

A System for Enabling Facility Management to Achieve Deterministic Energy Behaviour in the Smart Grid Era

Dejan Ilić, Stamatis Karnouskos, Per Goncalves Da Silva and Sarah Detzler

SAP AG, Vincenz-Prießnitz-str. 1, D-76131, Karlsruhe, Germany

Keywords: Energy Management, Smart Grid, Facility Management, Energy Forecast, Energy Storage, Energy Trading.

Abstract: The vision of the Smart Grid empowers a variety of innovative approaches for flexible energy management that fuse the business goals with the asset monitoring and control offered by the Internet of Things. The facility management domain can benefit from these advances by building upon Smart Grid energy services thereby realizing new business opportunities that make the best out of its assets. Due to the increasing integration of highly dynamic assets in future buildings, short-term deterministic behaviour is difficult. However with the availability of controlled variable storage, and futuristic services such as energy trading, errors in prediction can be absorbed internally or traded with the ultimate aim of “making the best” out of the assets and situations. The latter has the potential to enable facility managers to reach strategic objectives and potentially use assets more effectively by seizing new business opportunities. In this work we propose an architecture, describe its key components and depict in scenarios its usage with the goal of enabling facility management to take informed business decisions by following enterprise strategies as well as considering the volatility of the available energy excess or shortage.

1 MOTIVATION

The combination of deregulated energy markets and prevalence of modern Information and Communication Technologies (ICT) on the electricity infrastructure is paving the way towards the Smart Grid (European Commission, 2012; BDI, 2010). According to the Smart Grid vision (Yu et al., 2011), improved energy management may stem from the near real-time bidirectional communication between, and within, stakeholders. Today, several projects are under-way (Giordano et al., 2013) that apply innovative concepts to realize different aspects of this vision. The realisation of this vision heavily relies on the prevalence of the Internet of Things (International Telecommunication Union, 2005), which introduces intelligent networked devices (such as sensors and actuators) to everyday objects, house-hold appliances, industrial systems, etc. and leads to the fusion of the physical and virtual worlds (acatech, 2011).

Among ongoing research and development projects (Giordano et al., 2013), there are efforts towards better grid management, integration of smart-houses (Karnouskos, 2013) and smart-buildings, accommodation of intermittent energy resources including Electric Vehicles (EV), demand-

response schemes (Mathieu et al., 2011), local energy markets for business interactions (Ilić et al., 2012), etc. Through the shift towards integration of small, highly-distributed, energy production and storage capabilities, not only will new stakeholders will (European Commission, 2012), but even the current ones may assume new roles. Combining advanced information-driven services (Karnouskos et al., 2012) with these new capabilities will give rise to new infrastructures (Karnouskos, 2011) eager for innovative business opportunities.

Future on-premise capabilities, such as on-site energy generation or EV fleets (Tomić and Kempton, 2007), will provide industrial facilities with new business and management opportunities (Kanchev et al., 2011). Since a typical industrial building can be seen as an ecosystem (Carosio et al., 2013), its internal (e.g. building infrastructure) as well as the new extended components (e.g. EVs, storage etc.) can cooperate to improve energy management (Palensky and Dietrich, 2011). This in turn can enable new forms of business interaction with other stakeholders that are either currently impossible, or incur high integration costs. Of particular interest is a facility's ability to keep-up with previously planned (Blank and Lehnhoff, 2013), or forecast (Vonk et al., 2012), levels of

energy consumption and/or production, and its flexibility in adjusting to new situations while trying to minimize costs, or increase revenue for its owners (Korpaas et al., 2003). By providing a reliable prediction, such a facility could generate revenue through effective participation in, for instance, local energy markets (Goncalves Da Silva et al., 2014), or demand response programs (US DoE, 2006).

Forecasting the electricity consumption and/or production behaviour of a building will of course lead to errors (Mathieu et al., 2011) internally; however these may not need to be propagated to external stakeholders as it is done today. The challenge is on how to leverage the facility's capabilities (Teleke et al., 2009) and external interactions in order to benefit the enterprise. More specifically, how the existing and new assets that are under the control of the facility management can be empowered with Smart Grid technologies and services, and be effectively used to address any energy shortage or excess caused by the on-site prediction errors (Pinson et al., 2009).

To address this problem, we propose a system that takes advantage of existing (including temporal) assets and Smart Grid services, and enables facility management to actively adjust its energy consumption/production behaviour as seen by external stakeholders, while adhering to its internal goals and strategies. The proposed system considers a stakeholder with variable storage and energy trading capabilities, which may be the norm in the years to come. We describe several management strategies that can be realized with this system to demonstrate its capabilities. Although individual aspects may exist in ongoing research work, the proposed system combines several of them together i.e. forecasting, storage and trading, with clear applications in facility management (i.e. buildings) and with a down to earth design that may enable it to be productively used in the short-term.

2 SYSTEM ARCHITECTURE

2.1 Overview

The proposed system is modular and designed to empower the collaboration of the independently operating sub-systems, as well as the homogenization of their functionalities in a mash-up end-user application. As depicted in FMC notation (www.fmc-modeling.org) in 1, we can distinguish the interactions of the end-user via the cockpit, the main systems involved in the back-end i.e. energy load forecast (ELF), variable energy storage (VES) and energy trading (ET), as well as the reliance on external par-

ties such as an energy market or an external energy stakeholder.




Figure 1: System architecture overview.

The *Energy Load Forecast (ELF)* is the system responsible for forecasting the energy signature based on historical smart metering data (residing in the metering system) as well as real-time data acquired by the infrastructure (Karnouskos et al., 2011). Its results form the basis for the decision making process of how to handle the excess or shortage of energy predicted.

The *Variable Energy Storage (VES)* consists of managing the available “storage” in the enterprise (Ilić et al., 2013). The latter may include static as well as dynamic energy storage (such as a fleet of EVs). The VES is also envisioned to have the capability of managing processes that could store or re-feed energy, such as, the rescheduling of a process.

The *Energy Trading (ET)* is able to trade energy on smart city marketplaces, that is, intelligently buying or selling energy depending on the needs of the overall system (Goncalves Da Silva et al., 2013).

The *Energy Management (EM)* is a coordinating entity which enables the collaboration among the different sub-systems, in our case ELF, VES and ET, while in parallel taking the decisions on the actions to be enforced. Based on the enterprise goals and strategies set by the facility manager, it may dynamically decide between the portions of energy that can be “stored” in the VES or traded by ET in an electricity market.

The *Cockpit* is the user interface (UI) that the end-user, i.e. the facility manager dealing with the energy related aspects, interacts with. The cockpit is envisioned as a mash-up application depicting key aspects of the status of the underlying infrastructure, including enterprise related key performance indicators. It can depict in real-time all information related to the utilization of the storage, the energy forecast-

ing as well as the achieved energy accuracy, the energy traded and related costs, the currently available and followed energy management strategies etc. The cockpit is considered to be easily realised as a web application hosted in the cloud, easily accessible via the browser e.g. of a mobile device or laptop.

Finally, we have to note that the envisioned system can communicate with external parties and services such as an *energy market* and *external energy stakeholders* in order to expand its capabilities. This also implies the role of being part of a larger ecosystem and the capability of being easily integrated in its business processes; for instance the goals pursued by the facility management could be adjusted to reflect dynamically changing enterprise needs.

2.2 Energy Load Forecast (ELF)

Forecasting is a well known component of every energy management system. Imbalances provoked due energy load forecast errors may result in a shortage or excess of energy that must be accommodated, e.g. in form of charging or discharging a battery. The ELF requires the availability of the actual energy load $y[n]$ of a stakeholder (an interval n of size T) in the past, i.e. its smart metering data and potentially other information such as weather data, asset specific behaviour or participation in processes, etc. With the availability of $y[n]$, an interval self-forecast $\hat{y}[n]$ can be reported with minimum offset of Δ , thus always reporting $\hat{y}[n - \Delta]$ at interval n . The reporting as such can be observed in 2.




Figure 2: Forecasting on different horizons and intervals to improve the forecast accuracy.

ELF utilizes advanced forecasting algorithms that continuously provide accurate predictions $\hat{y}[n]$ over smaller horizons. Its accuracy depends on multiple conditions such as the applied forecast algorithm, the required horizon etc. In this work, ELF provides a forecast of the system for any horizon h in the future, so the continuous load forecast is done for an interval $\hat{y}[n - h]$. It is expected that many components of the system will require different horizons, as $h \approx \Delta$ might not be of interest. 2 depicts the accuracy of $\hat{y}[n]$ as

being higher than $\hat{y}[n]$ (since $h < \Delta$). The ELF configuration is expected to be done internally based on historic accuracy of the achieved performance.

2.3 Variable Energy Storage (VES)

The VES component combines both static and dynamic storage into one (virtual) unit of capacity. A static storage has constant capacity if performance degradation is not considered. In contrast, dynamic storage is composed of multiple (potentially mobile) units that are at some point in time connected to the grid (e.g. EVs). While static storage can charge or discharge in dependence to its actual state of charge (SOC), these dynamic units are the actual energy-flexible components, when of course connected to the grid. This flexibility is gained by controlling the amount of energy that they charge or discharge as well as rescheduling such activities over an interval n of length T . As it will not always be possible to compensate the exact energy needed, e.g. due to technical restrictions, on every reschedule request, the error that should be absorbed $s[n]$, is not expected to be fully addressed, but reduced to what is actually stored $s[n]$. This gap can be however addressed within the variable storage as a whole, since it is combined from its dynamic and its static part (which does not have the same temporal restrictions). A potential usage of the VES might be to use its dynamic part to compensate the closest value possible, while the static part can correct the uncompensated part of the error by charging or discharging the amount of energy needed. However, the exact usage may depend on various other technical or financial constraints, and is out of the scope of this work.

A stakeholder owned EV fleet (for which it is assumed that the facility management has full control over) is a good example of the dynamic part of a variable storage, while it is limited by scheduling and vehicle restrictions. For the rescheduling step, different priorities will need to be satisfied in order to ensure that these EVs are always within the fleet requirements. As such, any EV fleet can be used to calculate the maximal shiftable load to positive $\Delta s^+[n]$ and negative $\Delta s^-[n]$. For this calculation it has to be considered, that EVs can only vary their charging between the maximal and minimal power, or interrupt the charging completely. Within these limits, the fleet can react on energy shortage or surplus at stakeholder's premise, e.g. by interacting with an energy market or even compensating forecast errors by rescheduling or shifting loads in an interval n . Therefore in case of an energy demand change, discharging of EVs would be a secondary option, while reschedul-

ing has precedence, since no losses are made due to the storage efficiency.

2.4 Energy Trading (ET)

Local energy markets may emerge as a scalable methodology for controlling the levels of consumption and production on the grid (Ilić et al., 2012; Goncalves Da Silva et al., 2013), in particularly as a response to the increasing deployment of distributed energy resources (e.g. PV panels, wind farms, μ CHP generators, etc.). Within the proposed architecture, a energy local market is considered as an opportunity for a stakeholder not only to maintain its predictability, but to also, in some cases, better utilize and capitalize on its storage facilities. With that in mind, the ET system component interfaces with the local market to buy/sell energy by applying different trading strategies, such as (Cliff and Bruten, 2000; Vytelingum et al., 2010).

The stakeholder calculates, on an interval basis, the energy trading target $\tau[n]$ based on its internal strategies and goals. For instance, the trading targets could be based on the forecasting errors provided by the ELF. A limit price, $\tau_p[n]$, for either buying or selling is optionally set with each target to indicate the maximum (minimum) buying (selling) price for interval n . If the pricing information for a particular interval is undefined, the ET will trade aggressively on the market to ensure that the targets are met, so $\tau[n]$ presents the net quantity traded by the ET with the interval. Otherwise, each target can only be met within the bounds of the its pricing constraints.

Current targets can be updated as new information is made available to the stakeholder. In such cases, the ET updates its market position to meet the new targets. For instance, if the target is set to $\tau[n] = 50 Wh$, of which current trading is $\tau[n] = 20 Wh$, when a new target of $-30 Wh$ is received, the ET should then sell $\tau[n] = -50 Wh$ to meet the new target. The performance of the ET can be tracked by requesting the total traded quantities $\tau[n]$. Furthermore, for purposes of a cockpit (thus assistance to an operator), the ET provides interfaces to access the overall market information, as prices $p[n]$ and trading volumes per time interval.

2.5 Energy Management (EM)

The facility manager, as illustrated in 1, interacts with the system via a cockpit. An example of such a cockpit and information it offers is depicted in 3. The facility manager can consume the (real-time) information depicted and by calibrating or setting the over-

all goals can exercise high-level control over the infrastructure. Such goals could be the optimization of the infrastructure reaction to the energy surplus or shortage reported by ELF towards economic objectives such as minimization of cost, or other corporate social responsibility related ones, e.g. maximization of usage of green electricity or even simpler ones such as making sure that the EVs of the employees are fully-charged by the end of their workday.




Figure 3: The envisioned system cockpit as a mash-up application.

The transformation of user goals (calibrated via the cockpit) to strategies are processed by the EM, which takes into consideration all other constraints of the system and takes the overall decision on the appropriate strategies to be followed. EM acts also as a communication broker among the different parts of the system as it holds the system-wide knowledge that is not available to the individual parts i.e. the ELF, ET, VES, enabling the latter scaling or extension of the system with other components or variations of the existing ones.

EM acts as the coordinator and decision engine, which communicates with ELF, ET and VES, and provides them with the operational context info. As an example, in a scenario where the EM is informed about the energy surplus available due to a forecast error, it may decide to redirect part of it towards charging the EVs while another part may be redirected to the ET (by charging schedule adjustment) in order to be traded to the market (because the price is high or can not be covered wholly by the VES). More example scenarios will be depicted in 3.

3 ENERGY MANAGEMENT STRATEGIES

3.1 Overview

The system proposed, whose main components are illustrated in 1, is flexible enough to accommodate several envisioned scenarios, depending on the goals set by the user, the available at time capabilities, and actions to be enforced. The scenarios we will focus upon, are in no way exhaustive, but serve to provide some understanding of the potential strategies that could be followed by the facility management. Our aim is to showcase the system’s flexibility, which is a key part of realizing agile enterprises in the future.




Figure 4: General view of the activity involving the architecture components.

A general view of the workflow is depicted in 4. The user input is acquired, which together with the forecast error and the underlying status and constraints of the subsystems, are used to reach a decision for either trading or storing of energy (or both). Some key strategies will be discussed in more detail i.e. 3.2, 3.3, and 3.4.

Generally, each envisioned strategy may not involve all parts of the system, as this depends on the actual constraints imposed at the time of the decision making. This also signals that an organization does not have to wait until all of the architecture parts are deployed and become operational to start realising (a

limited set of) energy management strategies. As an example, the ELF and the VES could be realized today, while the ET could be realized some years later when energy markets are available at smart city level and it makes economic sense for the facility managers to participate in them. Hence, the system architecture accommodates the “migration” i.e. incremental evolution of the infrastructure towards the fully-fledged Smart Grid vision.

3.2 Storing Energy not Traded by ET

The decision making process (as depicted in 4) may consider a strategy that is described as follows: *after the estimation of the energy error within an interval by the ELF, try to trade the difference via the ET and differ any non-traded energy to the VES for storage.* The workflow of such strategy is illustrated in 5. ET accommodated $\tau[n]$ for the interval in question (potentially even at different prices p), and VES is contacted in order to absorb the remaining $s[n]$.




Figure 5: Strategy activity: 3.2.

The actual outcome of the trading done by ET depends on the real-market conditions (law of demand and supply) and hence strategy adaptation might be needed over time e.g. acting alone or as part of a larger group (Goncalves Da Silva et al., 2014). As the ET might not be able to fully trade the energy needed to balance the forecast error $\tau[n] = y[n] - \hat{y}[n]$, a part of it remained non-traded. The traded quantity $\tau[n]$ is then communicated back to the EM, which instructs VES to accommodate the remaining $s[n] = y[n] - y[n] + \tau[n]$ energy. This process leads to a new

state where the error is minimized as a “best effort” procedure is followed by ET (interaction with external stakeholders) and VES (internal stakeholder) to minimize its impact.

Depending on the business motivation, this strategy may be followed when the return of investment (ROI) by selling the energy on the marketplace is high. This may be a result of high prices on the energy market, inability or no need of storing the energy internally, etc. The actual decision-making process will be dynamic and the exact fine-tuning is not considered here.

3.3 Trading Energy not Accommodated by VES

In compliance with the decision making process depicted in 4, here we focus on a strategy that can be described as follows: *after calculation of the energy due to the incurring error by ELF, try to accommodate the excess or shortage of energy via the VES (by charging/discharging) and for the remaining part not accommodated by the VES, use the ET.* In this strategy, the ET acts as a mitigating agent for any part of the error that could not be absorbed by the VES, which is shown in 6. In detail, the EM acquires the forecasting error ($\hat{s} = \hat{y} - \hat{y}$) from the ELF and informs the VES, which attempts to accommodate the imbalances introduced by the errors, and informs the EM of any amount that could not be accommodated due to its internal constraints (s). These amounts are then given to the ET to be mitigated on the market ($\hat{\tau} = \hat{y} - y + s$).




Figure 6: Strategy activity: 3.3.

This strategy is expected to be used when the enterprise has the capability to store energy extensively for its own use. For instance a significant number of EVs at the disposal of facility management, means that the VES can rely on storing energy there and acquiring it back again when needed. Even if the energy

is not needed during the day for tackling imbalances, the EVs are charged and the energy can be used for the enterprise’s processes in the future; an action that enhances better planning of energy-relevant actions. If a local marketplace is available, the ET tries to trade the remaining energy (or acquire in case of shortage) in order to make ends meet.

3.4 Trading Available VES Capacity

Independent of the strategies followed by the facility management, the available assets might not be fully utilized. Hence additional actions can be run, in compliance to the decision making process as depicted in 4 and in parallel to the existing strategies. As an example, consider a strategy that is described as follows: *check the available of VES flexibility and trade the additional shiftable energy with the help of the ET in the local energy market.*

This example goes beyond the traditional process of trying to cover the energy imbalances and targets clearly the optimization area of the available assets. The VES might be underutilized while trying to cover the occurring imbalances, and that “unused” capacity could be transformed into economic benefit for the company. The VES may have its own models for estimating capacity that will not be used, and hence can “release” it for further usage. Then the EM may consider this capacity and buy or sell energy with the help of ET in order to generate additional revenue for the company. The system is flexible enough to accommodate such actions. However, to avoid conflicts or side-effects, additional analysis on the resource utilization is needed which is not part of the investigation presented here.

Since the VES is trying to compensate the error produced by the ELF, a certain state of charge (SOC) will be achieved. Based on the actual flexibility levels with consideration of SOC, the VES can offer a certain capacity within an interval n for charging/discharging in order to increase the enterprise’s revenue. Instead of only offering the capacity which is left from the error compensation, the VES may calculate the maximal and minimal shiftable energy $\Delta s^+[n]$ and $\Delta s^- [n]$. This offered capacity can be then traded and benefit from current price p . In another twist, the VES may decide that the economic benefits of trading the capacity are greater than that of being used a storage and act accordingly, which is blending the borders of the other two strategies discussed.




Figure 7: Strategy activity: 3.4.

4 DISCUSSION

Several considerations on the main components i.e. ELF, VES, ET need to be adequately addressed. Forecasting done by the ELF, cannot only be based solely on historical data, but needs to include real-time information. To this end, the Internet of Things coupled with the Cloud (Curry et al., 2013; Karnouskos, 2013) and the vast resources for analytics will help. Additionally, more specific knowledge of the processes involved, their scheduling at enterprise level, as well as their potential interdependencies may lead to better forecasting and planning. For our environment this assumes access to the building processes per se, the impact of the employee usage of its facilities (which extends to the EV usage), as well as potentially external factors such as weather conditions etc. The appropriate combination of such intelligent algorithms with (real-time) fine-grained data may enable the better adjustment of the infrastructure behaviour prediction.

Another key part of the system, the VES, demonstrates that the temporal storage availability e.g. coming from an EV fleet, can be used to acquire additional benefits for the enterprise. Although charging/discharging EV batteries or even rescheduling their charging may sound promising, at the moment few, if any, companies have fleets large enough for these envisioned concepts to be practically applicable. Furthermore, the communication protocols of the EVs or the charging stations might also pose challenges as they might not allow a secondary actor to

set the charging schedule for them and always try to charge as fast as possible.

Such constraints would tremendously lower the possibility of interaction with other components in a smart way. The latter gives some indication that company-controlled fleets are the right target group for such concepts as the one presented here. Even then, due to constraints (physical or otherwise), fleets might not always provide the exact flexibility that is needed. The latter can be mitigated through the addition of static storage, or another buffer-like component, that can compensate the missing flexibility.

The ET demonstrates how interacting with other local stakeholders can not only aid the facility in reducing its forecast error, but also create additional opportunities through energy and storage capacity trading. Although local energy markets are a hot topic (Ilić et al., 2012; Lamparter et al., 2010; Ding et al., 2013) in Smart Grid research, how to effectively make use and interact with them is still unclear. As no such market currently exists, operational assumptions were made; however in a real-world assessment the underlying trading behaviour must be anchored in a clear understanding of the market's rules and protocols. Additionally, in order for the ET to meet the a wide range of strategies, such as the ones described in 3, it must be able to adequately handle dynamically changing trading goals in conjunction with market-forecasts and enterprise's needs.

Generally, we consider that there is added value if such systems are operational and would assist towards informed and automated decision-making processes in facility management domain. Their realization however, will need to be assessed and fine-tuned in real-world trials once the necessary Smart Grid services envisioned are mainstream.

5 CONCLUSION

The ability to capitalize on new business opportunities is vital for the survival of modern enterprises. To that extent, fully utilizing all the capabilities offered by its assets is pivotal. In the Smart Grid era, the facility management can take sophisticated decisions related to energy management, by including innovative technologies and concepts. The system we have presented builds upon the orchestration (by the EM) of three key independent components (i.e. ELF, VES, ET). We have provided insights on its basic components, their operation as well as their usage in a variety of strategies. In the latter we have also discussed the benefits for the enterprise as well as the roadblocks that might hinder their realization. Ultimately, the proposed sys-

tem enables on-the-fly decision-making that empowers the facility managers to better meet their goals in the emerging Smart Grid era.

ACKNOWLEDGEMENTS

The authors would like to thank the European Commission, the German Federal Ministry of Economics and Technology, and the partners of the projects SmartKYE (www.SmartKYE.eu) and iZEUS (www.izeus.de) for their support and fruitful discussions.

REFERENCES

- acatech (2011). Cyber-Physical Systems: Driving force for innovation in mobility, health, energy and production. Technical report, acatech – National Academy of Science and Engineering.
- BDI (2010). Internet of Energy: ICT for energy markets of the future. Federation of German Industries (BDI) publication No. 439.
- Blank, M. and Lehnhoff, S. (2013). Assessing reliability of distributed units with respect to the provision of ancillary services. In *IEEE 11th International Conference on Industrial Informatics (INDIN)*, Bochum, Germany.
- Carosio, S., Hannus, M., Mastrodonato, C., Delponte, E., Cavallaro, A., Cricchio, F., Karnouskos, S., Pereira-Carlos, J., Rodriguez, C. B., Nilsson, O., Seppä, I. P., Sasin, T., Zach, J., van Beurden, H., and Anderson, T. (2013). *ICT Roadmap for Energy-Efficient Buildings – Research and Actions*. EU Project ICT4E2B Consortium.
- Cliff, D. and Bruten, J. (2000). Less than human: Simple adaptive trading agents for CDa markets. In Holly, S., editor, *Computation in Economics, Finance and Engineering: Economic Systems*, pages 117–122, Oxford. Pergamon.
- Curry, E., O’Donnell, J., Corry, E., Hasan, S., Keane, M., and O’Riain, S. (2013). Linking building data in the cloud: Integrating cross-domain building data using linked data. *Advanced Engineering Informatics*, 27(2):206 – 219.
- Ding, Y., Alexander Neumann, M., Goncalves da Silva, P., Beigl, M., and Zhang, L. (2013). A control loop approach for integrating the future decentralized power markets and grids. In *IEEE International Conference on Smart Grid Communications (SmartGridComm)*.
- European Commission (2012). SmartGrids SRA 2035 – Strategic Research Agenda. Technical report, European Technology Platform SmartGrids, European Commission.
- Giordano, V., Meletiou, A., Covrig, C. F., Mengolini, A., Ardelean, M., Fulli, G., Jiménez, M. S., and Filiou, C. (2013). Smart Grid projects in Europe: Lessons learned and current developments 2012 update. Joint Research Center of the European Commission, JRC79219.
- Goncalves Da Silva, P., Ilic, D., and Karnouskos, S. (2014). The impact of smart grid prosumer grouping on forecasting accuracy and its benefits for local electricity market trading. *IEEE Transactions on Smart Grid*, 5(1):402–410.
- Goncalves Da Silva, P., Karnouskos, S., and Ilic, D. (2013). Evaluation of the scalability of an energy market for smart grid neighbourhoods. In *IEEE 11th International Conference on Industrial Informatics (INDIN)*, Bochum, Germany.
- Ilić, D., Goncalves Da Silva, P., Karnouskos, S., and Griese-mer, M. (2012). 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.
- Ilić, D., Karnouskos, S., and Goncalves Da Silva, P. (2013). Improving load forecast in prosumer clusters by varying energy storage size. In *IEEE Grenoble PowerTech 2013*, Grenoble, France.
- International Telecommunication Union (2005). ITU Internet Report 2005: The Internet of Things. Technical report, (ITU).
- Kanchev, H., Lu, D., Colas, F., Lazarov, V., and Francois, B. (2011). Energy management and operational planning of a microgrid with a PV-based active generator for smart grid applications. *Industrial Electronics, IEEE Transactions on*, 58(10):4583–4592.
- Karnouskos, S. (2011). 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.
- Karnouskos, S. (2013). Smart Houses in the Smart Grid and the search for value-added services in the Cloud of Things era. In *Industrial Technology (ICIT), 2013 IEEE International Conference on*, pages 2016–2021.
- Karnouskos, S., Goncalves Da Silva, P., and Ilić, D. (2011). Assessment of high-performance smart metering for the web service enabled smart grid. In *Second ACM/SPEC International Conference on Performance Engineering (ICPE’11)*, Karlsruhe, Germany.
- Karnouskos, S., Goncalves Da Silva, P., and Ilić, D. (2012). Energy services for the smart grid city. In *6th IEEE International Conference on Digital Ecosystem Technologies – Complex Environment Engineering (IEEE DEST-CEE)*, Campione d’Italia, Italy.
- Korpaas, M., Holen, A. T., and Hildrum, R. (2003). Operation and sizing of energy storage for wind power plants in a market system. *International Journal of Electrical Power & Energy Systems*, 25(8):599–606. 14th Power Systems Computation Conference.
- Lamparter, S., Becher, S., and Fischer, J.-G. (2010). An agent-based market platform for smart grids. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: Industry track*, AAMAS ’10, pages 1689–1696, Richland, SC.

- International Foundation for Autonomous Agents and Multiagent Systems.
- Mathieu, J., Price, P., Kiliccote, S., and Piette, M. (2011). Quantifying changes in building electricity use, with application to demand response. *Smart Grid, IEEE Transactions on*, 2(3):507–518.
- Palensky, P. and Dietrich, D. (2011). Demand side management: Demand response, intelligent energy systems, and smart loads. *Industrial Informatics, IEEE Transactions on*, 7(3):381–388.
- Pinson, P., Papaefthymiou, G., Klockl, B., and Verboomen, J. (2009). Dynamic sizing of energy storage for hedging wind power forecast uncertainty. In *Power Energy Society General Meeting, 2009. PES '09. IEEE*, pages 1–8.
- Teleke, S., Baran, M., Huang, A., Bhattacharya, S., and Anderson, L. (2009). Control strategies for battery energy storage for wind farm dispatching. *Energy Conversion, IEEE Transactions on*, 24(3):725–732.
- Tomić, J. and Kempton, W. (2007). Using fleets of electric-drive vehicles for grid support. *Journal of Power Sources*, 168(2):459–468.
- US DoE (2006). Benefits of demand response in electricity markets and recommendations for achieving them. Technical report, US Department of Energy.
- Vonk, B. M. J., Nguyen, P., Grond, M., Sloopweg, J., and Kling, W. (2012). Improving short-term load forecasting for a local energy storage system. In *Universities Power Engineering Conference (UPEC), 2012 47th International*, pages 1–6.
- Vytelingum, P., Ramchurn, S. D., Voice, T., Rogers, A., and Jennings, N. R. (2010). Trading agents for the smart electricity grid. In *9th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2010), Toronto, Canada*, pages 897–904.
- Yu, X., Cecati, C., Dillon, T., and Simões, M. (2011). The new frontier of smart grids. *Industrial Electronics Magazine, IEEE*, 5(3):49–63.