems for monitoring and controlling them, while making efficient use of existing assets. The task of controlling large numbers of flexible loads, e.g., air conditioning systems in buildings, electric vehicles, and small-scale energy storage systems, while guaranteeing overall system stability, is highly demanding in terms of computational complexity, required data communication and data storage. In the Smart Grid space, the challenge faced by EI researchers is to develop and carefully evaluate new ideas and actual system components enabling Smart Grid systems that are scalable, efficient, reliable, and secure.

Organization of the Seminar

The week-long workshop plan was as follows. Day 1 introduced the attendees to each other, and set the stage through invited tutorial presentations and brainstorming sessions. Day 2 was spent in breakouts focused on identifying the research challenges and opportunities, organized by application area such as Smart Buildings or Energy Grids, based on attendee interest and expertise. On Day 2, we also held the first out of two presentation session, where participants could give a short overview about their current research. Day 3 was used to assess the workshop at mid-stream, conduct group discussions, and make necessary corrections. Initial writing assignments, to document the discussions of the breakout sessions, were made on this day, as well. Day 4 consisted of a second round of breakouts focusing on enablers and crosscutting issues (e.g., data management, system design patterns, and human machine interaction) and the second participants' presentation session. Work on completing the report draft continued on that day. The last day consisted of the reviewing of the report draft, and through group discussion, identify the summary findings and recommendations.

2 Table of Contents

3 Overview of Keynote Talks

3.1 A Composable Method for the Real-Time Control of Active Distribution Networks with Explicit Power Setpoints

Jean-Yves Le Boudec and Mario Paolone (EPFL, CH)

The classic approach to the control of medium and low voltage distribution networks involves a combination of both frequency and voltage controls at different time scales. With the increased penetration of stochastic resources, distributed generation and demand response, it shows severe limitations both in the optimal/safe operation of these networks, as well as in aggregating the network resources for upper-layer power systems.

To overcome this difficulty, we propose a radically different control philosophy, which enables low and medium voltage distribution networks as well as and their resources to directly communicate with each other in order to define explicit real-time setpoints for active/reactive power absorptions or injections. We discuss a protocol for the explicit control of power flows and voltage, combined with a recursive abstraction framework. The method is composable, i.e., subsystems can be aggregated into abstract models that hide most of their internal complexity.

Within this control framework we specifically analyze the case of a low-overhead decentralized Demand Response (DR) control mechanism, henceforth called Grid Explicit Congestion Notification (GECN), intended for deployment by distribution network operators (DNOs) to provide ancillary services to the grid by a seamless control of a large population of elastic appliances.

Contrary to classic DR approaches, the proposed scheme aims to continuously support the grid needs in terms of voltage control by broadcasting low-bit rate control signals on a fast time scale (i.e., every few seconds). Overall, the proposed DR mechanism is designed to i) indirectly reveal storage capabilities of end-customers and ii) have a negligible impact on the end-customer.

3.2 Bringing Distributed Energy Resources to Market

Christoph Goebel and Hans-Arno Jacobsen (Technische Universität München, DE)

The first part of this talk was meant to be a tutorial on current challenges in the operation of power systems induced by renewable integration. The second part provided an overview of our current research in this area.

In the first part, we introduced the different stakeholders involved in electricity generation, transmission, and consumption and how coping with the individual challenges they face requires the use of innovative information technology. We argued that better monitoring and control of renewables, grids, and the demand side competes with more traditional ways to deal with renewable integration challenges, e.g., aggressive grid expansion or construction of highly flexible power plants. In the long-term, with several developments happening at the same time, e.g., declining prices of batteries and cheap control infrastructures, the intelligent dispatch of distributed energy resources by aggregators could prevail in this competition. The task of such aggregators is to "make the most" out of the potential of various distributed energy resources while respecting the individual resource and wholesale market constraints.

In the second part, we presented several research challenges in the area aggregator systems. Among other things, these challenges include predictive capabilities, multi-resource/multi-purpose dispatch, profit maximization of aggregators, new optimization techniques, distributed system designs to achieve scalability and fault tolerance, and optimal data storage schemes for representing resource schedules and system states. We presented recent work in the area multi-purpose dispatch and new optimization techniques in more detail. We closed by motivating new research in the Smart Grid field, but also repeated the necessity for putting more common effort into the collection and consolidation of public data for research purposes and the development of open source tools, e.g., test beds and simulation environments. As starting points, we presented two such efforts we recently initiated at TUM.

3.3 The Software-Defined Building: A Technical Approach for Smart Buildings

Randy H. Katz (University of California – Berkeley, US)

The built environment is the human-made surroundings that provide the setting for human activity, ranging from buildings to cities and including their supporting infrastructures, such as for water or energy. Buildings are the starting point for Smart Cities. They represent a large component of modern economies. They are where we spend 90% of our time, consume 70% of our electricity (40% of our primary energy consumption) and generate 40% of our GHG emissions. The aggregate energy expenditure for the U.S. buildings stock exceeds $400 billion and construction represents more than 10% of the entire U.S. GDP. Over the last several decades, information-processing abilities have grown enormously; yet thus far this has had little effect on the built environment.

Given the widespread deployment of digital controls in commercial buildings, we can consider them to be Cyber-Physical Systems, and we can enhance their functionality through software programmability. The analogy is how phones are (almost) infinitely extended through network-connected applications. Traditional building facilities are stovepipes; e.g., HVAC is separated from lighting controls, even though there is a correlation between lit and conditioned spaces. Awareness of occupant location and activities offers new opportunities for the building to become aware of and respond to occupant needs, faster and with more efficiency. Occupant devices integrated with the building frees it from the limitations of mechanical and rigidly placed sensors and controls, and awareness of environmental factors like weather and sun orientation can be exploited to better control interior spaces. The building's information processing capabilities can become cloud integrated, providing unlimited capacity and the ability to extend control to fleets of buildings, neighborhoods or complete cities.

The key research need is the development of a new category of operational software, a Building Operating System, a City Operating System, etc., to provide a common foundation for abstracting and managing the resources of the Built Environment, providing integration with user devices and external information sources, data analytical processing, and enabling new kinds of control, information presentation, and planning applications. The advantage will be demonstrated by quantifiable improvements in such figures of merit as energy efficiency (including agile and intelligent interaction with the electric grid for energy flex), reduced cost of ownership (maintenance and management) and importantly, improved occupant satisfaction (comfort, indoor air quality, aesthetics, information transparency, e.g., how activities translate into energy consumption or savings; this can increase productivity of the workforce and sales for retail spaces), as well as enhanced controllability, agility, and extensibility via an open platform and technology ecosystem to achieve these goals.

3.4 Industrial Perspectives on Smart Buildings and IoT Impact

Milan Milenkovic (Intel – Santa Clara, US)

Buildings are one of the primary users of energy (40% of all energy and 70% of electricity according to US Department of Energy and other international agencies), but some 20-40% of that is wasted due to inefficiencies. We advocate an Internet of Things (IoT) based approach to solving this problem that reduces costs and provides numerous benefits. Commercial buildings today have a variety of automated systems to monitor and control different aspects of their behavior, including: HVAC and lighting (usually BMS), energy consumption, lifts, security and access, fire alarm, water management, parking, landscaping and irrigation, audio visual, digital signage. Coordinated behaviors of such systems can result in increased efficiencies, safety, and occupant comfort, such as monitoring of occupancy to dynamically adjust heating/cooling and lighting, or automatically moving elevator cars to ground floor for safety in case of imminent power failures. Most of the currently deployed legacy building control systems are isolated proprietary systems that do not interoperate because that would require prohibitively expensive custom interfacing.

Use of IoT in building control systems brings standards, lower costs and well known benefits of Internet – connectivity and interoperability, scalability, tested tools and design practices, faster/cheaper development and interoperability. This tutorial presents an end-to-end IoT system architecture and argues for interoperable sensor/control data and meta-data definition to facilitate collaboration among domains and building control systems in particular. We also introduce examples of building deployments with IoT gateways in commercial buildings (usually BMS) and residential (retrofit example) and rooftop HVAC units.

Smart Grid needs to interface with building control systems. They are a key load and, when instrumented, can provide detailed feedback on current and projected electricity usage, so production can be more adaptive and the two systems can balance consumption and production using detailed usage information, estimates (based on accurate building models and past behavior), and interact/balance in real-time using techniques like demand response.

3.5 Energy Management in Smart Homes

Florian Allerding, Birger Becker, and Hartmut Schmeck (KIT – Karlsruher Institut für Technologie, DE)

Motivated by the challenges of the energy transition ("Energiewende") like fluctuating power generation, uncertainty of supply and inherent decentralization, this talk presents approaches to smart home management systems with a focus on energy management. While some of these (commercially available) systems address monitoring and remote control only, others include learning user habits.

As a particular example of a research prototype running in some test locations, the Organic Smart Home of KIT is presented, which includes an Energy Management Panel for visualizing the current energy situation and for discovering or specifying user preference for the operation times of appliances. An observer/controller-based architecture then optimizes the operation times of smart home appliances and other energy relevant devices, complying with the degrees of freedom specified by the residents of the home and considering also external information about time varying power tariffs and potential power limits.

Furthermore, an outlook is given on extensions to regional energy management by organizing a large number of devices (like appliances, CHPs, electric vehicles, heat pumps) into a pool in order to provide a cascaded form of responses for coping with spontaneous deviations from power schedules. Concluding, a number of questions is presented which are relevant for the design of Energy Management Systems in Smart Homes.

4 Overview of Participant's Talks

4.1 Organic Smart Home Energy Management and Building "Operating System"

Florian Allerding (KIT – Karlsruher Institut für Technologie, DE)

The Organic Smart Home is a flexible, generic "operating system" for Smart Buildings in real-world applications, which is already in use in households and office buildings. The major contribution is the design of a "plug & play"-type Evolutionary Algorithm for optimizing and management distributed generation, storage and consumption using a sub-problem based approach. Relevant power consuming or producing components identify themselves as sub-problems by providing an abstract specification of their genotype, an evaluation function and a back transformation from an optimized genotype to specific control commands. The generic optimization respects technical constraints as well as external signals like variable energy tariffs.

4.2 Direct Control of Demand Flexibility: Applicability of Batch Reinforcement Learning

Bert Claessens (Vito – Antwerp, BE)

In this talk, I presented recent work that contributes to the application of Batch Reinforcement Learning (RL) to demand response. In contrast to conventional model-based approaches, batch RL techniques do not require a system identification step, which makes them more suitable for a large-scale implementation. This talk discussed how fitted Q-iteration, a standard batch RL technique, can be extended to the situation where a forecast of the exogenous data is provided. In general, batch RL techniques do not rely on expert knowledge on the system dynamics or the solution. However, if some expert knowledge is provided, it can be incorporated by using our novel policy adjustment method. Finally, we tackled the challenge of finding an open-loop schedule required to participate in the day-ahead market. We proposed a model-free Monte-Carlo estimator method that uses a metric to construct artificial trajectories and we illustrate this method by finding the day-ahead schedule of a heat-pump thermostat. Our experiments showed that batch RL techniques provide a valuable alternative to model-based controllers and that they can be used to construct both closed-loop and open-loop policies.

4.3 Load Prediction of Non-Controllable Household Devices

Christoph Doblander (TU München, DE)

While many devices in the household could be controlled and the energy usage is known, the majority of the energy is consumed by devices which are not controllable. To increase self-sufficiency, we propose an architecture where the energy consumption of non-controllable loads is predicted and taken as an input for a control algorithm which controls the other devices to maximize for self sufficiency. One benefit of such a system is financial, since the incentive to feed rooftop solar energy back into the grid are declining.

We evaluate multiple machine learning algorithms, support vector machines, naive bayes and benchmark against persistence. We also evaluate the prediction error reduction when additionally features are extracted from the time series. The predictions are done on a sliding window, e.g., every minute, the load of the next 15 minutes is predicted based on the last 15 minutes, hence supporting a scenario in which a control algorithm continuously optimizes the actuation of the controllable devices.

The data was collected within the field trial of the "PeerEnergyCloud" project. Roughly 30 households were equipped with up to 7 plug meters. The results suggest that additional feature extraction reduces the prediction error and the benchmark persistence can be beaten. The evaluation of the prediction algorithms was done on 20 different households which allows to derive significant conclusions. The prediction error was measured by the Mean Absolute Percentage Error (MAPE) to compare it to existing literature

4.4 Are Energy Markets Efficient? The Case of Real and Virtual Storage

Nicolas Gast (INRIA – Grenoble, FR)

The electrical grid of the future will require more storage to compensate for the intermittency of distributed generators (such as solar, wind, combined heat and power). Storage will be real (batteries, water reservoirs) or virtual (demand response). In this talk, we analyzed the impact of storage of the real electricity markets, using several stochastic models. We showed that there exists a market price such that selfish users are provided with incentives to control their appliances in a socially optimal way. However, by setting these prices, users have an incentive to install a sub-optimal quantity of storage.

4.5 Loose Coupling Approach to Demand Response for Distribution Networks

Kai Heussen (Technical University of Denmark – Lyngby, DK)

In this talk, I reviewed our recent research in the area of congestion control in distribution grids. In particular, I focused on the necessary coordination between aggregators and distribution grid operators. The insights of our research have led to the development of FLECH, a market place for flexibility.

4.6 Information Systems and Science for Energy (ISS4E) at the University of Waterloo

Srinivasan Keshav (University of Waterloo, CA)

The current grid suffers from several problems, ranging from a high carbon footprint to very coarse control of loads. These can be solved using three key technologies: solar energy, energy storage, and the Internet of Things. In this talk, I discussed the difficult challenges that need to be solved using these technologies, such as meeting stochastic loads using stochastic generation and the need for control over multiple time scales. I then outlined three approaches being taken by the ISS4E group at Waterloo (http://iss4e.ca) to meet these challenges: a) using the Internet as an inspiration for Smart Grid architecture b) analysis of Smart Grid data sets and c) using Internet technologies for smart sensing and control.

4.7 Integrated Simulation of Power and ICT Systems

Johanna Myrzik (TU Dortmund, DE)

Co-simulation of ICT and power systems becomes increasingly important to develop and test Wide-Area Monitoring, Protection, and Control (WAMPAC) applications. In this talk, I presented INSPIRE, an integrated simulation of power and ICT systems for real-time evaluation. I provided insights into the architecture of INSPIRE and presented selected simulation results obtained in different control scenarios.

4.8 Introduction of the Grid4EU Project

Peter Noglik (ABB AG Forschungszentrum Deutschland – Ladenburg, DE)

In this talk, I gave a short overview of a large EU project, Grid4EU, which focuses on the large scale demonstration of advanced Smart Grid solutions with wide replication and scalability potential for Europe. It focuses on how distribution system operators can dynamically manage electricity supply and demand, which is crucial for integration of large amounts of renewable energy.

4.9 A Business Model for Scalable Demand Response

Anthony Papavasiliou (University of Louvain, BE)

I introduced ColorPower, a business model for scalable residential demand response (DR). The ColorPower software enables precision control of the consumer's demand flexibility, which it dispatches prioritized by customer preference. Demand management impact is divided into green (any time), yellow (peak periods), and red (emergencies). I discussed results from several experiments applying ColorPower, including automatic emergency DR and precision DR shaping.

4.10 Agent-Based Smart * Management Platform with Plug & Play

Yvonne-Anne Pignolet (ABB Corporate Research – Baden-Dättwil, CH)

Smart * Management System (SMS) to control and coordinate residential, commercial and industrial sites need to be flexible and support devices entering and leaving the network due to malfunctions, mobility or when new components are added. In these environments the burden of administration quickly overwhelms any potential benefit if the devices require explicit configuration in order to work together. Moreover, the embedded systems used in these scenarios have very severe requirements in terms of costs and sizes, offering very limited resources to any application running on them.

In this talk, we present a platform that requires no human intervention for the configuration and simplifies the addition of new devices. The platform enables the interaction of independent (distributed) agents with a publish subscribe architecture; different agents do different things in different places. This permits the system to grow as needed. In addition heterogeneous technologies are supported with an abstraction layer hiding the specific requirements of each appliance and offers a uniform interface to the agents above it. The user benefits from two technologies: SmartScript and SmartEnvironment. The first allows the user to write powerful building automation rules in a simple language, at a really high abstraction level. Users don't need to know how the appliances work, they only specify what they want the system to do. The second technology needs even less interaction with the user because it learns from the habits of the user and subsequently controls the appliances to increase the comfort and to reduce the energy consumption automatically according to the user behavior it observed. The whole architecture has been proven to work with inexpensive devices with low power consumption and very constrained HW resources.

4.11 Distributed Optimization in Smart Grids

Jose Rivera (TU München, DE)

The introduction of the Smart Grid will allow the active control of several devices: Electric vehicles, distributed storage units, distributed generation units and smart home appliances. This poses new challenges for operators of large power systems: How can they actively control large numbers of devices in a scalable, reliable and efficient way? This talk explored the contributions that distributed optimization can make towards answering these questions.

4.12 Low-Voltage Grid Control over Heterogeneous Communication Networks

Hans-Peter Schwefel (FTW Forschungszentrum Telekommunikation Wien GmbH, AT)

Voltage control in the Low-Voltage (LV) distribution grid can be performed by a hierarchical control architecture, in which a controller placed on the secondary substations, called Low-Voltage Grid Controller LVGC, communicates setpoints to the local controllers on assets in the corresponding parts of the LV grid.

In order to minimize the communication overhead and to maximize asset utilization, such LVGC is designed to be passive while the voltage is within certain boundaries, and only becomes activated when a sensor measures and communicates an exceedance of such voltage band. Example results from co-simulation of an example grid with photo-voltaic assets show that such voltage control is effective to reduce duration and extent of voltage events in the low-voltage grid. The design of the information exchange between assets and LVGC, however, has a strong impact on the performance of the controller: When using an adaptive monitoring framework to optimize the time instances at which asset information is requested by the controller, control performance can be significantly improved and robustness to communication network delays can be achieved. The adaptive monitoring scheme thereby uses as optimization target an information quality metric, which can be efficiently calculated based on asset dynamics and the (measured) communication network delays.

In the final part of the talk, additional challenges of control use-cases in the low-voltage and medium voltage grid were outlined and solution approaches followed up by the FP7 Research Project SmartC2Net were introduced: adaptive control approaches, adaptive grid and network monitoring, ICT capability analysis, communication network reconfiguration, and assessment approaches via analytic models, co-simulation models, and coupling of different lab test beds.

4.13 Revealing Household Characteristics from Smart Meter Data

Thorsten Staake (Universität Bamberg, DE)

Utilities are currently deploying smart electricity meters in millions of households worldwide to collect fine-grained electricity consumption data. We present an approach to automatically analyzing this data to enable personalized and scalable energy efficiency programs for private households. In particular, we develop and evaluate a system that uses supervised machine learning techniques to automatically estimate specific "characteristics" of a household from its electricity consumption. The characteristics are related to a household's socio-economic status, its dwelling, or its appliance stock.

We evaluate our approach by analyzing smart meter data collected from 4232 households in Ireland at a 30-min granularity over a period of 1.5 years. Our analysis shows that revealing characteristics from smart meter data is feasible, as our method achieves an accuracy of more than 70% over all households for many of the characteristics and even exceeds 80% for some of the characteristics.

The findings are applicable to all smart metering systems without making changes to the measurement infrastructure. The inferred knowledge paves the way for targeted energy efficiency programs and other services that benefit from improved customer insights. On the basis of these promising results, the paper discusses the potential for utilities as well as policy and privacy implications.

4.14 Smart Metering: What Drives the Impact of Behavior-specific Feedback

Verena Tiefenbeck (ETH Zürich, CH)

Transparency of consumption is supposed to foster energy efficient behavior, but the conservation effect in smart metering trials has been smaller than expected. I presented the results of a study involving 697 Swiss households on the impact of real-time hot water consumption feedback using a new metering device. We observed a stable average reduction of 23% of both energy and water use in the shower.

4.15 EV Fast Charging on German Highways

Victor del Razo (TU München, DE)

The limited driving range of electric vehicles makes them a sub-optimal alternative for long distance trips, particularly on highways, where higher driving speed has a negative effect on power consumption. In this talk, we discussed alternatives for using ICT for reducing the overall driving time, while keeping the required additional infrastructure at a minimum. Through a route optimization and charging time reservation system we reduce the trip duration and make the energy demand from power stations more predictable.

5 Reports from the Breakout Groups

5.1 Smart Grid Data Analytics

Bert Claessens, Nicolas Gast, Christoph Goebel, Mario Paolone, Anthony Papavasiliou, Jose Rivera, Joachim Sokol, Andreas Veit, and Holger Ziekow

Ecosystem

The smart (electric) grid ecosystem encompasses many stakeholders with well-defined tasks and objectives. These stakeholders include conventional and renewable power generation companies (gencos), electricity retailers, industry, aggregators, con- and prosumers (end customers on low voltage side), transmission system operators (TSOs), distribution system operators (DSOs), and market operators, technology providers, as well as organizations that play several stakeholder roles at the same time. The integration of renewable power generation into the electric grid leads to new challenges and opportunities in operational control, long-term capacity planning, and business strategy (including new business models). Information technology will play a major role in coping with these challenges as well as taking advantage of new business opportunities. This report focuses in particular on the role of measurement data and how current and future ICT applications of different stakeholders may take advantage of more and more of this data becoming available from different sources. The contents of this report are structured along the stakeholder axis. Applications including state-of-the-art and opportunities are therefore described on the stakeholder level.

Renewable Power Generation Companies

Renewable power generation companies include, e.g., wind farms, solar PV plants, and hydro power plants. In general, renewable gencos are interested in maximizing their power output, which in contrast to conventional generation depends on variable environmental conditions. This application is called maximum power point tracking. For instance, the blades of wind turbines can be adjusted to extract the maximum a. c. power from the wind based on the characteristics of the turbine. The mechanisms for maximum power point tracking could be further improved by data analytics, e.g., to coordinate the control of single generation units interconnected in larger wind farms or solar PV installations (advanced maximum power point tracking applications).

Wind and solar power generation is highly variable. Thus, if renewable gencos have to fully participate in electricity markets, which penalize deviations from scheduled market bids, more accurate predictions of their power output will become economically advantageous. While renewable gencos, similar to other gencos, monitor their total power production over time (market participation requirement), they often don't correlate this data with other potentially available data, e.g., weather data and forecasts. Moreover, many of them do not forecast output since they do not have economic incentives to do so if the receive a guarantee that their entire production is fed into the grid (which is the case under the current subsidy scheme in Germany, for instance). Output prediction is therefore an important data-based application that will become more important as soon as renewable gencos are forced to become regular market participants. Another important application based on the data collected from sensors in wind and solar power generation units is predictive maintenance, which allows gencos to predict the possible failure of generation equipment and therefore opens up cost savings opportunities in the maintenance area.

Current data sources:
- Metering and monitoring data from generators
- Generator configuration data

Current applications:
- Operational control
- Maximum power point tracking
- Health monitoring

Future data sources:
- Local weather data and forecasts
- Additional monitoring data from generators
- Historical configuration and power output data on the generation unit level

Future applications:
- Power output predictions on multiple time scales for market participation
- Advanced maximum power point tracking
- Predictive maintenance

Data volume estimate: from small to medium

Data velocity estimate: from small to medium

Transmission Grid Operators

Transmission grid operators are in charge of operating transmission grids (high voltage) within the secure region to prevent instability that could lead to blackouts. They own sophisticated models of the transmission grid which they use to infer the current grid state (e.g., current on all lines, voltage angles, etc.) using the available real-time data from generators, substations, high-voltage transmission lines, etc. This data is transferred to the

TSO's back-end system, which can happen at very high data rates based on the type of sensors deployed on the transmission grid level. For instance, phasor measurement units can transmit updated measurements at rates of up to 10 kHz.

To extrapolate the system state into the future and react to potential threats, TSOs access additional data sources, in particular weather data and market-cleared schedules of dispatchable generators. Their portfolio of countermeasures includes reconfiguration of the grid structure (e.g., by disconnecting transmission lines), activation of reserves, re-dispatch of generators, and load shedding. The more data can be accessed (higher granularity, special resolution, etc.), the higher the probability that the accuracy of statistical models will improve and enable more accurate extrapolations of the system state. Apart from operational control, TSOs are involved in long-term capacity planning. They are responsible for extending the transmission grid's capacity and provide reserves in close coordination with regulatory institutions. This capacity planning process is based on historical operation data as well as longer term models describing the evolution of the underlying variables.

Current data sources:
- Metering and monitoring data on the transmission level
- Weather data and forecasts
- Historical system states (currently produced at up to 30Hz)

Current applications:
- Real-time state estimation and visualization
- Contingency analysis (n-1)
- Decision support for operational control including reconfiguration, activation of reserves, generator re-dispatch, load shedding
- Reserve provision (usually via auctions)
- Decision support for long-term capacity planning

Future data sources:
- More detailed weather data and forecasts
- Additional monitoring data from generators, substations, and transmission lines using PMUs

Future applications:
- Advanced state estimation (probabilistic, multiple time scales, etc.)
- Advanced decision support
- Advanced long-term planning based on models and historical data

Data volume estimate: from low to high

Data velocity estimate: constantly high

Distribution Grid Operators

Distribution system operators are responsible for assuring the power quality and supply security in power distribution systems. They are currently able to monitor relevant metrics at substations, but have little visibility downstream, i.e., about the conditions at the end consumers. The can perform voltage regulation at the substation level by switching, transformer tap changes, or reactive power injections, but usually cannot remote-control any of the elements further downstream (e.g., protection or voltage regulators), which mostly operate independently. Distributed generation, load flexibility, and distributed energy storage will in the long term complicate the traditional tasks of DSOs described above.

Current data sources:
- Substation monitoring data

Current applications:
- Voltage regulation using tap transformers
- Switching at substations
- Long-term capacity planning (transformers, power lines, etc.)

Future data sources:
- Smart meter data
- Data from RTUs and PMUs deployed on the distribution level
- Local weather data and forecasts
- Detailed power output data from distributed generation, especially solar PV

Future applications:
- State estimation for distribution grids (probabilistic, multiple time scales, etc.)
- Advanced decision support from distribution system operation, in particular voltage regulation
- Advanced long-term planning based on models and historical data

Data volume estimate: from low to high

Data velocity estimate: from low to high

Aggregators

Aggregators bring the flexibility of their customers (industrial loads, thermal loads, bio gas plants, etc.) to market by buying the right to control them within certain limits. Due to the ongoing integration of variable renewable power generation into the grid, more flexibility is needed, which will eventually be reflected in more short-term trading and higher prices for flexibility, e.g., reserves. Once the large and obvious sources of flexibility (e.g., cooling houses, large commercial buildings, industrial loads with high flexibility) have been brought to market, smaller and less obvious resources may be accessed (e.g., smaller buildings, EVs, solar-attached storage). The data analytics requirements will therefore increase: The better aggregators can predict the capability of resources over time, the more efficiently they can dispatch these resources to maximize profits.

Current data sources:
- Resource monitoring data (e.g., temperature, power consumption, etc.)
- Resource meta data and models (e.g., battery capacity, charging/discharging rates, etc.)
- Market data (prices, bids, etc.)

Current applications:
- Market participation (reserve and spot markets) via resource dispatch

Future data sources:
- More detailed and accurate resource monitoring data, data from new types of resources (e.g., EVs, residential batteries, HVAC systems, etc.)
- Historical resource monitoring data and controls
- Weather data and forecasts

Future applications:
- Use of more accurate resource models for more profitable control / more efficient use of resource pool
- Advanced resource scheduling techniques applicable to large pools of heterogeneous resources (stochastic optimization, etc.)

Data volume estimate: from low to high

Data velocity estimate: constantly medium

Retailers

The business model of electricity retailers is to buy energy on the markets and sell it end consumers. While electricity prices on the wholesale markets vary over time, end consumers usually pay a fixed price. They can only make a profit if they manage to buy their electricity in a smart way and negotiate sufficiently low transmission and distribution prices with TSOs and DSOs. Imbalances of supply and demand decrease their profit since they result in penalties. Therefore retailers are interested in accurate predictions of electricity prices and demand. Electricity prices and demand will in the future depend more heavily on weather conditions, thus retailers may be interested in corresponding prediction services.

Current data sources:
- Meter data (nowadays usually measured 1-2 times a year)
- Customer data (address, payments, contract, etc.)
- Wholesale market prices

Current applications:
- Data analytics for market price and demand forecasting (OTC, spot market)
- Marketing, sales, billing

Future data sources:
- Smart meter data
- Weather data and forecasts

Future applications:
- Data analytics for market price and demand forecasting
- Data analytics for marketing and sales based on smart meter data
- More profitable market participation based on better supply and demand forecasting

Data volume estimate: from low to medium

Data velocity estimate: constantly low

Prosumers

Prosumers are end consumers of electricity that actively manage their consumption and may generate electricity on their own. As more accurate and detailed data on their own electricity usage becomes available to them, they can take advantage of innovative systems for consumption feedback and automatic energy saving mechanisms (NEST being a good example for such a system). If retailers offer time of use tariffs, such systems can also be used to shift load and save some money on the electricity bill. With more and more distributed generation and storage being deployed in the future, advanced energy management systems will enable even lower energy provision costs for prosumers. The higher the amount of relevant data that these systems can access gets, the more efficiently they can fulfill their purpose. Relevant data includes in particular data on the bounds of user comfort (e.g., temperature), user behavior (e.g., user at home or not), and environmental conditions (weather data and forecasts). It will be interesting to see if services can be established that transfer knowledge extracted from the pooled data of a given population of prosumers to new prosumers installing energy management systems in their homes.

Current data sources:
- Smart meter data
- Smart home sensors (dis-aggregated electricity demand, temperature sensors, etc.)

Current applications:
- Consumption feedback, e.g., to achieve higher energy efficiency
- Savings on electricity bill via load shifting (if time of use tariffs are available)
- Smart home control and automation (e.g., NEST)

Future data sources:
- Data from "personal sensors" – (e.g., smartphones, smart wristbands, etc.)
- Local weather data and predictions
- Derived preference data (rules), possibly based on populations of prosumers
- Derived activity data (rules), possibly based on populations of prosumers

Future applications:
- Advanced building energy management systems to achieve energy efficiency, load flexibility, and higher degree of autarky

Data volume estimate: from low to medium

Data velocity estimate: constantly medium

5.2 Smart Grid Communications

Jean-Yves Le Boudec, Srinivasan Keshav, Hermann de Meer, Florian Michahelles, Peter Noglik, Victor del Razo, and Kai Strunz

Summary

Communications are at the heart of the Smart Grid: sensing the status of grid assets and loads and controlling them depends critically on the availability of an underlying communication infrastructure. This infrastructure needs to span a wide geographic scale across multiple continents, yet provide a high density of communication endpoints, such as tens or hundreds of sensors in a single room.

Smart Grid communications are diverse in many dimensions: Smart Grid communications takes place over both private networks and public networks. Most communications today is over private networks, for security and reliability, but we expect that this will shift eventually to VPN communications over a public network. Communications use both wired and wireless links. The links are attached to a diverse set of endpoints – from tiny sensors to multi-million-dollar PMUs. Endpoints have a diverse set of performance requirements technologies used for communication links are also diverse, and depend on the geographic scale of communication.

Recommendations

Communication network technology is mature. Nevertheless, we have identified several avenues for research, that we discuss next. In the short term, we believe that there is the need for increasing the scale of communications at an economically-affordable cost per carried bit. This could be used, for instance, to carry sensor values over IP multicast or to obtain smart meter readings from millions of endpoints. There is also a need for reliable, very low-latency communication between controllers and sensors and controllers and actuators. In the mid-term, we see a strong need for integration of legacy wired/wireless technologies into IP for both WAN and LAN applications. We also see the need to upgrade and maintain the communication infrastructure and achieve reliability in communications, such as by authenticated sensor readings and tamper-proof communication. An interesting challenge is to assure communication despite loss of electrical power, so that the communication network can be used to restart the grid.

There is a need for research into achieving low-latency, low jitter, reliable, tamper-proof communication for sensing and control. In the long term, we believe that the research focus should be on infrastructure security and dependability. This would require solutions to problems such as securing devices despite physical access by attackers, and ensuring that emergency "back doors" to communication equipment does not compromise security. A more ambitious goal is to add intelligence to the communication network so that we evolve from a Smart Grid to a Semantic Smart Grid.

Appendix

5.2.1 Geographic Scale and Endpoints

Table 1 Taxonomy of Smart Grid communication.

	WAN/MAN – Public	WAN/MAN – Private	LAN (Private)
Wired	RTUs HEMS (Gateway) EV charger PV inverters	PMU Protection devices RTUs STN Sensors PV inverters	HEMS (Gateway) EV chargers PV inverters Sensors Storage management
Wireless		Sensors PV inverters	

A taxonomy of Smart Grid communication can be structured in different ways. The way chosen here is according to the geographical scale of communication and the connection endpoints. In addition, we distinguish between public/private and wired/wireless communication.

On geographical level we have have two primary scales:

- Wide area / Metropolitan area networks
- Local area networks

WAN/MAN are usually used to connect to endpoints which already have collected and/or aggregated data. Typical endpoints are, for example, RTUs (Remote Terminal Units) in primary and secondary substations and HEMS (Home Energy Management Systems), EV

Chargers, PMU (Phasor measurement units) and so forth. The communication media can be wireless as well as wired with different flavors. The choice of network technology depends on the requirements of the application and will be discussed later. As of today, a considerable part of the communication is done over private networks which are built for dedicated usage. This closed communication is increasingly becoming public. To secure it against cyber attacks, usually a VPN with encryption is used. There is a wide range of well-defined and accepted protocols which are used for the purpose.

For local area networks, all networks are private. The connection endpoints are here the data concentrators, like HEMS, as well the sensors and actuators. Actuators are not only, for example, blind controls, but also EV-charging systems. Even in a private network, encryption has become more and more advanced to meet cyber security requirements.

Many communication protocols are established in the market. Some of them are well defined and accepted like IEC61850 or KNX, but there are also many proprietary protocols, which makes it very difficult to mix endpoints from different vendors.

5.2.1.1 Requirements

This section summarizes the communication requirements of future Smart Grid networks.

Interoperability. The Smart Grid will be composed of multiple grid systems which have to share and exchange information. Thus, a common understanding of exchanged information and interfaces has to be established, even between equipment bought from multiple vendors.

Ability to upgrade/maintain. In order to cope with future needs and requirements the Smart Grid has to be flexible to incorporate evolving technologies. Development and adoption of standards could help to avoid customized efforts but to maximize utilization for all users.

Reliability (despite grid failure). Adding information technologies to the power-grid should improve the reliability of the grid, limit the extent of breakdowns, accelerate recovery from failures, and establish self-healing of nodes. Additionally, as the grid goes down the communication system has to remain active in order to take control of the grid and manage the recovery.

"Low" cost. Despite the merits and expectations of an evolving Smart Grid, its implementation and operation have still to be cost-efficient. While there is an understanding that today's energy prices are too low, it's unclear how consumers and commercial sector depending on power would react on significant price increases. Introductions of Smart Grid have to be incremental in order to balance costs, collect experience, minimize failures and develop business models, incentive models and education of customers correspondingly.

Delay bounds. Total delays of data in a Smart Grid's control must not exceed certain delay limits in order to the requirement of a reliably Smart Grid. Thus, power grid control may need its own dedicated private networks or should be prioritized on public networks.

Throughput needed. A Smart Grid's control commands must have a high rate of successful delivery across communication channels being used. Communication channels, public or private, have to be designed in order to provide enough bandwidth for transmitting also the maximum of control communication (e.g., in emergencies) successfully.

Error rate bounds. In the case of bandwidth requirements exceeding the network availability, error rates should be minimized in order not to exceed the delay for a successful grid control.

Clock/time distribution (PTP) (synchronicity). Clocks of all components involved in the control of the Smart Grid have to synchronized in order to allow for an efficient control in the distributed network.

Privacy. Privacy has to be protected of both the Smart Grid operators and its users connected by public networks. The operators' managing information has to be protected from the users as well as the individual user-specific characteristics of using the grid from the operators.

Security. The smart power grid will be controlled by an information communication system whose confidentiality, integrity and availability has to be protected. The security measures have to be applied against system failures, use control errors and external events.

5.2.1.2 Technologies

Communication network technologies are mature and well understood. For completeness, these are summarized in the table below.

Table 2 Communication network technologies.

	WAN/MAN	LAN (Private)
Wired	Fiber Optics PLC A number of proprietary technologies	Ethernet BACNet and Others Powerline PRP
Wireless	3G/4G Proprietary	WiFi Zigbee Bluetooth WiFi-Direct Many proprietary technologies

5.2.1.3 Research Opportunities

Near-term research goals (2 years out):
- Scalable IP multicast that works for wide area (short-to-mid-term)
- Scalable and interoperable communication paths (for smart meters)
- Reliable (multi-path), practical and cost-effective communication
- Guaranteed very low latencies in WAN/MAN for real-time control applications

Mid-term research goals (5 years out):
- Integration of legacy wired/wireless technologies into IP for both WAN and LAN
- Ability to upgrade/maintain communication infrastructure
- Authenticated sensor readings
- Making communication tamper proof
- Reliability of communication despite power network failures

Long-term research goals (10 years out):
- Securing devices despite physical access by attackers (booting, etc.)
- Going from communication to semantic SG
- Emergency back doors that are safe
- Localization and integration of pervasive sensors in secure way (long-term)

5.3 Smart Grid Control

Bert Claessens , Frank Eliassen, Christoph Goebel, Kai Heussen, S. Keshav, Hermann de Meer, Johanna Myrzik, and Mario Paolone

Summary

Control is pervasive in the existing grid and will play an even more critical role in the future Smart Grid. Control actions, the players taking these actions, their objectives, and the control mechanisms themselves are diverse, complex, and sometimes mutually conflicting. Control actions today span continents (such as with HVDC interconnects) and 12 orders of magnitude in terms of time-scales of control; from milliseconds to decades (if we may interpret planning as a form of control). They are taken by entities such as government bodies and market regulators, as well as by transmission system operators (TSOs) and distribution system operators (DSOs). The elements being controlled include equipment such as load tap changers and PV inverters but also some of the entities themselves (for example: the establishment of a grid code by a government is one way for them to control TSOs; a TSO requests demand response via an aggregator who is responsible for the actual unit control).

Not surprisingly, the objectives of control are also diverse, ranging from supply security and greenhouse gas mitigation to technical frequency and voltage regulation. These objectives are achieved by a number of mechanisms, including day-ahead and hourly markets, establishment of regulatory legislation, changing transformer taps, and topology reconfiguration through sectioning. One outcome of our work is a comprehensive analysis of the numerous control mechanisms in common use today, which can be found in the appendix.

Recommendations

Based on our analysis, we make a number of recommendations for research directions in the area of Smart Grid control. In the short term, we suggest the study of novel control policies for decentralized, policy-based control of voltage & frequency, and coming up with better models for demand-response capabilities of loads (i.e., characterize their flexibility). It would also be interesting to pin-point and eliminate inefficiencies in market design, to adapt to characteristics of changing energy resources and uncertainty.

In the short-to-medium term, we suggest a number of research areas. A few are discussed here, details of the rest can be found in the Appendix. We suggest research into better prediction of loads and stochastic generation, especially at short time-scales. We would also like to predict not just loads and generation but also characterize the uncertainty and variability in this forecast and integrating improved forecasts into state estimation, decision support and control systems. We suggest studying innovative bid types in market that take into account energy constraints as well as load flexibility. We also advocate research into optimal rules for storage operation and better state estimation. In the medium term, we believe that it is critical to allow SCADA to scale to much larger sets of inputs and processing frame-rate (a “SCADA on steroids”).

In the medium-to-long term, we believe that we need to study the supervision and design of dependable self-reconfiguring control systems that can act semi-autonomously on behalf of the system operators. We also believe that it is important to study how uncertainty can be quantified in market clearing to represent actual cost of uncertainty and promote the value of flexibility resources.

In the long term, we recommend studying the merging of control-based & data-driven control in distribution sub-systems. Another (more radical) idea would be to design and operate fully-local and autonomous micro-grids that can operate entirely independent from the grid. It may perhaps be possible to eliminate the grid altogether, if resource availability in each microgrid exceeds the benefits from grid connection. This may be the best way to allow renewables integration at low cost and without impact on system stability. Even with grid connection, such a solution offers higher availability than the present grid when faced with systemic failure.

Appendix

The following points briefly expand on the summary, adding some important details.

5.3.1 Geographic Scale

- Inter-continental interconnection
- Super grid (HVDC)
- Country level
- Regional
- Primary substation
- MV feeder
- LV grid
- Single building
- Single room

5.3.2 Time Scales

- Decades
- Investment horizon (1-2 years)
- Seasonal (6 months)
- Day ahead
- Intra-day (hourly)
- Tertiary (balancing)
- Secondary (60 s)
- Primary (0-10 s, droop and inertia)
- Protection (100 ms)

5.3.3 Players

- P0. International bodies (IEC, IEEE), equipment suppliers
- P1. Continent wide regulators (ENTSOE, FERC, ...)
- P2. SuperTSOs[1], market operators, electricity authorities (government)
- P3a. Gencos, balancing operators, aggregators
- P3b. Pure TSOs, DSOs
- P4. Retailers, aggregators, [micro-grid operators]
- P5. Industrial, commercial, and residential consumers, prosumers, aggregators
- P6. Individuals, plant operators

[1] By SuperTSO, we mean an entity that provides both transmission and market making services, such as those found in Germany. In contrast a Pure TSO provides only high voltage transmission.

5.3.4 Control Objectives

- O1. Supply security
- O2. Greenhouse gas mitigation
- O3. Energy affordability
- O4. Risk assessment and management
- O5. Minimizing COE / max. profit
- O6. Congestion management (transmission line planning)
- O7. Supply security – reserve contracts
- O8. Balance (day-ahead and faster)
- O9. Frequency stabilization
- O10. Voltage stabilization
- O11. Rotor angle stabilization
- O12. Protection
- O13. Intra-day portfolio balancing
- O14. Management of performance requirements for ancillary services

5.3.5 Control Elements

- E1. Power plants / generators
- E2. Transmission line switches and topology
- E3. Grid inverters
- E4. Reactive compensators
- E5. FACTS and universal power flow controllers
- E6. HVDC point-to-point
- E7. OLTC
- E8. Controllable load
- E9. DG elements
- E10. Inverters
- E11. Energy storage devices
- E12. Asynchronous generators
- E13. Reclosers/switches
- E14. Protections
- E15. Power conditioners

5.3.6 Mechanisms

A mechanism is characterized as: "PLAYER meets (CONTROL) OBJECTIVE by controlling ELEMENT/PLAYER using MECHANISM"

5.3.7 Detailed Recommendations

Time-scales: S = Short, M = Medium, L = Long

- Improve market design (S-L)
- Fully local micro-grid, total decoupling (L)
- Novel control policies for decentralized policy based control of voltage and frequency (S)
- Innovative bid types in markets: energy constraint, flexibility (S-M)
- Quantifying uncertainty as part of clearing strategy (M-L)
- Performance quantification and service requirements of new ancillary services (S-M)

Table 3 Smart Grid Control Mechanisms.

Player	Objective	Element/ Players	Mechanism
P0	O1, O2, O9, O10, O11, O12	E1, P0	Standards, making GHG measurable
P1, P2	O1-O8	P2, P3, P4, P5, P6	Grid code, FIT and subsidies, market regulation
P2	O4-O12	P3, E2-7, E10, E11	Bid types (product types), dispatching
P3a	O5, O13, O14	E8, E9	DR, droop and compound control, dispatch
P3b	O4, O9-12	E2-7, E9-11, P3a	Real-time monitoring through SCADA, droop/comp control, ancillary dispatch, reconfiguration, state estimation, prediction, dispatch aggregators
P4	O5	E8, E9	Monitoring, prediction, arbitrage
P5, P6	O1, O2, O5, O10	E9, E12	Turn on/off, set preference for DR, choose tariffs, choose aggregators, install PV, install energy storage

- High-frame rate optimal control (decisions per second) (S-M)
- Scalable SCADA systems – SCADA on steroids (S-M)
- Supervision of autonomous control systems (M-L)
- Better state estimation (real-time) (S-M)
- OPF and reconfiguration – optimal solution (S-M)
- Better prediction
- Ultra-short-term (S-M)
- Predict uncertainty
- Reconfigurable control system (M-L)
- Self-organization
- Model DR capabilities
- Unified modeling framework (S-M)
- Merging of control-based and data-driven control in DSS (Distributed Storage Systems)
- Controller conflict detection at all levels (S-M)
- Optimal rules for storage operation (S-M)
- Minimizing number of sensors in Smart Grid (S-M)
- Synchronicity of control structures and asynchronous control architectures (M-L)

5.4 Smart Commercial Buildings

Florian Allerding, Birger Becker, Christoph Doblander, Frank Eliassen, Manuel Götz, Hanno Hildmann, Hans-Arno Jacobsen, Peter Nogilik, Jiri Rojicek, Mischa Schmidt, Verena Tiefenbeck, Anwar Ul Haq, and Holger Ziekow

Today's Problems

Many of today's commercial buildings are equipped with some level of instrumentation and automation, addressable through more or less sophisticated installed Building Management

Systems (BMS). However, for advanced analytics and optimized control, information about the location of devices and systems is needed, too. While ontologies and standards like Industry Foundation Classes (IFC) exist in some buildings there is often no consolidated mapping of a building's structure, its systems and the devices. The process of creating and maintaining this Building Information model (BIM) is labor intensive.

A further problem in current building operation constitutes the often encountered contradiction of control actions, e.g., cooling and heating being active simultaneously. This is a symptom caused by a deeper rooted lack of orchestration of different building systems. Often individual energy systems are optimized for their purpose, but building control and management does not take complex interactions into account.

Also, over the lifetime of a building, there may be multiple retrofits, additions, upgrades of the building structure, its use and/or its instrumentation. Often equipment from multiple vendors is installed, giving rise to potential compatibility issues and proper documentation of installed systems is often found to be lacking. Very often this puts the BMS which integrates the different pieces of equipment into a powerful position where in essence the owner is in a BMS vendor lock-in situation.

Another challenge is the high degree of heterogeneity of commercial buildings; many of them are customized to individual functional or specific geographic requirements. In combination with a high degree of complexity of the supporting building infrastructure (e.g., HVAC or security system) and a variety of standards and communication protocols used by different vendors, solutions are difficult to transfer from one building to another. In addition to that, continuous training of the building operation staff is required to ensure that technical upgrades and new functionalities of the building are understood and not tampered with.

State-of-the-Art

There are available a number of commercial services offering energy audits, fault detection, diagnostics, and other information necessary to make energy-efficient and money-saving business decisions. Those services create value for building owners, but based on proprietary platforms and, as such, not open for further improvements by the community.

5.4.1 Automating Energy Audits, Fault Detection, and Diagnostics

Today, energy audits and diagnostics largely rely on experts that analyze the captured data with limited tool support. Inefficiencies and system faults are detected, e.g., through manual analysis of descriptive statistics and visualizations of sensor data. The rather high degree of human participation in this process limits the scalability of energy services.

5.4.2 Plug & Play Configuration

Setting up building sensors/actuators still requires a large degree of manual configuration. This includes the definition of communication endpoints as well as capturing the context and semantics of each sensor/actuator. An additional challenge is, that the configuration and context of sensors may change over time, and has to be reflected in the system.

5.4.3 Building Commodity Interactions

While building simulation models take interactions among different building systems into account, e.g., the additional warmth created by lighting when switched on, so that additional heating may not be required to reach comfort parameters, BMS currently do not consider this kind of interplay.

Research Challenges

5.4.4 Automating Energy Audits, Fault Detection, and Diagnostics

A challenge for future research is to increase the degree of automation in audit, fault detection, and diagnostic processes. New analytics methods need to be developed that automatically detect inefficiencies and faults as well as assist in defining corresponding actions. The prospect to increase the degree of automation can make corresponding energy services available at a larger scale and to more users. Energy experts may still be part of the process but may be able to server more installations though better analytics support. The automation of this process may benefit from better integrating BIM data with information about the building instrumentation.

5.4.5 Plug & Play Configuration

Future research should address solutions that support the process configuring sensors and the maintenance of this configuration. This goes beyond automatically establishing communication channels but includes mechanisms to capture as well as maintain the context and semantics of an installation. Ideally, adding/changing or removing a sensor in an application should be limited to executing the physical deployment and not require any further human interaction with the system. In this context also the location within the building of the new added sensor should be automatically identified.

5.4.6 Multi-tenancy

In commercial buildings it is common to find multiple tenants (organizations). BMS should support, e.g., via a virtualization concept, a separation of concerns and allow tenants to individually access and control systems of their individual concern, while resolving situations where tenants have conflicting interests causing inefficiencies. On a more individual level, individuals have different preferences, e.g., with respect to thermal comfort or lighting. As a result, to the extent that these preferences are compatible with other individuals' preferences, building occupants should have the possibility to adjust settings for sub-spaces (e.g., individual offices). Ideally, the operation system learns individual preferences and automatically balances local system states between conflicting individual preferences.

5.4.7 IT/Cyber Security

In any case, the integrity of current and future BMS must be ensured. It must be possible to identify a sensor/actor and make sure that this device is exactly the device that the system thinks it is. Furthermore encryption is must. Any physical access to sensors/actors must be detected by the system and not lead into unrecognized exchange of device with unpredictable system behavior.

5.4.8 Building Commodity Interactions

BMS need to be enhanced to optimize overall energy efficiency within the operational constraints by exploiting building commodity interaction effects. Moreover, buildings should be enabled to adapt their needs of the different energy forms to the utility requirements including power, gas, district heating etc. by exploiting these interactions.

5.4.9 Research Opportunities

Measuring energy demand on a plug-level requires today the usage of smart plugs. These are currently in the price-range of EUR 50-70. Even if there price will drop in the following years they will not scale economically. Hence, an alternative way to determine the energy use of individual devices needs to be found. Disaggregation approaches have come up recently promising to identify household appliances based on their characteristic fingerprint. Analyzing the aggregated overall consumption (e.g., at the Smart Meter) will allow the assignment of energy usage to each consumer. These methods need to become reliable and accurate.

5.4.10 Research Goals

Near-term research goals (2 years out): Anomaly detection today relies on a number of data points when many buildings are not well instrumented. To introduce methods utilizing virtual data points or working on less data points is a near term task. Also, enabling BMS to leverage building commodity interactions to optimize overall energy efficiency can and should be near-term research goal.

Mid-term research goals (5 years out): A medium-term research goal should be to enable the use of building commodity interactions to adapt building energy needs to utility DR requirements. Furthermore, the logic of current building systems usually are scoped for one tenant and hence can only optimize for this tenant. We therefore expect that future research will create system that optimize the building under considerations of all its inhabitants, which could be another mid-term research goal. For instance, the energy consumption will optimized under consideration of the building as a whole, and not, e.g., apartment by apartment. However, this optimization comes with challenges like mediating between conflicting goals and the increased complexity of the optimization problem. While supporting multi-tenancy, privacy and security are required to be maintained.

Long-term research goals (10 years out): In the long run (by 2025–2030), smart buildings will evolve to become semi-autonomous smart entities (as compared to today's rigid and inflexible shells which merely contain increasing numbers of smart elements) with varying degrees of autonomy. Such entities will be able to interact with the outside world in the same way an organism (comprised of a host of interconnected elements, all of which may have requirements and restrictions of their own and which are strongly dependent on each other and the organism) interacts with its environment. Abstractly speaking, the whole will become more than the sum of its parts. This may initially be restricted in the sense that the building manages and assigns resources between its parts, however the ability to do so in an independent manner and driven by objectives that are not necessarily shared with its internal elements will already constitute an improved level of "smartness" over the current state of affairs.

We envision buildings which effectively act as a layer between the different internal entities, resources and capabilities; and, in addition, between their internal entities and virtually all aspects of the outside world. The building-wide management and control of resources, ranging from parking slots to renewable energies, bandwidth, storage and supplies up to human know how and manpower will enable a dramatic increase in the efficiency of usage of these resources and as well as transparency, reliability, and accountability.

Shifting the jurisdiction over resources from the stakeholders to a smart semi-independent entity will enable the dynamic usage of resources, which today are still assigned permanently to one owner. This will both lower the cost for the resources incurred by the respective user, as well as increase the utilization thereof. As this evolution of buildings progresses across cityscapes, the ability to share resources and capacities between buildings and or their stakeholders will play an integral part in the evolution of our urban environments from relatively static institutions to truly smart cities.

5.5 Smart Residential Buildings

Florian Allerding, Birger Becker, Manuel Görtz, Jiri Rojicek, Hartmut Schmeck, Thorsten Staake, Kai Strunz, Verena Tiefenbeck, Anwar Ul Haq, and Andreas Veit

Problem Statement

Smart residential buildings share a number of properties with commercial buildings, but they also have distinct special characteristics. They refer to private living spaces in houses or apartment buildings, owned or rented by private persons, commonly referred to as smart homes. Investment decisions on infrastructure in residential buildings are quite often based on emotional or convenience arguments in contrast to being rational or business case driven. A typical motivation for investing into smart infrastructure may be the intention to demonstrate a “green” attitude. The users interacting with the smart residential building are the residents themselves and will thus usually be technically rather illiterate, hence needing out-of-the-box solutions which are mostly self-explaining. Nevertheless, in particular with respect to trustworthiness, installed services should provide interfaces for optional control on demand and they should inform transparently about their functionality and operational state. A private resident usually will not invest into completely integrated smart home solutions, but rather, over time, get a collection of verticals providing functionality for a special use case. The size of components in a residential building usually is much smaller than in a commercial scenario. The same is true for the expected lifetime, ranging from 3 to 5 years in a home to at least a decade in commercial buildings. Finally, a critical aspect with respect to the acceptability of smart home infrastructure is the aspect of security and privacy protection. In particular, security should be guaranteed by the systems while privacy aspects have to be under complete user control. This also needs guarantees from the verticals for compliance with privacy protection preferences. In simple words, a private resident wants to get perfect service support without being bothered by technical details, but having the option of complete control.

The objectives connected to smart home services depend very much on the resident’s individual attitude. Most users are interested to increase comfort by building automation, with less priority on cost-effectiveness. To maximize the benefit of smart homes for users and operators, the smart home verticals are expected to provide value-added services, e.g., the maximization of self-supply or an improvement of maintenance aspects. An essential objective of a private resident may be to achieve independence from public infrastructure.

Nevertheless, users will even be able to provide services such as demand response themselves, potentially for making money with it or just for the feeling of doing something good for society. In any case, the integration of smart home technology into a building must not compromise the whole system's reliability.

State-of-the-Art

Currently, there are several players on the market offering specialized smart home solutions, ranging from home automation over ambient assisted living to energy monitoring and control. Examples include Dropcam, Nest, Netatmo, or Plugwise. They are "verticals", only rarely capable to interact with each other. Commonly, smart phones or existing home gateways are used as their hub to stream data into the vendor's cloud and to interact with devices in the home. Some new players appear on the market delivering "smart home" platforms for (mostly wireless) connection of household components, sensors, and actuators. These platforms offer the potential for interaction between different devices and their data, respectively. Quite often, they use a single open or proprietary near-field communication protocol (such as ZigBee, Z-Wave or Homematic) for connecting the devices. This heterogeneity of communication protocols is an essential barrier for home users, as the individual devices (household components, sensors, and actuators) are usually not compatible with each other. However, there are a few commercial approaches to combining several technologies into an integrated approach (like EE-Bus, Qivicon, or Homee). In connection with smart metering, some utilities (e.g., EnBW, RWE, Vattenfall) are offering metering data analytics and recommendation services. On a research level, several management platforms or "operating systems" for smart homes are emerging. Some examples are Organic Smart Home (OSH), OGEMA, FPAI, Eclipse Smart Home, ABASG, or Open Energy Monitor, offering platforms for monitoring, management and control functions for energy scenarios. A widely missing feature is plug-and-play as known from computers today. One approach to realize this feature for a management platform uses ZeroConf. Furthermore, security and privacy aspects are often neglected which can generate additional resistance of the customer to buy into the systems. Since private users are widely sensitive to energy efficiency labeling (like A++), additional "smart home readiness" – labels might positively influence buying decisions.

Outlook

Over the next 15 years, substantial progress in both the adoption as well as the capabilities of smart homes is to be expected. Future applications for smart homes require two key enablers: A sufficient level of infrastructure/instrumentation and well-established standards. In the light of EU directive 2009/72/EC, we expect a substantial increase in the deployment of electricity (and gas) smart meters in most EU countries in the short- to mid-term. In the mid-term (around 2022), we expect different home automation providers to establish themselves. Rather than classical energy providers, likely candidates include telecom providers and entertainment companies as they already have physical presence within residential buildings. While these home automation systems will be able to integrate some parts of the smart infrastructure, we expect to still see a broad range of independent verticals for surveillance, HVAC, load monitoring as well as entertainment and content.

In the long run (until 2030), we expect two key developments for smart homes. Firstly, we see a strong integration and advances of the smart capabilities of residential buildings, and secondly, we see an increased independence of smart residential areas from external infrastructure, like the electricity transmission grid. In-home infrastructure will be created which is able to accommodate various elements like PV panels, batteries, thermal storage, or EVs. This includes monitoring, analytics, control and a coordinated interplay of all elements involved. Regarding smart capabilities, we expect that the technology enabling intelligent features of smart homes will be almost invisible to the user. A key aspect will be the ease of use for the residents. Smart home systems will be able to observe the users' feelings and preferences with non-intrusive sensors in a variety of devices such as phones and wearables. The invisible intelligence of the smart home will be able to autonomously understand the user's needs and adapt its services accordingly. Although the residents will not be required to actively tell the home automation about their preferences, they will be able to overwrite autonomous control, if necessary. In fact, these systems will not only be able to adapt to the needs of single users, but also to understand and negotiate between the needs of groups of people in a room. Further, the home automation system will be able to automatically detect and integrate new devices that enter the building and know or learn their best utilization. Regarding the increased independence from external infrastructure, energy supply will decrease in importance. This is due to much more efficient building insulation and usage of distance heat from, say, factories and compute centers. Further, there will be significantly increased energy "harvesting" from within residential areas.

This increased independence of smart residential areas with a focus on decentralized systems could make high voltage power transmission obsolete, which would be a strongly disruptive development. Nevertheless, while this outlook refers to highly industrialized countries, heavily increasing energy demand in other parts of the world might lead to different scenarios, where the need for smart residential building automation is even stronger, in particular with a focus on energy conservation and demand optimization.

Research Challenges

The gap between state of the art and prospects of future smart home technologies translates into several research challenges at the intersection between ICT, energy systems, human behavior, and the policy framework. These challenges include the collection of data from the user and its environment as well as robust mechanisms to translate the data into stimulating and insightful information. Furthermore, the data has to serve as actionable input for automatic control systems in order to meet the highly individual user preferences including comfort, safety, and sustainability. Working towards these objectives requires considerable progress in the following fields:

To collect the data, a multitude of sensors and other data sources need to cooperate. These sensors need to be non-intrusive, lightweight, and energy-autarkic to neither burden the user nor cause high costs for installation and operation. A particular challenge is that the systems will often need to become part of existing infrastructures, and have to adhere to a multitude of domain-specific standards and characteristics.

Domain-specific machine learning techniques will be the second corner stone of smart buildings. The algorithms need to be capable of dealing with data time series of varying depth and quality. One major concern is the large variety of data characteristics influencing the accuracy of control and predictions.

Human Machine Interaction is another vital aspect. Looking at the state of the art of both, current products and research prototypes, it becomes apparent that the energy informatics community has to still go a long way to build systems that are easy to use for the general public and at the same time effectively motivate a desired behavior. In this domain, a closer cooperation with psychologists and behavioral economists might help to establish the tools and methods that trigger behavioral change.

Since integration of a large variety of verticals is essential, a core challenge relates to the design and dissemination of an adequate infrastructure operating system. It will have to provide typical services known from computer operating systems like resource scheduling and allocation, software updating, resilience and access control to name a few.

The multitude of highly personal data collected in smart homes imposes a challenge for data management and privacy. A framework is needed that provides the analysis of sensor data in a way that preserves users' privacy and maintains security. Similarly, the smart infrastructure needs to be well-protected against attempts to compromise the user's safety. This requires robust methods of authentication for users who want to access the building system as well as control processes for providers of software services for the building's infrastructure.

Life cycle analysis with respect to resource efficiency is another important aspect. It touches all fields from producing, shipping, operating, and recycling the growing number of smart devices.

Last but not least, whether residential building automation systems become a success will be largely determined by the underlying business cases and their attractiveness for service providers. Since much depends on the availability of consumption data (e.g., from smart metering), policy makers will have to find a delicate balance between limiting the use of data and privacy protection. The energy informatics community can contribute to these considerations by providing methods for effective data usage control.

5.6 Smart Transportation

Longbo Huang, Randy Katz, Victor del Razo, and Hans-Peter Schwefel

Breakout Focus

This breakout session focused on the research issues in Smart Transportation. In addition to commercial and residential buildings, transportation represents the largest source of energy consumption in industrialized societies (approximately 40%) and the largest source of greenhouse gas emissions, with critical implications for the earth's climate. The breakout considered both personal vehicles (cars) and public transportation (including bikes). There are many incentives for smarter transportation systems, including reduced energy consumption and GHG production, improved air quality with implications for health and quality of life, reduced societal externalities associated with the very high cost of car ownership, and the promise of reducing accidents and road deaths.

State-of-the-Art

The current state-of-the-art is largely dedicated user applications for assessing and reporting traffic congestion and estimated time of arrival to their cell phone users. Accuracy and timeliness can be an issue, resulting in the driver finding herself in congested traffic before it can be avoided. Such applications maintain proprietary data silos, with no sharing of trip data even though this is in principle owned by the users. Rerouting recommendations are not coordinated among such applications, which can lead to failure to improve congestion through rerouting suggestions. Generally the tools to help users out of private vehicles and on to public transport are limited in their effectiveness and usability.

Potential Game Changers

There are several potential game changers that are likely to challenge the current state of transportation systems. The first is the more IT-centered entrants in the vehicle sector, such as Tesla, Google, and Apple. The second is the rise of the so-called Sharing Economy, characterized by such firms as BikeShare, ZipCar, Uber, and Lyft. The third is the pervasive ability to track the location of virtually every vehicle on the road, whether it is publically available such as for public transit, or implicitly collected by such end user applications as Google Maps and Waze. A finally is disruption is the electrification of the vehicle fleet and the implications this has for the Grid and ready access to charging services.

Computer Science Challenges

The overarching computer science challenges in transportation can be stated as follows. The goal is to turn data about trips and transportation usage, collected from many sources and across many time scales, into actionable information for infrastructure and vehicle operators as well as passengers. The time scales vary from seconds, for safety and accident avoidance, to minutes to days for route options, including behavioral and economic incentives for alternative routing, load balancing, and road congestion avoidance, to months or years, for infrastructure planning and provisioning (e.g., charging stations, bike stations, bus routing, etc.). This will require a distributed and decentralized operational architecture to collect data, process it, and infer and decide at scale across a region and at the appropriate time scale. Such an architecture must be sensitive to concerns of information ownership, relevant business models, and privacy/security considerations.

Example Research Challenges

One identified challenge was EV range extension via route planning and charging station reservation, with capacity planning to inform charging station placement. Another was shared transportation resource planning and placement via demand and trip awareness (e.g., how many bikes and where to place them). This raises the important question of who owns mobility data and what is the business model for how it is collected and used. Privacy issues must be understood in such circumstances. A final example challenge was the definition and implementation of vehicle-to-vehicle and vehicle-to-infrastructure communications systems, to enable safety considerations, accident avoidance, and which span vehicles and railways, across dedicated or shared infrastructures.

Research Opportunities

The near-term research opportunity is to explore overlay architectures that allow the combination of multiple traffic sources for more accurate and timely congestion detection.

The medium-term research opportunities are to investigate how new vehicle sharing models impact the transportation system. Open for investigation is extensions of the overlay architecture to allow for intelligent rerouting with load balancing of traffic. Within such an architecture, introduce incentives to change operator and/or passenger behavior, such as migrate towards using higher density, more energy efficient transit modes such as public transit or shift modify travel times to avoid congestion. Large user studies should be undertaken. Finally, the study of the effect of new technologies, such as autonomous vehicles like cars or UAVs, on logistics systems should be undertaken.

The long-term research opportunities include investigations of the implications of high-penetration electrification of the vehicle fleet, such as, charging station infrastructure provisioning and placement. Also the value of vehicle trip data for managing the Smart Grid for EV charging should be investigated. Also ripe for inquiry is the impact of self-driving vehicles on the overall transportation system, from the perspective of fundamental changes to the existing car ownership model. These include effects migration to a largely shared vehicle fleet, with implications for road occupancy and parking, and the avoidance of new road and parking construction. Other implications to be studied include accident avoidance and road safety; ride sharing, trip scheduling, and road congestion avoidance; dynamic locating of shared vehicles to where they are likely to be needed within a city; and so on. The interface to Smart Cities should also be explored, including incorporating vehicles as full participants in the Internet of Things.

5.7 Data Crosscut

Manuel Görtz, Randy Katz, Srinivasan Keshav, Jiri Rojicek, Hans-Peter Schwefel, and Anwar Ul Haq

Breakout Focus

This breakout session focused on data issues within and across Smart Infrastructures. A particular focus of discussion was the Smart Grid, and the exchange of data between (aggregated) loads and entities that supply that load, across time scales and geographic regions. Discussion included the processes for collection, cleaning, processing, and curation of data, and how it forms inputs to decision making. Furthermore, there was an awareness that other sources, such as weather, events (e.g., World Cup), transportation and mobility, social networks, and so on, provide potentially useful indicators of human activity that influence energy demand and ultimately could be an input to the energy system.

State-of-the-Art

The current state-of-the-art was assessed starting with smart metering of end loads. While in some jurisdictions considerable effort has been dedicated to deploying smart meters, generally smart meter data is unused and not particularly useful. This is due in part to consumer resistance to dynamic pricing (the original motivation for smart meters) and a general mistrust of the utility operator. Generally, transmission and distribution operators have good quality forecasting tools to predict demand and provision the grid for the current state of grid architecture. Fine time grain data is not needed, since agile control has not yet been deployed. Since consumer level load data is not actually used, there is in fact no real privacy concerns. At the level of large-scale aggregation, operators are deploying better measurement infrastructure, in the form of PMUs (Phase Management Units), to better manage their networks in the face of increasingly complex load and supply dynamics, and grid interconnections. Nevertheless, these devices are expensive and generally limited to the transmission grid.

Potential Game Changers

There are several potential game changers that are likely to challenge the current state of data collection, processing, and use in the Grid. First is the emerging disaggregation of the Energy Network. The Energy Network is evolving from an integrated end-to-end system, to one that features looser connections between TSOs and DSOs, to one that may eventually become made up of semi-independent microgrids. The implication is the end of the "Law of Large Numbers" that allows statistical multiplexing to hide time variations in individual generation supplies and loads. The second game changer is greater penetration of edge PV generation and EV charging, yielding even higher time variation in supplies and loads. This will drive the requirement for information about instantaneous energy supply and consumption on finer time scales and smaller regions. The third trend is migration of intelligence towards the edge, with a greater prevalence of microgrids (e.g., office parks, campuses, industrial parks, shopping centers, etc.). We foresee a future Grid in which large aggregates of supplies and loads are replaced with smaller aggregates, managed with real-time intelligence for more localized control. We believe that natural unit of aggregation is likely to be at the level low voltage distribution (e.g., secondary transformers) supplying loads to something on the order of 10–50 homes.

Research Opportunities

The near-term research opportunities are to lower measurement and analysis costs by developing more rugged technologies for monitoring and performance analysis of the Grid. There is also an immediate need to improve customer awareness of energy usage, provide a better user experience, and deliver an improved perceived value of collecting edge energy usage data. This will require better tools and visualizations for consumers and other edge customers (e.g., building and campus facilities managers) to understand their detailed energy usage data. Finally, there is an opportunity to collect existing and to create new datasets of energy usage that can be made available to the research community for analyses at larger scales and over greater diversity.

The medium-term research opportunities are to use these new tools and data sets to better understand individual load profiles, e.g., to level of individual appliances usage, from house level load data (NILM: non-intrusive load monitoring). The data architecture that links the monitoring capability at the secondary substation level to the home load should be designed, prototyped, and evaluated in the context of dynamic Grid control. This architecture must be developed in a way that is sensitive to privacy issues, in part by using the appropriate level of aggregation and data partitioning to avoid tracking of individual behavior. A data processing architecture needs to be developed for characterizing and classifying of loads (clustering), tracking of data transformations and its long term archiving (curation), data placement and storage (collection), processing, and dissemination (up-sharing of aggregated and sampled data). Larger scale building and home data sets, including metadata, should be collected and made available for further study. Further analysis of the effects of human activity indicators on aggregated energy loads and microgrid coupled supply and load behavior should be undertaken, and control architectures and algorithms developed.

The long-term research opportunities is to design and demonstrate effective data-informed control of highly dynamic disaggregated loads and generation assets in a disaggregated Grid environment while understanding how automated exchange of data exchange across societal infrastructures can lead to better, more agile control algorithms.

5.8 Design Patterns and Paradigms for Smart Infrastructure

Frank Eliassen, Kai Heussen, Hanno Hildmann, Peter Noglik, Jose Rivera, and Hartmut Schmeck

Breakout Focus

The group discussed design patterns for the areas of Smart Buildings and Smart Grids. Buildings and electricity networks are primarily a built physical infrastructure with historically little embedded ICT, functionality of conventionally engineered buildings and electricity infrastructure has been conceived the same (obviously) inflexible structures as the physical structures that enable them: i.e., shelter, heat, or reliable electricity supply. The purpose of a deeper embedding and higher sophistication of software technology within these infrastructures would be to offer enhanced functionality and reconfigurability that has been achieved with software defined systems in other domains.

Design patterns inform high level choices that have to be made at an early stage in a system development process. A stereotypical choice is whether to take a centralized versus a decentralized approach or whether to use hierarchical structures – and if so, whether to do so strictly or loosely. A central management structure may be replaced by (or at least supplemented for some aspects) by self-* approaches so as to push some of the management overhead closer to the device layer. Monolithic architectures are receiving increasing competition from modular systems. These are of interest as composability allows for wider application of systems due to enhanced customization and facilitates the compartmentalization of problems.

The area of application is vast, but two aspects take a central role: the definition of suitable software platforms and the re-formulation of system control architecture. Engineering processes that consider jointly physical structure and interdependencies, control structures and software systems have to be developed. Engineering methods and tools are needed to support these considerations for off-line design work, run-time updates, and, particularly in context of electricity grids, also for re-engineering of the operational system (running systems).

Regarding the evaluation criteria, performance requirements and time constraints are as relevant as considerations for feasibility, resilience, reliability and stability, robustness and, last but not least, simplicity. Criteria that seem to have a high impact on the acceptance of a system range from disruptability (i.e., whether a platform or architecture can be amended while in use or whether parts of can be taken down without bringing the system to a halt), over versatility and vulnerability to trustworthiness. In addition, explainability and the identification and/or assignment of liability are of high importance to ensure wide acceptance of an approach.

A larger number of fields contributes to the state-of-the-art (see below), but common properties that seem to span across most of these fields are reconfigurability, adaptiveness and robustness, the ability to predict models and the ability to implement parts of the system in a distributed manner.

State-of-the-Art

- Control theory
- Trade-off matrix for different approaches

- Self *
- Autonomic computing
- Organic computing
- Multi-Agent Systems (MAS)
- Machine learning (learning systems)
- Emergence and emergent effects of local actions and interactions
- Power system control
- Software requirements engineering
- Cyber-physical simulations

Research Challenges and Opportunities

One of the main challenges we have identified is the need to consolidate in a coherent picture of the vast amount of knowledge available on this topic. It is often the case that several disciplines have developed the subject independently of each other, and many times unaware of the results achieved in other disciplines. For instance, we need to couple control engineering requirements with software/systems engineering principles, we need requirements engineering. Merging mathematical control theory with the self* is another important component, failsafe engineering. For this we will need to develop engineering process for self* systems and create the theoretical basis to understand them. The consolidation of this knowledge represents a key research challenge and will have a major impact in the design of future energy systems. We consider that a common benchmark platform that addresses the needs of different disciplines provides an opportunity to consolidate the different research efforts.

Several aspects like resilience, reliability versatility, and liability are also a major challenges and will play a key role in acceptance of these systems. There is a multitude of options we need to develop, e.g., we need to design self-reconfigurable cyber-physical architectures, we need architectures for real-time distributed MAS, we need to be able to do runtime deployment and, last but not least, we need to factor in the aspect of IT security. We consider that there will be a multitude of options, such that the community of researchers and practitioners will need a trade-off matrix of the different approaches.

Near-term research goals (2 years out):

- Trade-off matrix for different approaches
- Combine control engineering and software/systems engineering
- Engineering process for self* systems
- IT-security
- Runtime deployment
- Requirements engineering

Mid- and long-term research goals (5–10 years out):

- Merge self* with mathematical control theory
- Failsafe engineering
- Design rules

5.9 Human Machine Interaction in Energy Informatics

Birger Becker, Christoph Doblander, Johanna Myrzik, Thorsten Staake, and Verena Tiefenbeck

Problem Statement

Human Computer and Human Machine Interaction (HMI) concepts without doubt account for a large share of the success of many ICT products and services. Based on research results and learnings from practical deployments, industry has come a long way to better reflect the needs of the individuals using their systems. However, HMI in the energy domain faces specific challenges that are rooted in particular in the dynamic and multifaceted constraints and requirements, including the need to deal with time-critical changes on the supply side, the large number of energy-consuming applications and actions, and the often difficult to predict and highly situation-specific user intentions. Both, balancing energy supply and demand as well as energy efficiency increasingly has to follow supply without putting a too large burden on the end-user. Hence the degree of automation and prediction has to increase while preserving the users' freedom to intervene and overrule control systems.

Additional, interrelated challenges of HMI in the energy domain include:

- The relatively low average level of consumer interest in energy topics. In fact, energy itself is often neither visible nor of primary concern for the user since energy related expenses represent only a small share of a consumer's wallet: Moreover, energy conservation might develop into a pronounced social norm, but today's limited visibility of related actions limits the effort the user is willing to invest.
- Unclear benefits of smart energy system from many user's perspective. This is mainly the case due to the high level of comfort and reliability of the existing systems as well as due to the complexity of relationships of cause, action, and outcome of changes in energy systems.
- Short lived interest in energy dashboards. Closely related to the aforementioned issues, motivational "cues" are needed achieve a sustained system usage. Such "cues" are often not present in today's, engineering-oriented designs.
- Limited trust in utility companies. Systems (and the concepts behind marketing them) face the additional challenge that users do often not understand the motives of their energy providers behind offering energy efficiency products. This leads to distrust, which needs to be mitigated.
- Little tolerance for wrong system assumptions (e.g., cold shower water in the morning). Related to the reliable but inefficient existing system, user expectations are high. In most cases, developers will simply have to work to meet the requirements.
- Last but not least, user preference are inherently dynamic and situation specific (e.g., temperature preferences are not stable over time). Thus, prediction algorithms must be very accurate, or the system must offer a convenient way to allow for adjustments by the user.

State-of-the-Art

HMI concepts are used in different domains such as safety, security, and control. Energy-specific HMI concepts are needed to provide energy management systems with interfaces for user interaction. Systems that are available on the market are often domain- or application-specific, only vertically integrated (e.g., from a temperature and a dedicated motion sensor down to the heating control), and cannot be integrated with other systems. An example for an innovative system could be the NEST thermostat: The HMI has an innovative design, it is structured very well, and it is easy to use even for non IT-affine users. However, its "self-learning" features currently receive much criticism.

HMI approaches on a prototype level are able to provide flexible and configurable visualization and the capability for "complexity only on demand". This means that users can choose between individual levels of detail for visualization. Additionally, the status of self-adaptable systems will become visible for the user through the HMI.

In current research approaches for energy specific HMI, three different types of interaction are often distinguished: (i) visualization, (ii) parametrization, and (iii) configuration. In a first step, the state of the connected system will be visualized by current HMI solutions on a different degree of detail, so that the user is able to get information about the operation of single appliances or even of the overall system. Selective visualization may cause a change of user behavior. Additionally, parametrization of HMI allows the user to interact with the system during its operation concerning the individual demand. In this way, the individual parameterization enables the HMI to communicate personal preferences (e.g., the degree of freedom regarding the on-times of a washing machine) to the system. Configuration is the third interactive step of HMI, which allows the initial adaptation of the system to the real environment.

Vision

In the medium term, we expect an increased interest in energy management system, in particular among prosumers who are able to put the information in context with the energy production of their own systems. The increased interest is fueled by running out feed-in subsidies from the government, shifting goals towards self-supply. This requires HMI solutions which adapts to the behavior of the user. For this purpose, high reliability of the predictions and understanding of the user preferences are needed.

Information from these systems can be embedded into object-specific displays or mobile devices e.g., smart watches or phones. However, many other application compete for user attention, and particular attention has to be paid not to overwhelm the user. Through meaningful information display, the complexity can be decreased.

Motivated by the increasing share of intermittent generation, the time criticality and flexibility requires highly adaptive systems. Most of the decisions should be done by the system automatically, however the user should be always able to intervene. Cross domain interpretation of sensor data can be brought into future HMI applications. In general, HMI has the potential to increase trust in complex automation systems.

Research Challenges

The vision formulated before translates into a number of challenges for HMI research. This is especially true since several specific characteristics of energy supply and demand need to be considered: Energy – despite its enormous value for our society – is relatively cheap given the value it provides. It is a low-involvement good, with consumers not per se requiring energy but the services it enables (e.g., heat, light, telecommunication, etc.). At the same time, energy use is spread over a very large number of activities, and many devices need to be activated long before they can provide their service. These aspects add to the supply-side challenges and make the control – and with it the HMI – difficult. Among the many research challenges in HMI, the following bear a special reference to energy-related aspects:

- A large number of energy-consuming devices and activities will require many sensors and other data sources to arrive at a complete picture of energy demand, the state of the environment, and the user's objectives. This in turn translates into the need to effectively deploy and maintain many sensors, and to retrieve and combine the data from multiple sources. Many of the data sources will serve multiple purposes – they may been originally installed to increase comfort of safety, not energy efficiency – and the multi-faceted use will add to the complexity of their integration in energy systems. Energy informatics is thus probably one of the most advanced application domains for interoperability concepts.
- Raw data on energy supply, demand, the environment, and last but not least the user needs to be processed in order to reveal their underlying patterns. The insights from machine learning which range from electricity prices to the mood of the inhabitant are often required to trigger a target behavior (i.e., enable behavioral interventions) and render possible adaptive control systems that do not require any user interaction. Predictions in the energy domain are especially challenging as influencing factors are numerous, dynamic, and related to many application-specific characteristics. Yet predictions are important for many processes with high latencies. Even problems that may at first sight appear to be not hard – such as, for example, predicting the time of the following day an apartment is empty – are indeed very challenging, yet solving them could help to considerably conserve energy on heating.
- Energy systems often need to consider a large number of constraints that are highly user specific (internal: desired temperature, time, mood; external: weather, prices, etc.). Yet the expectations regarding the level of comfort are high, so states achieved by "smart systems" that are perceived as further away from the optimum than the outcome of conventional techniques (that are often comfortable yet energy intense) are hardly accepted by consumers. Optimization problems are hard in many energy applications, with extensive future work needed to arrive at suitable approaches for both the building and the transportation sector.
- A challenge for HMI that especially holds in the energy context is that complex information must be boiled down to a few easy-to-understand key figures in order to make the interaction between user and system feasible. Other than in health or nutrition, energy information is less tangible for the user. Nevertheless, complex or hidden relationships between behavior, energy use, and the consequences thereof need to be conveyed to the vast majority of non-energy literates including those who are not interested in becoming energy experts.Yet, a complete picture must be available upon demand (e.g., in case of failure), and the visualization interfaces must thus be configurable. Covering the balance between simplicity and in-depth information poses new challenges to the interface design.
- In order to trigger a target behavior regarding a non-visible, low-cost, and low-involvement good requires a very solid understanding of the behavioral concepts underlying human behavior. Unlike in health, fitness, or ostentatious consumption where the (perceived) benefits mighty be immediately felt, motivation to conserve energy is often smaller. An important research challenge certainly is to further develop interventions from time-invariant problems to tackle time-variant challenges (e.g., load shifting). Both visualization and interventions are more difficult to realize since many actions are time-dependent and inherently dynamic: There might be more than enough electricity available one day at noon and a shortage the next day at the same time. Timing is important. Interruptions need to be placed (of an automated service or a user action), without being annoying.

Moreover, the use of concepts from psychology and behavioral economy (e.g., social norms, goals, competition, etc.) should be given due consideration when working on HMI. Relying on rational choice models and monetary incentives – as mostly done today – is simply not sufficient.

The ultimate challenge probably is to make HMI disappear while making the residential building a place that – without consuming resources of its inhabitants – balances the requirements of all stakeholders. It will be a long and interesting way to work towards this goal.

5.10 Smart Cities

Kai Heussen, Hanno Hildmann, Longbo Huang, Milan Milenkovic, Johanna Myrzik, Jose Rivera, Misha Schmidt, Hans-Peter Schwefel, Joachim Sokol, and Randy Katz

Breakout Focus

Motto: "The battle for sustainability will be won or lost in (smart) cities."

Rather than attempting to (re)define the term "smart city", the team opted to outline its salient characteristic as an informed, data-driven, management of city and its operation for the benefit of citizens, and in accordance with their expressed consensus/preferences. The informed part comes from a combination of technologies, predominantly ICT and IOT based, including new and existing sensors, data crowd sourcing via social networks (existing and purpose built) for collecting citizen inputs in terms of preferences and observations, and opening of data from existing legacy city systems for integrated services. The vision is to create a city platform that integrates all of this information into a holistic view and allows creation of new applications and services to improve quality of life in the city and to streamline and increase efficiency of its operations. Over time, the city platform is expected to evolve to real-time sensing and reaction to events, to perform optimizations in accordance with policies derived from citizen preferences and to alleviate emergency situations when they occur. The city platform itself is supposed to provide mechanisms for realization of desired operational objectives, rather than assuming or imposing specific policies for doing so. Objectives and concerns of a city are generally driven by the desire to improve its attractiveness and quality of life, including: implementation of suitable infrastructure, access to affordable healthcare, good educational resources, transportation, supply of food and goods, social equity, mobility offerings, safety etc. The smart-city platform is expected to aid in achieving many of these objectives by providing metrics and data-driven basis for real-time management in accordance with objectives and priorities established by its citizens and government.

With regard to the city's inhabitants' objectives it is going to be a multitude of quality-of-life enhancing aspects rather than a single "killer app" that fuels the continued support and investment in the smart city platforms. The process is likely to be a long-term (decade(s) rather than years) process that will move forward with great momentum.

Today's Problems

The problem space of today of cities is different in different regions and vary from city to city. Nevertheless, major trends and drivers can be identified, which create pressure for actions and solutions:

- Population growth
- Increasing urbanization
- Increasing life expectancy
- Aging society
- Water scarcity
- Traffic congestion
- Pollution (air, soil, water, noise, waste)
- Energy supply
- Global warming
- Social divide
- Lack of funds
- Increased safety and security constraints

The list is not comprehensive, but it illustrates the complexity, magnitude, and dependencies involved in the underlying economic and societal pressures that challenge the rapid transition and change of city KPIs.

State-of-the-Art

The state-of-the-art is characterized by disjointed legacy management systems that will require an interoperability layer and addition of sensing fabric to evolve into a city platform that we envision. While there is a flurry of activity in smart city pilots and engagements worldwide, there are no major deployments that would validate the benefits of a smart city platform that we envision or clarify the business case.

- Disparate, uncoordinated systems for administration and management
- Individual siloed systems and administration, e.g., energy, buildings, water, transportation, energy
- Legacy infrastructure
- Many smart city pilot blueprints for implementation and testing possibilities, but difficulties in converting to real impact on cities (that a deployment may have)
- Business case still unclear: new city funding or funding from existing budget categories (do something that is being done better/faster, more efficiently)

Innovation Areas

- Energy efficiency/conservation in buildings: commercial and residential
- Water usage/conservation
- Air quality monitoring
- Transportation: public, multimodal travel, commuting
- Shared transportation: bicycles, electric cars, short-term sharing cars
- Smart parking
- Smart lighting
- Crowd control, including prediction of human (group) behavior
- Security, crime (prevention)

- Emergency (response)
- Maintenance, repair (potholes, water breaks)
- Citizen participation: communicate with government and each other, crowdsourcing of relevant information
- (Business) attractiveness

Research Challenges, Barriers and Drivers

- IoT/ICT infrastructure cost and complexity
- Administrative and organizational "silos" in city administration
- Privacy concerns
- Security
- Digital divide
- Ownership of (existing) infrastructure
- Existing (age of) infrastructure
- Local culture (affinity: sharing, computers)
- Level of education
- Social and economic divide
- Demographics, age distribution
- Legislation
- Communications – mobile/cellular data communication for sensors
- Data ownership

Research Opportunities

Near-term research goals (2 years out):

- Concepts for integrated or federated (evolution of) smart city computing infrastructure/-platform
- Standards: data and meta-data interoperable formats
- Common set of (generic) use cases
- Cross-silo integration at service level for (collaboration) infrastructure
- Big data: storage and analytics, access to multi-domain data sources
- Open standards for IoT sensing
- Simulations / planning studies – demonstrate impact of changes (to be done in advance before engagement, multi-dimensional)

Mid-term research goals (5 years out):

- Demonstration of smart city platform and service-level interoperability
- Standards: data and meta-data interoperable formats understood, tested and (widely?) deployed
- Linkage between ICT and energy/water/transport sensing infrastructure
- Big data: holistic data processing (including social media and consumer) and big-data driven control and operator support
- (Common) reference architecture for end-to-end system (including interoperable data and metadata formats)
- Real-time stream processing vs. batch processing
- Distributed vs. centralized processing
- Legislation and security, privacy

- Energy footprint for IoT and data
- “Smart city emergency/disaster response training center”

Long-term research goals (10 years out):

- Smart-city platform that supports full interoperability and (third party) cross-domain application and service deployment
- Big data based real-time action/control (closed loop, policy driven)
- (Common) reference architecture for end-to-end system with distributed data processing, storage, and analytics

5.11 Infrastructure Operating System, Application Platforms, Stakeholder Interoperation, and Plug & Play Resource Management

Florian Allerding, Florian Michahelles, Milan Milenkovic, Victor del Razo, Joachim Sokol, and Holger Ziekow

Breakout Focus

The group interpreted the scope of this breakout session as an investigation of the benefits of an integrated platform that would support development of cross-domain services and applications. Such an integrated platform would combine Smart Grid and Smart Cities (with smart building and citizenry involved through social networking), to allow global optimizations and benefits. It was group consensus that integrated platforms provide sufficient benefit and are likely to be developed, despite the complexity inherent in such complex multidisciplinary endeavors. A first-order simplifying approximation is to achieve integration at the level of common application (data retrieval) APIs with interoperable data and meta-data formats. The second-level, deeper integration would include a modular implementation based on a common reference architecture, but that is a longer-term research challenge. In this discussion we focused on the data-plane aspect of the platform (data formats, flows, and application APIs). We recognized that a deeply integrated platform would also need to include a control plane definition (configuration, management, security) but decided to defer this discussion, given that data-plane integration is complex enough and it will likely come first.

Motivation

Current and expected advances in ICT have triggered the smartification of basic and commodity services, opening a door of opportunities and presenting us with a number of challenges. Traditional facility management evolves towards smart building management. The increase in renewable energy generation require a more actively managed grid, particularly at the distribution level. Cities begin a transformation towards integrated services. Smart transportation and smart infrastructure concepts emerge as means of providing users with added value.

The need for interaction between services, equipment, data sources etc. becomes a must. The level of complexity of such level of integration can only be solved by the active participation of all stakeholders. An ecosystem that facilitates interaction between these stakeholders and allows for competition and entry to new stakeholders becomes a requirement. ICT paradigms and technologies that would enable integration and service composition, like Service Oriented Architecture (SOA) already exist. The question is which of, and whether, these technologies meet the requirements of the future smart services.

This report aims at identifying the domains where potential smartification would benefit from massive integration. We take a look into the different stakeholders, potential new comers and their requirements and potential interactions. Finally, we discuss challenges and opportunities towards enabling the level of integration that we believe is required.

State-of-the-Art

In current practice, cross-domain interactions are rare and difficult since they require expensive custom interfacing between closed systems. As a first step towards development of an interoperable platform and data formats, our team started by identifying key stakeholders and domains and mapping possible useful interactions between them. We believe that the table and matrix below, albeit incomplete at present, provide a useful framework and structure for reasoning about such interactions. The table below summarizes the domains where we see a potential evolution towards smart* and horizontal integration as well as the identified stakeholders for energy and buildings domain:

Domains:
- Building
- Power and energy
- Cities
- Transportation
- Industrial consumption/production

Stakeholders:
- City management
- Utilities
- Operators
- Industrial customers
- Commercial customers
- Residential customers
- Equipment providers
- Service providers

Most of these domains are characterized by vertically integrated businesses. Resulting in limited interoperability and cooperation. Additionally stakeholders tend to intensively protect their technology, intellectual property and data. Therefore openness to integration is commonly faces resistance. Each stakeholder has different set of requirements in terms of what they need from others and what they are willing to share. These requirements are an important building block for any interaction. Existing sensor and devices involve relatively high integration efforts and its operation and management is not trivial. A number of standardization, open interfaces and common reference architecture efforts have already been started (e.g., open ADR, open SCADA, SGAM).

Challenges and Opportunities

Given the interest of the stakeholders of protecting their data and intellectual property, it is important to understand and define strategies for evaluating what can be shared, the potential benefits, and the methodological and technological alternatives for doing so. We don't believe that such systems will evolve to fully open systems. Rather, we see them interacting on a common interfaces and communication rules. This means that the problem is more of a system composition than a system design one. The group identified the following challenges:

- Interoperable platform, but not necessarily public data: privacy, security, intellectual property subject to business arrangements between stakeholders
- Definition of data and meta-data formats, interoperable across domains
- Definition of basic service-level APIs for data and meta-data retrieval, subject to security and privacy constraints
- Understanding security and privacy risks, level of required functionality
- Design of a common reference architecture
- Separation of control (security, management) and data planes (data and meta-data)
- Strategies and technologies for device and sensor management and configuration
- Plug & play, sensors and apps
- In the particular case of Smart Grid devices, availability of ICT resources and devices despite power grid failure
- Matching legacy and evolution
- Matching and tracking requirements of heterogeneous target groups (also in terms of SLA and KPIs)
- Different expectations by stakeholders in terms of product/technology life-cycle

While the team did not have enough time to characterize research topics into short-, medium- and long-term categories, we do believe that a reasonable proxy for those may be found in the recommendations from the Smart Cities breakout session. They describe cross-domain interactions a more limited set of domains and stakeholders, but can serve as an illustration for a broader scope covered in this section.

Participants

- Florian Allerding
KIT – Karlsruher Institut für Technologie, DE
- Birger Becker
FZI – Forschungszentrum Informatik, DE
- Bert Claessens
Vito – Antwerp, BE
- Hermann de Meer
Universität Passau, DE
- Victor del Razo
TU München, DE
- Christoph Doblander
TU München, DE
- Frank Eliassen
University of Oslo, NO
- Nicolas Gast
INRIA – Grenoble, FR
- Christoph Goebel
TU München, DE
- Manuel Görtz
AGT International – Darmstadt, DE
- Kai Heussen
Technical Univ. of Denmark – Lyngby, DK
- Hanno Hildmann
NEC Laboratories Europe – Heidelberg, DE
- Longbo Huang
Tsinghua University – Beijing, CN
- Hans-Arno Jacobsen
TU München, DE
- Randy H. Katz
University of California – Berkeley, US
- Srinivasan Keshav
University of Waterloo, CA
- Jean-Yves Le Boudec
EPFL – Lausanne, CH
- Florian Michahelles
Siemens Corporation – Berkeley, US
- Milan Milenkovic
Intel – Santa Clara, US
- Johanna Myrzik
TU Dortmund, DE
- Peter Noglik
ABB AG Forschungszentrum Deutschland – Ladenburg, DE
- Mario Paolone
EPFL – Lausanne, CH
- Anthony Papavasiliou
University of Louvain, BE
- Yvonne-Anne Pignolet
ABB Corporate Research – Baden-Dättwil, CH
- Jose Rivera
TU München, DE
- Jiri Rojicek
Honeywell Prague Laboratories, CZ
- Hartmut Schmeck
KIT – Karlsruher Institut für Technologie, DE
- Mischa Schmidt
NEC Laboratories Europe – Heidelberg, DE
- Hans-Peter Schwefel
FTW Forschungszentrum Telekommunikation Wien GmbH, AT
- Joachim Sokol
Siemens AG – München, DE
- Thorsten Staake
Universität Bamberg, DE
- Kai Strunz
TU Berlin, DE
- Verena Tiefenbeck
ETH Zürich, CH
- Anwar Ul Haq
TU München, DE
- Andreas Veit
Cornell University – Ithaca, US
- Holger Ziekow
AGT International – Darmstadt, DE