Energy Services for the Smart Grid City

Stamatis Karnouskos, Per Goncalves Da Silva and Dejan Ilic
SAP Research, Karlsruhe, Germany
Email: {stamatis.karnouskos,per.goncalves.da.silva,dejan.ilic}@sap.com

***Abstract*—The smart grid enabled city is an emerging complex system of systems where different stakeholders will have to strive towards achieving their goals while interacting with each-other. At parts of the city such as the districts, the energy signature and efforts towards its better energy efficiency will heavily depend on the utilization of availability and optimal usage of the local resources. The latter may be very dynamic and depend on several complex conditions such as weather, prosumer behavioural patterns, business interactions etc. An emergent city-wide behaviour appears when a number of simple entities i.e. the prosumers, operate in an environment, forming more complex behaviours as a collective. In order to empower the smart grid city, several energy services capturing the common needs of all stakeholders need to be made available to them. As a result of a (potentially) common platform that offers basic energy services, rapid development of applications can be realized without the need to start from scratch. Such energy services have been identified, analysed and implemented, in the context of a wider enterprise system architecture. An insight on their functionality, usage and development challenges and experiences is provided.**

I. Introduction

The smart grid vision [1] heavily depends on the increased energy data acquisition granularity for understanding better how energy is produced, consumed and where fine-grained and timely adjustments can be done. Hence, monitoring and control will increasingly play a key role for the future smart grid infrastructure and the applications that will depend on it [2]. It is also expected that the future energy monitoring and management systems will be in close cooperation with enterprise systems and heavily depend on IT technologies [3]. Monitoring and timely management becomes even more important as dynamic entities such as the prosumers i.e. users that can fit into the roles of both producer and consumer, are interacting with other traditional stakeholders.

Districts and neighbourhoods will also increasingly play a key role in the smart cities, as they are expected to be able to autonomously manage their energy resources e.g. a public lighting system, a shopping mall, a PV or wind farm etc. By offering a way to enable business oriented interaction among the stakeholders, one may achieve better energy management as well as enhance the procurement of energy from external providers. To achieve this, appropriate energy services must be in place, that will integrate information coming from highly distributed smart metering points in near real-time, process it, and provide an insight upon which appropriate decisions can be taken.

A platform providing several Internet-accessible services has been designed and implemented which, in turn, can be used to create mash-up applications that deliver customized functionality and additionally let the stakeholders interact in a market-based way [4]. Information coming from highly distributed smart metering points has been integrated, and made available in near real-time via various services. Decision making applications rely on them to provide sophisticated functionalities both on the consumer as well as the energy provider side. This is significant for the energy domain, as we propose to move away from traditional heavyweight monolithic applications, towards a more dynamic mash-up application development environment. The NOBEL project (www.ict-nobel.eu) is realising and trialing with real-users in the city of Alginet in Spain, such a set of energy services.

II. Architecture

The vision of lightweight and rapid mash-up application development, where Internet services are used as basis for creating dynamically customized applications at the end-user side is being increasingly adopted as the standard for popular large-scale online services e.g. in Facebook, Twitter, Amazon, etc. Adoption of such trend could indicate a paradigm change for the energy domain, where heavyweight monolithic applications are substituted by far more dynamic, real-time and interactive applications. The designed and implemented architecture offering such energy services is depicted in FMC notation (www.fmc-modeling.org) in Figure 1. As it can be seen there are several layers i.e. the device layer, the middleware, the enterprise services and mash-up applications. The embedded devices, such as smart meters, concentrators, and generally any prosumer device are representing the lowest layer in the architecture. A middleware layer is acting as an information acquisition and processing component that connects to and aggregates functionalities near to the point of action (e.g. the meter). On the service layer we have various enterprise services as part of the Integration and Energy Management system (IEM), that put the data into business context and provide sophisticated functionalities.

The IEM is the "heart" of our approach, which is composed of several services that provide the core of business services. They foster the Software-as-a-Service (SaaS) approach and are hosted in the "cloud". Typically, such services can be mashed up in order to provide key functionalities for applications such as an energy portal, mobile applications, and also a district monitoring and management centre. Apart from the middleware assisted communication and data processing for obvious reasons, the devices can also directly communicate with the IEM and vice versa. It is expected that any proprietary communication (e.g. with legacy meters) can be wrapped by

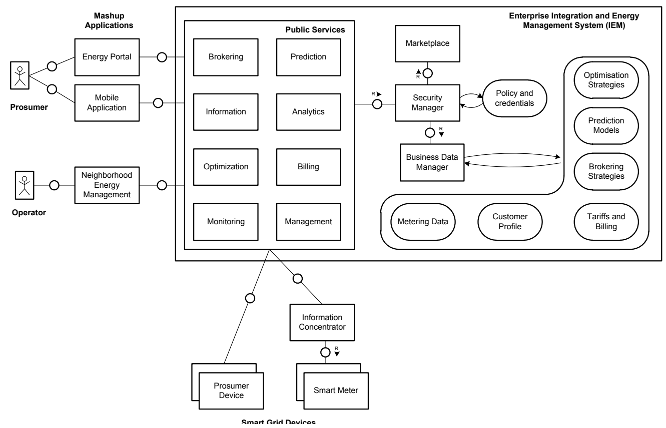

Figure 1. Overview of the System Architecture

gateways, or at most at the middleware level. At this stage, an event-driven interaction is done among all elements of the system, while striving for operational independence.

The IEM services have been fully implemented and tested with meters that have IPv4/IPv6 connectivity and can report their measurements to a concentrator or directly to a smart metering platform. However for practical reasons, ongoing validations assume either encapsulating legacy meters and their protocols or acquiring the data directly from the concentrator. The concentrator acts as a gateway for the communication to the meters (in case a non-IP proprietary protocol exist between them), as well as an information concentrator i.e. collecting and processing meter readings from the attached devices.

The IEM platform does not make use of typical heavyweight SOAP web services. Instead, a Representational State Transfer (REST) approach [5] is used in order simplify the API and its implementation, as well as to enable rapid application integration over the Internet. This imposes some architectural style selections e.g. client-server approach, stateless interactions, uniform interfaces and a layered system. Since all of the services will run on the "cloud", and may be hosted on various servers, we consider at design time, open interfaces accessible from anywhere via the web. As such, it is expected to obtain an easier integration, more flexibility and higher performance than the typical SOAP-based web services. All interactions among the key parts of the architecture depicted in Figure 1 are also done via REST and no strong requirements on co-location exist.To be able to adhere to potential scalability and performance issues, these could be hosted in dedicated servers specially customized for delivering maximum task-specific performance. An evaluation of these aspects is considered future work. Additionally, REST seems to be the technology of choice for APIs directed at sharing data easily (e.g. Facebook, Twitter, Amazon, and others use it). It is expected that, due to the volume of data collected in the smart metering era, users will be given a higher degree of control and access over it. Using a RESTful approach, may open the door for energy users to leverage their data against services provided by third parties.

Additionally, it should be easy enough to use existing enterprise services as building blocks in other more complex services in a mash-up way. The diversity of the involved entities leads to increased complexity in interacting and managing the approach. Challenging issues such as the overall system management and performance expectations were considered at design phase. Adhering also to uniformity, we try to provide interfaces that can be called in same fashion from any kind of device (bound to functionality), for instance, concentrator, or end user device, while the resulting output is intended for automated friendly integration (and not really humans) in the Internet of Things.

III. Realized Enterprise Services

Several services as depicted in Figure 1 have been designed and implemented. All of the IEM services have been implemented as *Java REST services* and are deployed in a *Glassfish 3.1 Application Server* (glassfish.java.net). The business data are stored in a *mySQL* DB (www.mysql.com). Specialized analytics and statistics are realized mostly on *R language* (www.r-project.org). All communication with the IEM is done via an encrypted channel i.e. *HTTPS* and a security (with role-based authorization and authentication) framework is in place based on *Apache Shiro* (shiro.apache.org). Additionally for performance reasons, all services interact using *Google Protocol Buffers* (code.google.com/p/protobuf/) which offer a highly efficient binary format.

A. Energy Monitoring

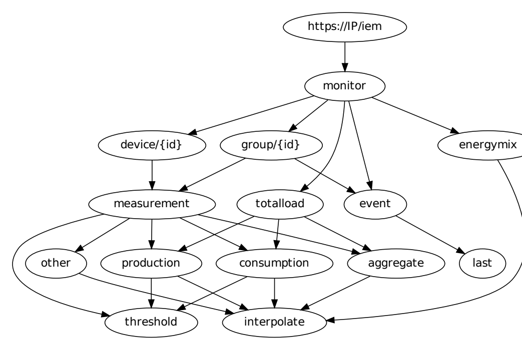

Figure 2. Energy Monitoring Service Overview

The *Energy Monitoring* service is responsible for acquiring and delivering data related to the energy consumption and/or production of a prosumer device. It offers a near real time view of the energy consumption/production as reported by the smart grid prosuming devices (e.g. PV, smart meter, electric car etc.). As this is a key functionality coupled with privacy concerns, it is made sure via the security framework, that only authorized users may access a subset of its functionality. For instance, users should only have access to their own production/consumption data, and group managers to the aggregated data of the group. The group functions enable monitoring of a group of devices, which provides extra flexibility and support for community-driven smart grid [6] behaviours. Figure 2 depicts a possible structure of the RESTful web service; one can derive the actual URIs of the service that can be called by following a path in the graph.

B. Energy Prediction

The main goal of the *Energy Prediction* service (structured as depicted in Figure 3) is to provide forecasts for energy consumption or production given a context e.g. historical information, weather prediction, prosumption device capabilities etc. Forecasts can be used by the users and operators to help with their electricity planing and trading activities; hence it would, for instance, enable operators to take advantage of opportunities, such as bidding at national energy markets, or even to comply with energy market regulations. A modular approach was taken where new algorithms can be integrated on the backend of the service without altering its API. Additionally practice has shown that sometimes users (who have better knowledge of their future behaviour) want to be able to adjust or further customize the prediction, hence such adjustments are also possible.

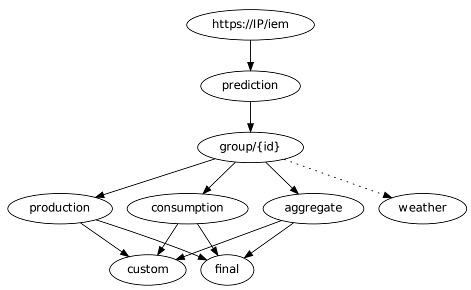

Figure 3. Prediction Service Overview

Of significant importance is again the group support, as an operator could via this service get real-time prediction of the energy production/consumption for one specific neighbourhood, building or even a dynamic group of prosumers e.g. formed based on social, financial, behavioural or other criteria. As the focus is on the architecture and functionality provision, the independence of the prediction algorithms realization and the availability of a stable service interface so that new prediction models can be plugged-in without breaking the interactions with the service's clients was carefully considered. The *R language* is heavily relied upon for the implementation of this service, as a multitude of algorithms and statistic models are already available via it.

C. Management

The *management* service (structured as depicted in Figure 4) facilitates the management of users and devices (e.g.smart meter, concentrator, etc.). Through it various stakeholders, for instance the operators, are able to manage groups or users, assign devices to them, maintain their information etc. The aim is to ease infrastructure management by making widespread changes easy and efficient. As an example, the energy threshold can be set up for individual users or for all users in the group e.g. a neighbourhood. While these services are mainly targeted at the operators, a small subset of them will be available to the users, enabling them to update personal data, or even other optional information to enable higher quality of services, such as load profile optimization, etc.

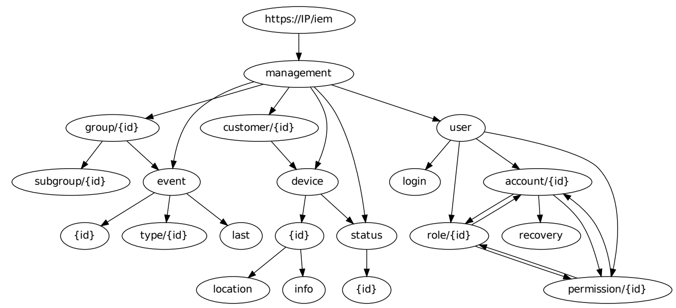

Figure 4. Management Service Overview

On the operator side, a full view of the infrastructure can be depicted, searched and managed. Hence, data provided via this service can be used by mash-up applications such as mobile applications and energy portals for end-user visualization. This service enhances traditional business processes, for instance, smart meter installation, connection/disconnection of the user, device status, etc. In the carried out experiments the goal is to have automatic information update especially from the dynamic aspects e.g. when a smart meter or concentrator connects to the infrastructure. This will enable full monitoring not only of devices joining or leaving the network, but also of their capabilities, services they might offer, as well as real-time status. Enhanced device management has the potential to empower business practices related to: approval process, procurement management, life cycle management, (re)deployment and disposal management. Additional value is expected to uncover savings through process improvement and support for strategic decision making, better asset management etc.

D. Energy Optimization

The *energy optimization* service (structured as depicted in Figure 5) provides the user with load profiles that are optimized for a particular set of constraints. As the constraints may vary depending e.g. on personal preferences, usage signature, contractual agreements etc. a multi-objective optimization approach needs to be followed. We consider as input data information coming from historical usage, predicted usage, contract information, and an optional list and schedule of load generating devices that operate within the user premises.

Given the user preferences and the available data, the optimization service aims at providing a suggested load profile that the user may want to adhere to, in order to attain more benefits, according to his goals and preferences. The optimization service may also provide additional information e.g. a disaggregated load profile that suggests which devices should be used and when they should be used (e.g. in case of shiftable loads that correspond to respectively flexibly adjusted activities). However, it should be noted that the quality of the output of such a service will highly depend on the quality

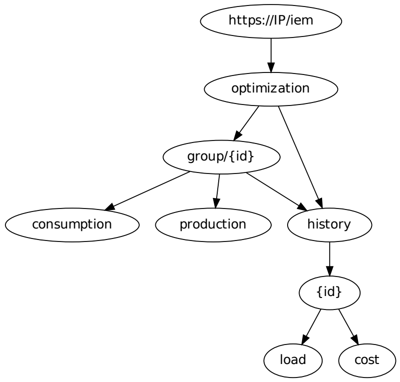

Figure 5. Optimization Service Overview

of available information. A typical case investigated within the NOBEL project is that of the optimization of energy by interacting with a public lighting system (and using it as an energy balancing partner).

Special focus is also given to optimizing larger groups of users and not only standalone ones. This is expected to assist the group operators to better understand their behaviour as well potentially use the disaggregation functions to further optimize based on the energy signature and capabilities of their group members [6]. Optimizations may have various performance indicators as key guiding ones e.g. targeting energy or power, the final cost the user would like to pay, special preferences (e.g. always consider "green" electricity) and market-driven energy procurement.

E. Billing

The next generation of billing services will include not only support for the energy traded with the local neighbourhood market [4], but also include a very detailed analysis on the energy behaviour and bill e.g. nature of electricity consumed, and later even mapping to the individual user devices (if wished). All of it of course should be done per user request in near real-time. Such a service, an overview of which is provided as a directional graph in Figure 6, may assist users to not only better understand but also plan their costs.

The user probably will not cover his whole electricity needs via what is traded in the local marketplace, but also have contract(s) with energy providers, which may also have variable tariffs. Providing real-time billing that also includes micro-contracting, dynamically changing tariffs and disaggregation of costs per device or user-defined activities (e.g. social event etc.), will be challenging.

F. Brokerage

The Brokerage service covers one key functionality envisioned for future smart grid neighbourhoods: the capability of being able to trade energy [4] in local marketplaces i.e. buy energy needed and also sell surplus. This service aims at enabling these brokering capabilities in order to allow users (or their surrogate agents [7]) to participate in the electricity marketplace [8]. As such, this service targets the operators,

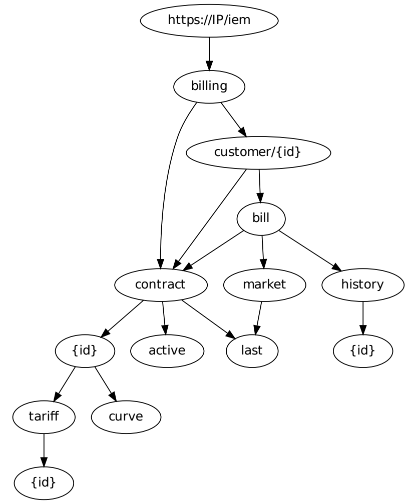

Figure 6. Billing Service Overview

for managing the market, the users, for actively buying/selling electricity, and also for their automated brokerage agents. A user's brokerage agent acts as a proxy for the user [9], and can be e.g. configured via a mobile device and act driven by the user's preferences, even at time when the user is unavailable. Although the brokering related service is designed with the goal to fulfil the energy market trading needs for the NOBEL project, it is considered that they are general enough to accommodate most market types. Aside from services enabling market participation, like placing an order, simple market analysis tools such as price and volume statistics, are also offered to aid the user or the automated agents in making trading decisions.

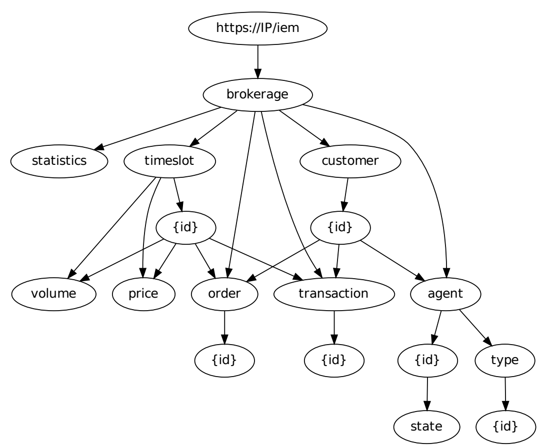

Figure 7. Brokerage Services Overview

The structure of the brokerage service is depicted in Figure 7. The service enables the participants to place orders, and retrieve information regarding transactions (successful bids), orders, and trading time slots. It also provides auxiliary functionalities such as statistics regarding trading price and volume, as well as the last traded price for each time slot. The `order` part of the service returns current market information on the order, as list of orders in a timeframe, or in a particular time slot. Transactions are made by matching, either fully or partially, the orders in the order book. They represent the contract made by each party to consume or deliver the specified amount of energy at the agreed price. The time slot service returns information regarding each time slot, for instance, the state of the time slot (open, closed, trading, etc.), as well as the start and end of the time slot, and its energy delivery timeframe. Additionally, under the "statistics" other services return more fine-grained information e.g. the "lastprice" returns a list of last traded prices for each time slot in a timeframe, the price curve returns a list of traded prices for a timeslot, or a timeframe, and the volume service returns a list of traded quantities per timeframe, or timeslot.

G. Other services

There are several other services which are under implementation. For instance the informational service is providing notifications to the users e.g. possible (urgent) messages from the energy provider, additional informational material such as network warnings, advertisements, news etc. The same service can be used to enable bidirectional interaction between the end-user and his provider e.g. for feedback, maintenance, communication of problems in his service etc. It is also considered, that in the future, the analytics service might be of key importance. This would enable for instance enterprise users to further customize analytics engines to do data mining on the available information gathered at IEM level.

Figure 8. A mash-up web application based on the IEM services

Without explicitly focusing on strong security aspects, state of the art approaches are used to cover the basic requirements for a real-world trial of the whole system that is taking place in the Spanish city of Alginet in 2012. Role-based access for authorization, login/password combinations for authentication and secure channels (via HTTPS) for all interactions of the IEM with external entities, as well as between the IEM services, are used.

IV. Discussion and Lessons Learned

During the process of creating multiple mash-up applications for a variety of stakeholders (residential end-users, utility etc.), valuable lessons, both technical as well as other related to design and social aspects, were acquired. To validate the concept and the implementation of the IEM services, a web application was created composed solely from a mash-up of the IEM services. An example of the GUI is depicted in Figure 8, while similar applications that take advantage of the same services are also built for mobile devices.

It would have been painful to create in the same timeframe and with the same functionality a monolithic application as it is traditional to do so in the energy domain. Without access to the energy data via open interfaces, the implementation overhead would be too costly to be realized, especially taking into account the multitude of end-user devices that were targeted. Hence the need for a layered approach, where the data acquired is processed and offered via open and long-lived interfaces is an absolute must. The provision of necessary functionalities over Internet services enables the creation of apps that can run any time, anywhere, and interact with user's data as well as additional services in the "cloud". As such it is crucial that basic energy services are defined and standardised in order to enable a common approach on harnessing the energy data and empower the next generation of energy applications.

If the sophisticated logic had to be separately developed for each application then it would probably lead to repetition of logic and functionality with a result of more fat clients (end-user applications). However, in this case, things were easier since a platform that offered basic common energy services to all users was developed. The common requirements for end-user applications were identified early and built into the platform. Currently, the initial basic set of services is being maintained, while the platform is extended with additional sophisticated functions. It is expected that in the future such issues will be negotiated by the application developers and the platform providers in order to ensure shared benefits.

Using technologies that would enable several other aspects e.g. lightweightness, high performance, backward and forward compatibility etc. were considered. All services are Internet based and use the HTTP as a complete application protocol, which also defines the semantics for the service behaviour (as followed by RESTful approaches). Additionally, to the REST style (with PUT/GET/DELETE/POST methods) the Google Protocol Buffers (extremely efficient binary format) for enhanced performance was used, since significant amounts of data had to traverse the network from the platform to the applications. Bulk data transfer was also used, instead of many smaller messages, to make the application more network friendly and decrease network round-trips.

Services may fail at any point of time (both on the client and server side), even when processing a request. Typical such incidents that occurred during the trials were examples where the network connection failed (disrupted communication link, time-outs), or even a case of server overload. It is therefore important to implement both at client and server strategies in order to detect and handle such failures.

Sometimes the server processing took extensive time (due to the nature of the request or server overload) and in the meantime the client either timed-out or was blocked. Hence, asynchronous behaviour was implemented (Request/Acknowledge) instead of synchronous (Request/Response) where possible, in order to avoid the client blocking. A typical case for this was when the user would request in detail depiction of his energy data for a long period of time e.g. 10 months. Here one can either consider sending the request and polling later for results (hence the responsibility is at the client side) or give a callback URI to the server so that he can push the results whenever ready. Both strategies will have to consider several aspects e.g. the frequency of polls as well as operational aspects e.g. if the application is behind a firewall the respective port must be open etc. Using a publisher/subscriber model might be difficult, especially in the case of mobile phones, which are generally behind a firewall for security reasons. One alternative, yet to be explored, is the use of `websockets`. This technology, which is available in HTML5, is somewhat between the request/response and publisher/subscribe models. The client can initiate a permanent connection to server, effectively giving it a direct route back to the client, and thus enabling it to deliver information when it is available.

An additional issue to consider is that for multiple data exchanges the applications should use one connection instead of establishing and destroying connections each time it is needed (this helps minimizing the latency). This has an impact on the server side since sufficient system resources must be available (HTTP is a connection oriented protocol which implies high availability of the server) and at sufficient capacity (e.g. CPU, memory, network bandwidth) to handle potentially large numbers of clients or spikes (when they send requests at very narrow timeslots).

Since the application developers do not have an overview of the number of energy services available and their API, it would be beneficial to have a Service Registry where one could search for services and their end-points. This will also ensure that whenever services are modified e.g. changing URI patterns and locations, the applications (clients) can discover them. Additionally this is considered beneficial for the infrastructure maintenance in order to add or remove services.

Security, trust and privacy are challenging issues especially associated with the emerging smart grid capabilities. Currently, the IEM does not focus on security, however, an effort was made to use standard approaches. For instance, only secure channels (all REST calls were made over HTTPS) were used, as well as basic HTTP authentication and authorization were used by all services provided by the platform. Clearly, precautions must be carefully considered for real-world deployments which might additionally include message signing and encryption of service hyperlinks.

V. Conclusion

Significant effort is invested towards creating innovative applications for the emerging SmartGrid. However, it seems like many of these efforts use the old modus operandi of the energy domain providers i.e. offer of a monolithic application residing on proprietary or provider-only data and with little thought about cross-application collaboration and information exchange. However, the future of SmartGrid applications resides on sophisticated web applications that can rely on a multi-source data and Internet provided basic services that can be easily customized for the specific end-user groups.

The approach presented here and the experiences acquired while creating a platform for energy services that can be used to empower end-user mash-up web-based applications have demonstrated the need for it as well as its promising potential. The goal is to enable developers to start from an abstract level of energy-providing services so that they can focus on rapid development of lightweight applications with innovative features that enable users to monitor energy, trade electricity, assess the energy efficiency of appliances and for the providers to manage their customer base. Possible innovative services that could benefit the future smart grid city have been identified and investigated. Their functionality in a wider vision has been presented and discussed in conjunction with a proposed implemented and assessed system architecture.

Acknowledgment

The authors would like to thank the partners of European Commission co-funded project NOBEL (www.ict-nobel.eu) for the fruitful discussions.

References

[1] "SmartGrids SRA 2035 – Strategic Research Agenda," European Technology Platform SmartGrids, European Commission, Tech. Rep., Mar. 2012. [Online]. Available: http://www.smartgrids.eu/documents/20120308_sra2012.pdf

[2] X. Yu, C. Cecati, T. Dillon, and M. Simões, "The new frontier of smart grids," *Industrial Electronics Magazine, IEEE*, vol. 5, no. 3, pp. 49 –63, sept. 2011.

[3] R. Katz, D. Culler, S. Sanders, S. Alspaugh, Y. Chen, S. Dawson-Haggerty, P. Dutta, M. He, X. Jiang, L. Keys, A. Krioukov, K. Lutz, J. Ortiz, P. Mohan, E. Reutzel, J. Taneja, J. Hsu, and S. Shankar, "An information-centric energy infrastructure: The berkeley view," *Sustainable Computing: Informatics and Systems*, 2011.

[4] S. Karnouskos, "Demand side management via prosumer interactions in a smart city energy marketplace," in *IEEE International Conference on Innovative Smart Grid Technologies (ISGT 2011), Manchester, UK*, Dec. 5–7 2011.

[5] R. T. Fielding, "Architectural styles and the design of network-based software architectures," Ph.D. dissertation, University of California, Irvine, Irvine, California, 2000. [Online]. Available: http://www.ics.uci.edu/~fielding/pubs/dissertation/

[6] S. Karnouskos, "Communityware smartgrid," in *21st International Conference and Exhibition on Electricity Distribution (CIRED 2011), Frankfurt, Germany*, 6-9 June 2011.

[7] ——, "Agent-based mediated control in smart grids," in *IEEE International Conference on Innovative Smart Grid Technologies (ISGT 2011), Manchester, UK*, Dec. 2011.

[8] D. Ilic, P. Goncalves Da Silva, S. Karnouskos, and M. Griesemer, "An energy market for trading electricity in smart grid neighbourhoods," in *6th IEEE International Conference on Digital Ecosystem Technologies – Complex Environment Engineering (IEEE DEST-CEE), Campione d'Italia, Italy*, Jun. 2012.

[9] P. Vytelingum, S. D. Ramchurn, T. Voice, A. Rogers, and N. R. Jennings, "Trading agents for the smart electricity grid," in *9th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2010), Toronto, Canada*, 10-14 May 2010, pp. 897–904.