

Simulation of a Smart Grid City with Software Agents

Stamatis Karnouskos, Thiago Nass de Holanda
SAP Research (www.sap.com)
Email: stamatis.karnouskos@sap.com

Abstract—In the future smart city, new information and communication technologies will enable a better management of the available resources. The future smart grid infrastructure is emerging as a complex system where fine-grained monitoring and control of energy generating and/or consuming entities within the electricity network is possible. This will result to better approaches that will boost energy efficiency. A simulation of a dynamic ecosystem such as the smart city, will enable us to test new concepts and resource-optimization approaches. Therefore we have analyzed, designed, and build a simulator based on software agents that attempts to create the dynamic behavior of a smart city. It simulates discrete heterogeneous devices that consume and/or produce energy, that are able to act autonomously and collaborate. The behavior of these devices and their groupings e.g. smart houses, has been modeled in order to map as near as possible the real behavior patterns of the respective physical objects.

Keywords-Simulation software, energy measurement, mobile agents, web services

I. INTRODUCTION

Advances in the areas of embedded systems, computing, and networking are leading us to an infrastructure composed of millions of devices. These devices will not simply convey information but compute on it, network, and form advanced collaborations [1]. This “Internet of Things” infrastructure will be strongly integrated with the surrounding environment, and additionally it will closely interact with the enterprise systems. The last will lead not only to further blur the line between the business and real world, but will change the way we design, deploy and use services. New opportunities will emerge for businesses, which can now closely collaborate with the real world.

In the future service-based Internet of Energy (IoE) [2], several alternative energy providers, legacy providers, virtual power plants, households etc. that are the elements of smart cities are interconnected. Via smart meters, one is able to interact with a service based infrastructure and perform actions such as selling and buying electricity. More advanced services are envisaged that will take advantage of the near real-time information flows among all participants [3].

In future smart cities, there will be several entities that will act interchangeably as consuming or producing devices (“prosumer” devices) e.g. electric cars. In this infrastructure the devices will be no more considered as black-boxes but will also get interconnected, which will provide fine-grained info with respect to their energy behavior. It is

also expected that they will provide their functionality as a service and be able to consume online services (Internet of Services) in order to better address their internal goals. Existing efforts in the emerging Internet of Things and Internet of Services, will be combined and be a crucial part in the envisioned Internet of Energy (commonly known also as smart grid or IntelliGrid). Several efforts are underway worldwide e.g. the European technology platform Smart-Grids (www.smartgrids.eu) and EPRI’s IntelliGrid initiative (intelligrid.epri.com).

II. OBJECTIVES

The smart grid infrastructure envisioned is highly dynamic, and by being able to tap to the information generated by its discrete items i.e. generating and consuming devices in real-time, new possibilities are emerging towards real-time adaptation, optimization and prediction. Today most of the tools assume static profiles matching a specific energy distribution, as we do not have the capability of drilling down to specific devices. Since this is expected to change in the near future, we consider a simulation of such an infrastructure, as the first step in order to use it as a basis for further testing of management algorithms, behaviors, business models [4] etc.

The objective of this work is to analyze, design, implement and evaluate an approach that simulates a dynamic infrastructure within the broader context of the envisioned smart city and smart grid infrastructure. As such some guiding requirements arise:

- have direct access to energy devices and their information
- simulate in a timely fashion (almost real-time) the energy consumed or produced
- being able to apply control functions to the devices and change their behavior
- provide service-oriented interaction capabilities for easy integration with enterprise applications
- consider devices or groups of them as autonomous goal-oriented acting entities
- generate real-time data and reports
- provide basic capabilities for negotiation, control, reactivity and adaptation
- be as near as possible to real-world device behavior pattern

As such, a dynamic infrastructure has to be created, that once up-and-running different algorithms and concepts can be tried out. One key goal is to have this infrastructure created as close as possible to a real-world model, in order to evaluate different scenarios. Thus, the overall energy consumption should be modeled in order to be within the data range that various studies have shown.

III. TECHNOLOGIES

In the simulator all entities have been implemented using the software agent paradigm [5] and more specifically the JADE [6] multi-agent system (MAS). MAS are often used by researchers in energy systems [7] in an effort to introduce intelligence and distributed management.

JADE is a popular MAS used for the development of distributed applications. The intelligence, the initiative, the information, the resources and the control can be fully distributed on mobile terminals as well as on computers in the fixed network. Each energy device has is mapped to one agent. JADE offers full support of the autonomous, adaptive and negotiation capabilities for agents, which enables us to implement the envisioned characteristics for the simulator. Finally the support of standards such as the FIPA (www.fipa.org), ensures that our approach is open enough for future interaction by using the Agent Communication Language (ACL) of FIPA with other multi-agent systems that offer such support.

As the simulator should provide open access to be controlled by external entities we have decided to use web services as a mean to interact externally with it. A requirement is that the results as well as selected management functions should be available to the enterprise applications and other entities. Using web services is the best way at the moment, as the enterprise world uses by de facto web services for interacting with any system. JADE provides a simple web service integration gateway which enables seamless and transparent connection with JADE platforms by offering bidirectional discovery and remote invocation of web services by JADE agents, and of JADE agent services by web services [6].

IV. SIMULATOR OVERVIEW

The architecture of envisioned system is composed by three layers. The top layer is the enterprise application layer and contains the business logic. This layer communicates with the simulator layer via web services. As such an interoperable interface is offered to any business application in order to directly access functionality of the simulator, change parameters in the runtime and also get results in a timely manner.

The middle layer is the logic and strategy and contains the main rules and configurations for the correct operation of the developed simulator layer during all the defined scenarios.

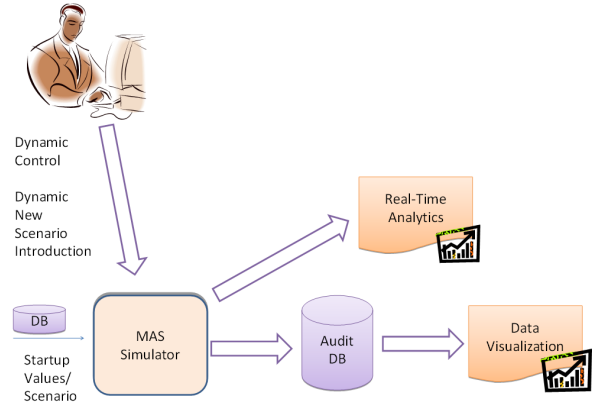


Figure 1. Simulator Overview

This layer communicates with the agent simulator layer via (local) method invocations.

The bottom layer is the simulator layer and contains all the agents which represent energy generating and/or consuming devices. The agents utilize the full communication capabilities offered by the agent platform and execute in order to offer the desired functionality. The interactions among the agents enable us to realize complex scenarios.

As each agent represents a real-world device on the grid (either standalone or collections of them), they have a different energy signature depending on the appliances and devices they simulate. We are able to track in real-time the populations of them, their status and create interactions among them.

An overview of the simulator system is seen in Fig. 1. The main application is the MAS-based simulator which gets its startup configuration (values and scenarios) from the database. The user has the power to dynamically control the simulator by interacting with the available agents. He can directly call the functions available on the agents as well as deploy or remove agents that interact with the rest during runtime. Being able to intervene at runtime gives us the possibility to run different scenarios standalone or in parallel as users on a very dynamic system. The real-time analytics show the current overall consumption and generation of energy. Data is also stored at the audit database and can later be used for more detailed analytics as well as future scenarios including possible prediction algorithms. The last could lead to the development and analysis of prediction models that use the simulator as a tool to test their behavior but is out of the scope of the current work.

V. SIMULATED ENTITIES

Effort has been invested to depict the real-world structure and behavior in the simulator, considering the smart city and smart grid era. As such we have devices that can move around e.g. electric cars, house appliances, groups of devices e.g. households as well as logical limits e.g. cities etc.

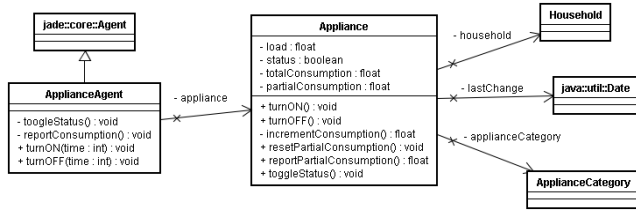


Figure 2. Appliance Implementation View

The main concepts behind the simulator structure with respect to the simulated entities include:

- **Households:** Each household belongs to one city and it is randomly placed within the city area. Households have a number of appliances that consume and/or produce energy.
- **Appliances:** Here mostly white goods are implied. Each appliance belongs to one household from where it consumes energy. The rate at which the energy is consumed is different for each appliance, within the average limits that several studies have found. Each appliance can have different states that affect its energy consumption profile. An appliance also belongs to an appliance category (as seen in Fig. 2) that holds information for the load distribution, presence and schedule for all the appliances of that category. The load is represented by a probability distribution function and indicates how much energy the appliance consumes over time.
- **Vehicles:** Electrical vehicles are expected to be integral part of smart cities and play the role of prosumer devices in the smart grid. The simulated vehicles are able to move to and between different cities and recharge from time to time. Each vehicle belongs to one household where it recharges.
- **Cities:** Each smart city starts with a number of households, vehicles and power stations. A city can run as a standalone or connected to others. Its area is placed at the origin of a coordinated system or as a neighbor to an existing reference city e.g. on north, south, east or west.
- **Power Stations:** The power station represents the energy generation part of the simulator. It holds the information of maximum capacity and relative usage. The maximum capacity is stored as a normal distribution. Every time a new power station is created it gets a new maximum capacity value according to this distribution and then keeps the same value for its lifetime.

In the implementation part, all of the entities are considered to be in binary status i.e. either in OFF or ON state and change between these states e.g. via their internal schedule or external command. The internal schedule of an appliance is stored at the appliance category inside the `timeON` and

| Category | Schedule time ON | Schedule time OFF | Power |
|--------------------------|------------------|-------------------|---------|
| Central Air Conditioning | 15 min | 35 min | 1,064 W |
| Room Air Conditioning | 15 min | 35 min | 221 W |
| Main Space Heating | 15 min | 35 min | 1,341 W |
| Furnace Fan | 15 min | 35 min | 190 W |
| Refrigerator | 15 min | 35 min | 471 W |
| Freezer | 15 min | 35 min | 395 W |
| Dishwasher | 1 h | 23 h | 1,403 W |
| Range top | 1 h | 23 h | 1,468 W |
| Oven | 1 h | 23 h | 1,205 W |
| Microwave Oven | 0.3 h | 23.7 h | 1,909 W |
| Water Heating | 30 min | 70 min | 971 W |
| Lighting | 8 h | 16 h | 322 W |
| TV | 8 h | 16 h | 47 W |
| PC and Printer | 4 h | 20 h | 263 W |
| VCR/DVD | 2 h | 22 h | 96 W |
| Clothes Dryer | 1 h | 23 h | 2,956 W |
| Clothes Washer | 1 h | 23 h | 329 W |

Table I
APPLIANCE SCHEDULE AND POWER OVERVIEW

`timeOFF` parameters; these represent the time interval the device will stay ON or OFF (according to its profile) and are stored as a normal distribution function for all the appliances of the same category. Each time an appliance changes its status it gets a new time interval from the appliance category. In this way the appliances do not have a time-fixed schedule but rather change their status according to their individual profile.

We mentioned that we wanted to stay as close to reality as possible. The information for the energy profiles was acquired from the 2001 Residential Energy Consumption Survey (RECS) done by the U.S. Department of Energy (DoE), that provides information on the use of energy in residential housing units in the United States. From all the 33 appliances that RECS reports [8] only 17 represent more than 1% of the total household consumption, therefore we focus on them as in any case altogether these appliances represent 85% of total household consumption.

The U.S. DoE survey does not provide information about the daily usage of the appliances and their power, but only has the presence penetration and yearly consumption info. Thus for this work it was assumed that all the days during the year have similar energy consumption distribution. This is one limitation of this work and should be expanded in future work in order to support different energy consumption profiles within the year e.g. summer and winter, different consumption for workdays, weekends as well as other location or demographic criteria related effects.

It was also assumed that the appliance runs on a continuous ON and OFF schedule which is also another limit. Future work should look towards expanding it using a time sensing approach, were for different time intervals within one day one appliance could have different probability of being ON or OFF, in order to better reproduce the appliance

real usage. The schedule for an appliance is based on also on the yearly appliance consumption as reported in [8]. The appliance power consumption is shown in Table I.

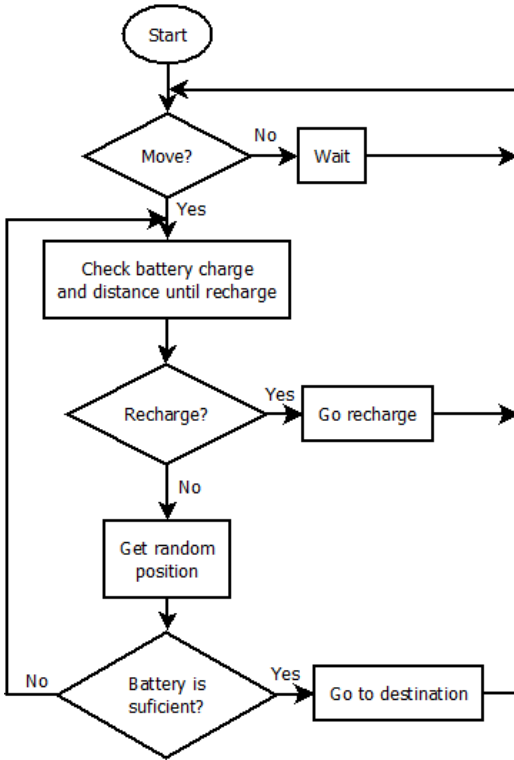


Figure 3. Vehicle Finite State Machine

The vehicle implementation class view is seen in Fig. 4. Its energy related information is acquired from the Advanced Vehicle Testing Activity (AVTA) of the U.S. Department of Energy. The primary goal of AVTA is to benchmark and validate the performance of special vehicles, including Plug-in Hybrid Electric and pure-Electric Vehicles [9].




Figure 4. Vehicle Implementation view

The behavior of the vehicle has been modeled using a finite state machine (FSM) provided by the JADE framework and depicted in the flowchart of Fig. 3. This is a cyclic behavior that keeps the vehicle moving as long as it has enough energy. As it moves the battery discharges and the probability to go for recharging increases. This probability is calculated at each new trip. By the end of each trip the

vehicle has a random probability to go to a new destination or just wait for some time. Before each new trip the vehicle verifies if it has enough energy to reach a destination and then go to recharge, in order to avoid a total battery discharge, which e.g. would mean the end of the trip for a pure electric vehicle.

Once the simulator is up and running, it generates energy consumption events that are stored in the database and sent to the real-time analytics application (via WS-Eventing) every second. This last tool consists of a GUI written in Java that depicts the real time data with respect to the overall energy generation and consumption from the simulator.

Another visualization tool has been also implemented that is more powerful than the first one and relies on the data available through the database. This tool can show charts even for scenarios that have already finished running, and therefore can help comparing the scenarios, but it takes more time to process the bulk database data stored than the real-time analytics tool which gets timely updates during the simulation run.

VI. EVALUATION

We have run a scenario simulation with 300 households evenly split between three smart cities. Each household started with appliances according with the appliance category presented on Table I. This resulted in 3840 appliances from all the 17 categories. The simulation included also 156 vehicles (52 at each city) and was running for more than 24 hours with all devices reporting their consumption every second to the database and the controller agent.

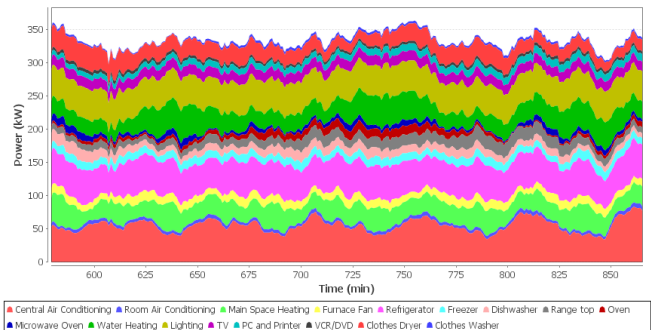


Figure 5. Power over time split in different appliance categories

Fig. 5 shows the power used by each appliance category over time. The time axis is in minutes and shows the relative time between the measurement and the simulation start time.

This scenario run generated 63×10^6 appliance consumption events, 8×10^6 vehicle consumption events and 480×10^3 power station generation events. Taking into account the number of appliances this makes an approximate average of 16×10^3 events per device. Dividing by the time the simulation was running, we result to 10 appliance consumption events per minute. The same information for

appliances, vehicles and power stations is shown in Table II. All this data combined used 7.5 GB of disk space on the database host at an approximate rate of 300 MB per hour of simulation.

| | Events | Devices | Events per device | Events per device per minute |
|----------------|------------|---------|-------------------|------------------------------|
| Appliances | 63,275,645 | 3,840 | 16,478 | 10.4 |
| Vehicles | 8,074,240 | 156 | 51,758 | 32.6 |
| Power Stations | 479,467 | 6 | 79,911 | 50.3 |

Table II
SCENARIO EVALUATION

As an example to what this simulated infrastructure has to offer we implemented a energy controller agent that continuously monitors the overall power generation and consumption measurements on the simulator. Any deviation beyond a specific predefined limit results to a control action. Assuming that there is too much consumption then e.g. the energy controller agent tries to turn OFF some devices and/or instantiate energy backups or start new generation resources. Alternatively if there is too much energy produced and fed into the network (e.g. due to wind from wind farms), then it tries to turn ON devices (e.g. refrigerators) and fully-charge the electric cars.

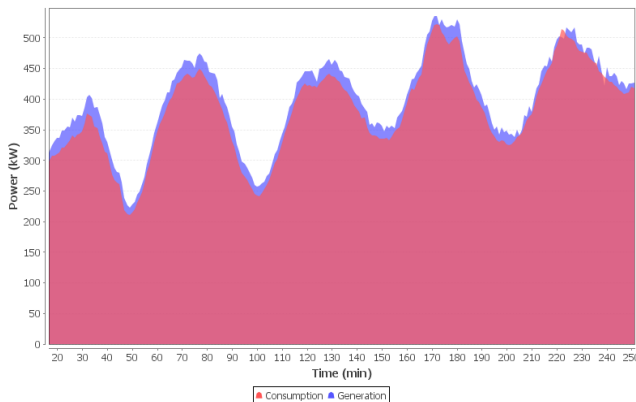


Figure 6. Consumption and generation chart

The result of this rather simple monitoring and control algorithm is to demonstrate how dynamic management can be applied and keep the generation and consumption within specific deviation limits, ideally very close. The results can be seen in Fig. 6 which shows the generation and consumption profiles over time for the specific simulation. From the figure it can be seen that the generation is always near to the consumption since this is was also the goal of the energy controller agent.

VII. DISCUSSION AND FUTURE DIRECTIONS

The work presented here is about the architecture, implementation and evaluation of a simulator based on software agents for simulating a dynamic smart city in the future smart grid infrastructure. This is a non-trivial task due to several requirements especially in the interaction among the various entities. The work presented here is complementary to our ongoing efforts in coupling simulations with real-world devices and connecting them to enterprise systems, however there are some notable differences. In our early efforts [10] we focused on connecting mobile agents with real web-service enabled devices; an approach that created an infrastructure that was also dynamically discoverable and where each device could interact via web services. In that train of thought we also validated that the Devices Profile for Web Services Protocol (DPWS) is fit with respect to scaling and reliability for the smart meters [1]. In this work we focus exclusively on the heterogeneity of devices and their consumption or generation profiles, as the main interest is on creating a smart city and focus on the interactions and dynamics of the infrastructure.

The appliances have been implemented using a continuous cyclic schedule. Further improvements can be done to allow the appliance to reflect different energy consumption profiles according to more fine grained parameters e.g. time/date/season etc. There should be an open and easy way to integrate various models that describe activity patterns of the devices, households etc in a smart city e.g. the pattern of occupants of a household [11].

Currently the appliances have been implemented as having two states i.e. ON and OFF but it is clear that these need to be expanded in the future in order to capture the multiple states future devices will support internally and will be of various levels. This is especially true for appliances that deal with complex tasks and change their energy consumption not deterministically over time to complete one task. Today only some of these devices e.g. white goods have various internal energy states but usually these are not exposed externally to the customer, and of course there is no external open management capability for other entities to interact with them. This is expected to change in the next years though, as devices are becoming more networked, with embedded intelligence and energy-aware.

The electric cars have been implemented only as entities that can charge and therefore take load of the grid but not vice versa. Thus they don't fully implement the prosumer profile that is envisioned in the future where electric cars can charge, store and when needed provide back electricity to the grid.

Since the focus of this simulation was to reproduce a residential scenario of energy consumption only the active load was taken into account. It is left as future work the extension for commercial and especially industrial consumer

scenarios, where the reactive load should also be considered. The reactive load is more relevant for this kind of consumers because they usually have a higher consumption and the electricity utility may impose constraints on how much reactive load might be available.

The electricity price was out of the scope of this work but future extensions of this concept, especially using a time varying price, could bring a new perspective and challenge for setting up the appliance schedule. This would result e.g. to appliances that are equipped with internal optimization algorithms as to complete all the tasks with the lowest running cost. There are already management functions available for the devices and the simulator. As such an external price-aware entity such as the PowerMatcher (www.powermatcher.net) [12] could take up the role of managing the devices with via the interfaces already provided and according to the prices available for the specific home or neighborhood.

VIII. CONCLUSIONS

In the context of this work, the motivation was to create a dynamic infrastructure that can partially simulate the behavior of the future smart grid city. Our aim was to realize a system that will be up and running and with which external entities can interact with. Software agents bear many of the social characteristics such as negotiation, autonomous decision making, collaborative interaction etc, and therefore the concepts have been implemented with them i.e. the JADE platform. The result is a highly dynamic infrastructure, broken down to its basic most energy consuming entities according to the energy profile surveys by U.S. DoE and therefore is near to the real-world model. Once up-and-running, enterprise applications and users can interact with it via the implemented management functions, but also new functionality can be introduced if one takes advantage of the software agent capabilities. Our goal to provide the basic infrastructure and tools to interact with it was therefore a success.

To demonstrate and validate the simulator, we have run several scenarios, where one can clearly see that the running simulation is very near to the model that was considered, which provides a very good tool for further research. Furthermore the real-time analytics as well as the historic data evaluation can both serve for users and providers in feeding them with timely info in order to ease monitoring of the infrastructure and enhance their decision making processes. Enterprise applications could evaluate scenarios in the simulator and see what the effects of actions could be, prior to applying them to the real infrastructure.

ACKNOWLEDGMENT

The authors would like to thank the European Commission and the partners of the SmartHouse/SmartGrid (www.smarthouse-smartgrid.eu), and CONET (www.cooperating-objects.eu) projects for their support.

REFERENCES

- [1] S. Karnouskos and A. Izmaylova, "Simulation of web service enabled smart meters in an event-based infrastructure," in *Proc. 7th IEEE International Conference on Industrial Informatics INDIN 2009*, Jun. 23–26, 2009, pp. 125–130.
- [2] S. Karnouskos and O. Terzidis, "Towards an information infrastructure for the future internet of energy," in *Kommunikation in Verteilten Systemen (KiVS 2007) Conference*. VDE Verlag, 26 Feb 2007 - 02 Mar 2007 2007.
- [3] H. Sui, H. Wang, M.-S. Lu, and W.-J. Lee, "An AMI system for the deregulated electricity markets," in *Proc. IEEE Industry Applications Society Annual Meeting IAS '08*, 5–9 Oct. 2008, pp. 1–5.
- [4] A. Vos, "Effective business models for demand response under the smart grid paradigm," in *Proc. IEEE/PES Power Systems Conference and Exposition PES '09*, Mar. 15–18, 2009, p. 1.
- [5] M. Wooldridge, *An Introduction to MultiAgent Systems*. John Wiley & Sons, 2009. [Online]. Available: <http://www.csc.liv.ac.uk/~mjw/pubs/imas/>
- [6] F. L. Bellifemine, G. Caire, and D. Greenwood, *Developing Multi-Agent Systems with JADE (Wiley Series in Agent Technology)*. Wiley, April 2007.
- [7] D. A. Cohen, "Gridagents: Intelligent agent applications for integration of distributed energy resources within distribution systems," in *Proc. IEEE Power and Energy Society General Meeting - Conversion and Delivery of Electrical Energy in the 21st Century*, Jul. 20–24, 2008, pp. 1–5.
- [8] U.S. Department of Energy. Electricity consumption by end use in u.s. households, 2001. Energy Information Administration, US Department of Energy. [Online]. Available: http://www.eia.doe.gov/emeu/repse/enduse/er01_us_tab1.html
- [9] U.S. Department of Energy Advanced Vehicle Testing Activity. Plug-in hybrid electric vehicles, vehicle testing reports. Idaho National Laboratory, US. [Online]. Available: <http://avt.inel.gov/phev.shtml>
- [10] S. Karnouskos and M. M. J. Tariq, "An agent-based simulation of SOA-Ready devices," in *Proc. Tenth International Conference on Computer Modeling and Simulation UKSIM 2008*, Apr. 1–3, 2008, pp. 330–335.
- [11] Y.-S. Chiou, "Deriving U.S. household energy consumption profiles from american time use survey data a bootstrap approach," in *11th International Building Performance Simulation Association Conference and Exhibition*, Glasgow, Scotland., 27-30 July 2009, pp. 151–158.
- [12] J. K. Kok, C. J. Warmer, and I. G. Kamphuis, "Powermatcher: multiagent control in the electricity infrastructure," in *AAMAS '05: Proceedings of the fourth international joint conference on Autonomous agents and multiagent systems*. New York, NY, USA: ACM, 2005, pp. 75–82.