

DEMAND SIDE MANAGEMENT IN A FIELD TEST: LESSONS LEARNED

Marco E. T. GERARDS

University of Twente – the Netherlands
m.e.t.gerards@utwente.nl

Johann L. HURINK

University of Twente – the Netherlands
j.l.hurink@utwente.nl

Roland HÜBNER

Innogy SE - Germany
roland.huebner@innogy.com

ABSTRACT

The SmartOperator project of RWE Deutschland AG combines multiple techniques to address power quality issues within the distribution grid: storage, tap changers and Demand Side Management (DSM). This paper focusses on the domestic DSM aspects of the project, which controls appliances such as batteries, whitegoods and EVs within houses. The DSM methodology operates in three steps: prediction, planning and online control. The main lessons learned during the course of the project are presented in this paper.

INTRODUCTION

Germany faces high PV penetration due to subsidies that stimulate renewable energy. The generation of PV is concentrated around noon and often unmatched by consumption, leading to high synchronized production peaks and degraded power quality.

Since the feed-in tariffs for PV surplus are often not very attractive, self-consumption is stimulated and devices such as electric vehicles (EVs) and heat pumps become increasingly popular. However, when such devices are used at times without PV production, this leads to high (often synchronized) consumption peaks.

To deal with the two above mentioned aspects, storage, tap changers, and Demand Side Management (DSM) [1,2,3] are proposed in the literature. Within the SmartOperator project [2,3] the interaction of these techniques is studied whereby a central controller, the *SmartOperator*, sends steering signals to the storage, tap changers and DSM capable houses. The developed system is already running for almost two years in a German village and is currently in its final evaluation stage. This paper presents the main lessons learned from the perspective of the DSM part of the system.

The remainder of the paper is organized as follows. The next section discusses the DSM approach called “Triana”, which is used within the SmartOperator project. The subsequent sections discuss seven lessons learned within the project. These lessons can help developers of new DSM systems to ease development and evaluation of their system. The paper ends with a short summary of the results and conclusions.

DSM APPROACH

The DSM part of the project consists of a centralized controller that coordinates the decisions on neighborhood

level and a Home Energy Controller (HEC). The HEC is based on Innogy SmartHome [9], to which the Triana DSM software is added. Triana steers the smart appliances (e.g., batteries, white goods, heat pumps, and EV chargers) depending on local PV production, the baseload of the house and steering signals from the SmartOperator.

To influence Triana, the SmartOperator sends five alternative steering signals to Triana. The first four steering signals are artificial energy price profiles (24 hours ahead, 15-minute resolution) that are used to influence the house load profile to resolve power quality issues within the neighborhood. The fifth steering signal instructs the HEC to be ‘selfish’, which is translated to maximize the self-consumption of the house and is the default mode of operation. Triana uses three steps to deploy the flexibility inside the house:

1. *Prediction*: Triana predicts loads and available flexibility 24 hours ahead.
2. *Planning*: Based on the predictions, five different plans (energy profiles) for the next 24 hours are made (one for each steering signal). Internally, Triana uses a resolution of 5 minutes. The plans are translated to 15 minutes and sent to the SmartOperator, which chooses the plan that best reflects the grid needs.
3. *Online control*: Triana controls the appliances based on the selected plan.

In this planning-based approach, the HEC makes five different promises in the form of five energy profiles to the SmartOperator. To make such profiles, obtaining good predictions and plans on house level is the main challenge.

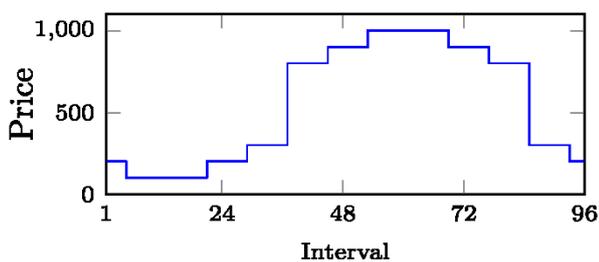
The next sections discuss the lessons learned from implementing and evaluating this DSM methodology. To illustrate the lessons learned, the battery is used as a running example.

RESPONSE TO PRICE SIGNALS

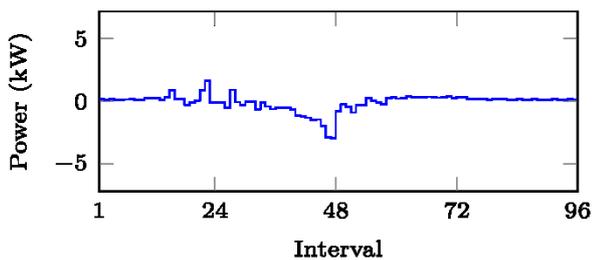
Artificial steering signals in the form of prices that vary over time are used to steer the appliances in the house in a certain direction. This is illustrated for the price profile in Figure 1a, which is used to influence a battery within a certain house of the SmartOperator project. The base load profile of this house is given by Figure 1b and Figure 1c shows the greedy response of a battery to these prices, i.e., the battery is charged at maximum power when the price is low and discharged at maximum power when the price is high. As this charging and discharging is done at maximum power only based on the price profile without

considering the remaining load profile of the house, large peaks in both directions may be added instead of preventing them.

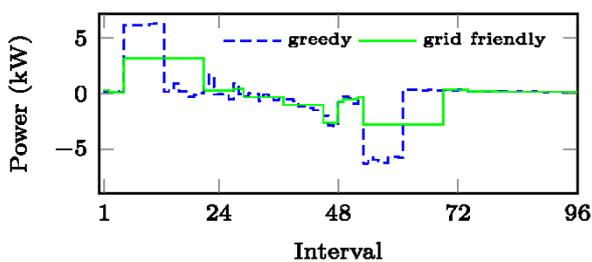
However, from a pricing perspective, it is in general not important *how* charging is distributed over consecutive intervals with the same price, as all such patterns lead to the same costs. The (solid) green graph in Figure 1c shows a (dis)charging pattern that still charges/discharges when prices are low/high, but in a way that it also minimizes the peak loads of the house. The aim of our DSM within the SmartOperator project was to create such behavior; i.e., to focus on price signals but also to take into account the resulting profiles.



(a) Artificial prices



(b) Measured house load (from project)



(c) Battery charging

Figure 1: Price-driven battery charging

To achieve this, we minimize the following objective function to steer the devices:

$$\sum_{t=1}^{96} p_t a_t + \epsilon (p_t + P_t)^2, \quad (1)$$

where p_1, \dots, p_{96} denotes the planned energy consumption profile of the steerable devices for the 96 upcoming 15-minute intervals (i.e., one day), a_1, \dots, a_{96} the

corresponding prices and P_1, \dots, P_{96} the given (predicted) house profile. The term $p_t a_t$ expresses the price related costs of the steerable device for interval t . When only the sum of these terms is minimized, this could lead to the dashed blue graph in Figure 1c. The second term, $\epsilon (p_t + P_t)^2$, takes the house load P_t into account and assigns higher costs to higher loads to achieve load flattening. We use an (extremely) small value ϵ to achieve that prices are optimized first and then peaks. Hence, when the price over several intervals is the same, this term ensures that the total load (house+battery) is spread evenly over these intervals. This leads to a desired situation as depicted in Figure 1c (solid green). We used an efficient and tailored algorithm [5] to solve this optimization problem within milliseconds.

This concept of integrating peak shaving in price optimization is also applied within several other device scheduling algorithms within the project, and leads to the first lesson learned.

Lesson 1: High load peaks resulting from time-based pricing can be mitigated by adding the following term to the objective function of the planning step: $\sum_t \epsilon (p_t + P_t)^2$, where ϵ is some very small constant.

UNCERTAINTY LEADS TO CONFLICTING OBJECTIVES

Two conflicting objectives are encountered within the SmartOperator project, which need to be addressed:

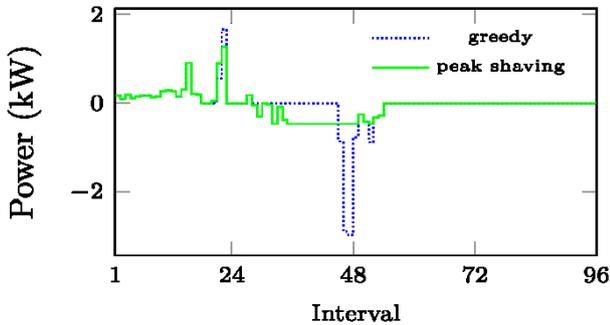
- DSO: Unlock flexibility and avoid peaks.
- Customer: Avoid discomfort and energy losses and increase self-consumption of energy.

Whereas with perfect information both objectives might be served, uncertainty may disrupt one of the objectives. This is best illustrated with battery charging, for which we simplify (without loss of generality) the DSO objective to peak shaving to ease the discussion.

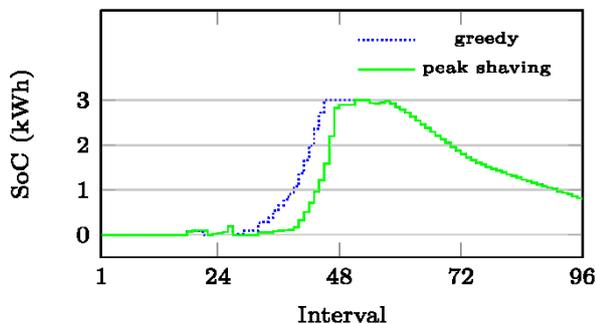
Commercially available battery systems for domestic use aim to maximize self-consumption and commonly do this by storing PV surplus and discharging when PV is insufficient to fulfil the demand. However, when the battery capacity is not large enough to store all PV production surplus, this may lead for the house profile of Figure 1a to the situation as depicted by the dotted blue graph in Figure 2a (assuming that the battery was empty at the start): a steep PV peak occurs at the moment the battery is full (around time interval 48) and stops charging.

In *theory* a plan for charging the battery can be calculated, which maximizes self-consumption in a way that such peaks are prevented, shown by the solid green graph in Figure 2a. This is achieved by charging less around interval 30 to have some room for charging around interval 48 (see also Figure 2b). This assumes that we should already know around interval 30 that there will be a PV surplus around interval 48. However, when this prediction is not accurate and it turns out that there is no PV around interval 48, the battery was not used optimally

for self-consumption since around interval 30 some surplus was fed to the grid. This means that reserving battery capacity for grid use may lead to higher costs for customers as self-consumption is economically beneficial for them.



(a) Battery power



(b) Battery SoC (in kWh)

Figure 2: Battery charging (3kWh) for self-consumption

Correspondingly, when the PV surplus around interval 48 is *higher* than was predicted around interval 30, there is no capacity to buffer the peak and the DSO objective is still not met. This suggests a duality: due to uncertainty either the customer or the DSO suffers. The DSO may resolve this by compensating the customer for its risk when it reserves capacity for a predicted PV peak, which may not occur.

Lesson 2: Uncertainty in predictions may lead to unavoidable conflicting objectives. Reserving capacity for an error margin results in an undesirable result for at least one stakeholder and therefore compensations should be introduced.

CONTROL OF APPLIANCES

During the control step of our DSM methodology, devices are controlled to realize the plans made during the planning step. However, the plans should not be directly translated to control actions, since the situation after the planning may have changed. Examples of such changes may be:

- The measured house load deviates from the prediction, and therefore (dis)charging the

battery according to the plan may not be the best decision given the actual measurements.

- It was planned to discharge the battery because of a predicted consumption peak but, due to weather changes, there is now a PV peak instead and thus it would not be good to discharge.

Instead of directly following the plan, we propose to discharge according to the desired optimization outcome that matches the given objective (e.g., load flattening). This means that we charge the battery such that the total house load becomes equal to the value $p_t + P_t$ of the planning. This means that we adapt, if possible, the value of p_t to changes of the value P_t . Note that $p_t + P_t$ is in general a constant level (referred to as “fill level”) for several consecutive intervals (see Figure 1c, in solid green) when prices are constant over these intervals.

The above concept may be implemented by constantly adapting the (dis)charging power to the measured house load. However we suggest a different, more engineering oriented, approach that asks some adaptation by the battery manufacturers. Nowadays many commercially available batteries allow charging when there is PV surplus and discharging when there is net demand. This means that the battery (dis)charges toward a total house load (incl. battery) of 0 W. When the manufacturers would make this level of 0 W configurable instead, the HEM just has to set it to the planned level $p_t + P_t$, meaning that continuous steering is no longer needed.

It has been shown that for other storage based devices, such as, e.g., electric vehicles [6], the same fill-level based approach can be used.

Lesson 3: Control is simplified and more robust when buffer-based devices can control their operation based on a given desired total house load (incl. the device itself), called a fill-level approach.

INFLUENCE OF PREDICTIONS

To determine when and how to deploy the flexibility of appliances inside a house and to be able to submit possible load profiles over time to the smart operator, predictions are needed and used. The first was already demonstrated in the previous sections, where predictions of $P_{1,\dots,P_{96}}$ of the base load of a house were used in the planning phase to flatten the overall load profile. However, predicting the base load profile is very difficult, since it is hard to predict if and when events occur (e.g., if and when the television is turned on). In other words, it implies predicting (the exact timing) of human behavior.

In the following we suggest a method that avoids to some extent the above mentioned difficulty. The method is based on the following observation. For optimal battery charging, the charging over some consecutive intervals mostly depends only on the battery capacity and the total house consumption in these intervals. This suggests that small changes to the house profile P_1, \dots, P_{96} , e.g., exchanging values for two consecutive intervals, does not

change the fill level (and thus the planned overall profile).

The above implies that the exact values beyond the first (i.e., P_2, \dots, P_{96}) are not crucial, but only their total energy volume is. Thus, we aim at a prediction wherein this volume is reasonably accurate. The first time interval, however, needs to be accurate since it is directly used. But this value is better predictable since it is about to commence. This concept does not only apply to batteries, but to many devices of which the consumed+produced energy volume is almost independent of control actions.

Lesson 4: When both the predictions for the first time interval and for the overall energy volume of the remaining intervals are accurate, near-optimal decisions can be made.

We use two prediction strategies, namely for the short term (i.e., for P_1) and the long term (i.e., for P_2, \dots, P_{96}). For the long term we take a linear combination of the prediction from the week before (or zeros when unavailable) and the measurements from the week before. In terms of energy volume this gives quite reasonable results. For the prediction of the first interval, we take a linear combination of the current measurement (reflecting current behavior) and the long term prediction for this interval.

UNSYNCHRONIZED MEASUREMENTS

In DSM, measurements from several sources, e.g., by smart meters, PV inverters, or flexible devices are used as input for decision making. Within the project, the smart meters measure every three minutes, the PV inverters every ten seconds, and the controllable devices use other measurement periods or report events. These measurement periods are unaligned and have different periods, which cannot be influenced due to different vendors.

When combining measurements all these variations need to be taken into account. By smoothening/averaging data over time before combining them, strong reactions to point measurements may be avoided to some extent. Nevertheless, aligning and increasing the measurement periods and frequencies may help to better reach this goal. In some cases measurements of the same quantity are available from different sources; e.g., the PV production may be measured at different rates by the PV inverter, the smart meter and the battery system. The DSM system needs to be able to combine the data from these sources. Rapid synchronized measurements and fast control are needed. Failing to accomplish the latter was partially the reason why a DSM system in another field test could not prevent a black out [4].

There is a strong link to “Lesson 2”: device controllers should synchronize measurements from different sources and control the device based on a fill level.

Lesson 5: Good, frequent and synchronized measurements and control are needed.

IMPLEMENTATION EFFICIENCY

In the SmartOperator project, Triana creates five new plans every 5 minutes. For the calculation of each such plan, we use an iterative algorithm [8] that internally invokes the device planning algorithms many times. This implies that efficient methods are needed and the data volume produced may get quite large.

To obtain such efficient implementations, we do not rely on standard solvers (e.g., ILP or convex program solvers), but we developed tailored algorithms that scale well (i.e., low computational complexity), are easy to implement and have a low computational footprint (i.e., low “constant factor”). For the battery implementation we used an algorithm that scales quadratically with the input size [5]. For devices with discrete (on/off) modes of operation we designed a tailored algorithm based on dynamic programming, which also scales quadratically with the input size [7]. These device planning algorithms can find the solution to problems with quadratic objectives, such as (1), and nontrivial constraints within a few milliseconds on the embedded system used.

The data volumes place, especially in the test phase, where a lot data needs to be stored for evaluation, a high burden on the system. To ease evaluation we store HEC data and decisions within a central MS SQL server. Despite the load on the system, this helped us at evaluating and improving the running system rapidly.

Lesson 6: Proper setup of data handling and communication is needed in combination with very efficient algorithms. For early evaluation and improvements in field tests, centralized data collection provides extreme benefits at cost of efficiency.

ROBUSTNESS AGAINST MISSING DATA

The system should be robustness against unavailability of the communication infrastructure and against missing data from the past. To achieve this, the system should be self-learning and should adapt to changing circumstances.

In the prediction of the house load, this is accomplished as follows. Since the prediction combines the prediction from the week before with the measured load from that week, old data remains present in the prediction. This implies that on the one side all old data remains to be used, although with decreasing weight, and on the other hand that the systems is able to adapt quite fast to changing behavioral patterns if the weight for the measured load of the last week is not chosen too small.

Without precautions, missing historical smart meter measurements or other data may have high impact on predictions in the following weeks. Therefore, when data is (partially) missing, we use subsequently data from the previous day or a default data set as fall back to fill in the missing data.

When communication with the smart operator malfunctions the HEM falls back to its default mode of

operation and default settings after some time (seven hours in our implementation). In the default mode, profile five (optimization of house self-consumption) is followed in the online control, and a default set of price profiles serve as base for future flexibility offers.

Lesson 7: DSM systems must be designed with the possibility of missing data, changing circumstances and down time of (part of) the system in mind.

SUMMARY AND CONCLUSIONS

This paper described the field test implementation of a DSM approach that is based on three steps: (load) prediction, (flexibility) planning and online control. To explain the used concepts and findings, a battery was used as running example, since this device offers the most flexibility and the achieved conclusions generalize to some extent also to other devices.

The DSM approach works in its base with (linear) price-based (artificial) steering signals. Since the response to prices may be extreme, we proposed to add a component to the objective that aims to spread the total load evenly over time intervals for which the price is the same.

Furthermore, we presented a way to accomplish this additional component in the planning step. The resulting plan at a house level consists of consecutive time intervals wherein the battery and other devices steer the total house load toward a certain *fill level*. This idea of using a fill level can be exploited in the design of future device controllers to ease control. With synchronization of measurement and control, additional robustness can be achieved.

Although the DSM methodology depends heavily on predictions, we argued that only the prediction for the upcoming interval needs to be rather precise and that it is sufficient to predict the overall energy volume of the remaining intervals to support the decision making process. When these values are correctly predicted, the DSM methodology becomes increasingly robust against prediction errors. Nevertheless, it remains unavoidable that uncertainty leads in some situations to conflicting interests, e.g., either higher peaks occur or less self-consumption is achieved. We argued that taking into account these uncertainties already in the planning phase by using error margin for the prediction, implies that one of the conflicting objectives suffers.

Finally, we argued that next to robustness against prediction errors, the system must also be designed with robustness against missing data and efficiency of the implementation in mind.

Acknowledgments

Part of this research has been supported by ‘Deutsche Bundesstiftung Umwelt’ within the project ‘Neubau einer Kindertagesstätte als Energiepuffer für regenerative Energien’.

REFERENCES

- [1] Nykamp, Stefan, et al. "The project" ElChe" Wetzlingen: storage as an alternative to grid reinforcements-experiences, benefits and challenges from a DSO point of." International ETG Congress 2015; Die Energiewende-Blueprints for the new energy age; Proceedings of. VDE, 2015.
- [2] Willing, Stefan, et al. "Improving quality of supply and usage of assets in distribution grids by introducing a "Smart Operator"." Electricity Distribution (CIRED 2013), 22nd International Conference and Exhibition on. IET, 2013.
- [3] Bakker, et al., "Controlling and optimizing of energy streams in local buildings in a field test," Electricity Distribution (CIRED 2013), 22nd International Conference and Exhibition on. IET, 2013.
- [4] G. Hoogsteen, et al., "Charging electric vehicles, baking pizzas and melting a fuse in Lochem", Proceeding CIRED 2016 conference, under review.
- [5] T. van der Klauw, et al. "Resource allocation problems in decentralized energy management", OR Spectrum, under review.
- [6] M. E. T. Gerards, et al. "Robust peak-shaving for a neighborhood with electric vehicles", Energies, vol. 9 (8), article 594.
- [7] M. E. T. Gerards, et al., "Planning of on/off devices with minimum run-times," in proceedings of ISGT-Europe (2016).
- [8] Molderink et al., "Management and control of domestic smart grid technology," IEEE Transactions on Smart Grid, vol. 1 (2), pp. 109-119.
- [9] Innogy SE, "www.rwe-smarthome.de", retrieved: 16-01-2017.