Preprints of the 19th World Congress
The International Federation of Automatic Control
Cape Town, South Africa. August 24-29, 2014

Scenario-based MPC for Energy Schedule Compliance with Demand Response ⋆

Theodor Borsche, Frauke Oldewurtel and Göran Andersson

ETH Zurich, Power Systems Laboratory
Physikstrasse 3, 8092 Zurich, Switzerland
borsche / oldewurtel / andersson @ eeh.ee.ethz.ch

Abstract: Demand Response (DR) is one way of providing more flexibility to the electric power system. Various approaches to enable control of small, thermostatically controlled loads such as air conditioners and Electric Water Heaters (EWHs) have been published in the past. This paper focuses on the optimal dispatch of these flexible loads. A Model Predictive Control is implemented to minimize total cost arising from both energy acquisition and schedule deviations. To handle uncertainty, scenarios of the future demand are considered. Results show the trade-off between using flexibility as a service and accrued energy cost. Detailed sensitivity analyses give insight on the dependency of DR utilization on load parameters.

Keywords: Energy management systems; Power systems; Predictive control; Stochastic programming.

1. INTRODUCTION

In electric power systems, some flexibility has to be available at all times to ensure the balance between production and consumption. Traditionally, this flexibility is provided by the production side. Power plants providing ancillary services such as frequency control are remunerated for the offered control power or energy. In liberalized power markets, the resulting cost is handed to those accountable for the balance deviation. In the European system, Balance Groups (BGs) are responsible for buying energy for their customers and submitting a day-ahead load schedule. All consumers and producers belong to such a balance group.

Recently, it has been much discussed how Demand Response (DR), that is an automated adjustment of consumption, can be used to either provide ancillary services, see (Kondoh et al., 2011), or support the schedule compliance of a BG, see (Vrettos et al., 2013). Different loads could be employed in such DR schemes, from industrial consumers to small residential appliances. Either price signals or direct control signals can be used for achieving the desired response. In our research, we focus on direct control of large aggregations of Thermostatically Controlled Loads (TCLs), such as Electric Water Heaters (EWHs), refrigerators or air-conditioners. While it was shown by (Callaway, 2011) that such an aggregation 1) can very closely track a reference trajectory and 2) due to the distributed topology is inherently reliable, the earnings or savings that can be achieved per participant are rather low. In order to make DR with household appliances economically viable, both investment cost and earnings have to be optimized. To reduce necessary investments, the infrastructure to control such loads has to be as simple and cost effective as possible. This can be achieved by using broadcasts, as in (Koch et al., 2011), and state estimation, confer (Mathieu et al., 2013) and (Borsche et al., 2013b).

To further increase profitability, the dispatch of the DR scheme should also be optimized. We differentiate between *day-ahead DR* and *intra-day DR*. In a day-ahead approach, flexible loads are shifted to low cost hours in a day-ahead schedule optimization. A scheme like this is already being widely used in Switzerland, where EWHs are heated up at fixed times during the night, making use of the usually low prices at those times, as well as reducing the peak load of the grid. In an *intra-day DR*, flexible loads are also used to perform an intra-day optimization based on real-time measurements and price information. The load behavior is thus altered from the plan made on the previous day.

We propose an intra-day DR scheme to minimize the money a BG has to spend on Balancing Energy (BE). This has one pitfall: it is possible to shift all flexible loads to low-cost hours. But if flexible loads are to be used for BG schedule compliance, they will also be activated during usually expensive peak hours. There is hence a trade-off between the opportunity to buy energy at low price, and the potential savings due to the ability to counteract some of the occurring schedule deviations.

The paper is organized as follows: Section 2 describes the DR and control topology envisioned. Section 3 details the control scheme for the intra-day DR. Section 4 describes the simulation setup used to test the control scheme, Section 5 concludes by giving and discussing detailed simulations results.

2. DEMAND RESPONSE AND CONTROL SCHEME

The DR scheme has to be able to handle the uncertainties associated with load forecasts and renewable generation. In previous research by (Borsche et al., 2013a),

⋆ This work was supported by the Swiss Commission for Technology and Innovation (CTI), as well as our project partners EKZ and Landis+Gyr.

an optimization problem for the day-ahead planning and energy acquisition under load uncertainty was proposed, taking into account the ability to shift flexible loads. To account for the uncertainty, demand scenarios based on the load forecast were created. The optimization then aims to minimize the expected cost over a set of some ten scenarios. In this paper, the intra-day dispatch of the DR scheme is investigated. Specifically, we try to answer the question: when and how should I utilize my flexible loads to achieve the maximal benefit for the BG? To answer this question, a hierarchical control scheme with three levels is proposed: 1) a day-ahead schedule optimization as described in (Borsche et al., 2013a); 2) a Model Predictive Control (MPC) for the intra-day decisions when to actually activate the DR; and 3) an inner control loop that tracks this reference r by sending broadcasts to the flexible loads. A set-up how this tracking can be achieved was described in (Borsche et al., 2013b), and we assume a similar system to be in place. Figure 1 summarizes the hierarchical control scheme, the focus of this paper lies on the scenario-based MPC.

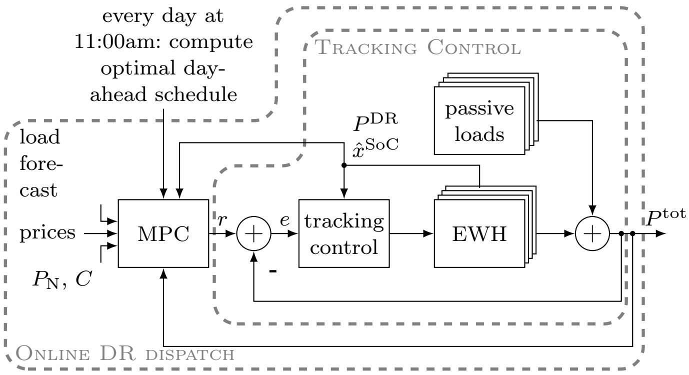

Fig. 1. Three-level control structure. A schedule for the BG is computed day-ahead. The MPC has real-time measurements of demand and up-to date forecast. Every 15 minutes, the MPC sends a new reference r to a tracking control, which in turn manages communication and control of the EWH aggregation.

3. SCENARIO-BASED MODEL PREDICTIVE CONTROL FOR INTRA-DAY DR DISPATCH

Notation Subscripts denote the time step, where t is the current time and k the number of time steps into the future. If some variable is depending on current observations, the index is appended by $|t$. Superscripts specify the concerning variable in more detail. If a variable depends not only on time but also on some scenario i, the superscript contains an i. Further we use the notation $\mathbb{N}_a^b$ for the set of all integers $\{[a, b]\}$.

Problem Formulation The goal of the MPC is co-minimization of energy cost and cost for BE. The MPC provides the current reference point r for the tracking control loop, confer Figure 1. The MPC receives the BG schedule $\overline{u}_k$ that was fixed day-ahead as input. Deviations from this schedule have to be covered by BE. The amount of BE consumed is denoted by $\epsilon^{\mathrm{s},i}$ for short positions and $\epsilon^{\mathrm{l},i}$ for long positions, and penalized by the price for positive BE c^+ and price for negative BE c^-, respectively.

The BG consists of both flexible loads, in our case EWHs, and passive loads. The passive loads can not be controlled and have an uncertain demand. This uncertainty is described by consumption scenarios $P^{\mathrm{BG},i}_{t+k|t}$. The electric power demand of the flexible loads $P^{\mathrm{DR},i}_{t+k|t}$ can be chosen within the constraints. (1e) captures the dependency between BE, demand and schedule. The MPC tries to set the flexible demand in such a way, that the expected cost for BE over all scenarios is minimized, see (1a). The first set-point r has to be a best-fit for all scenarios (1d). From stage two on-wards, the set-points may differ in order to account for future measurements.

In addition, the optimization guarantees that the State of Charge (SoC) and power limits are kept for each scenario (1c). The initial SoC is set to the current measurement or estimate of the aggregated SoC $\hat{x}^{\mathrm{SoC}}_{t|t}$, confer (1f). This value is received from the inner control loop. The final SoC is set to be higher than 50 % (1g). This terminal constraint is necessary, as the optimizer would otherwise let the SoC go against zero in order to minimize energy consumption and thus cost.

$$\min_{P^{\mathrm{DR},i}} \quad \sum_{k=0}^{N-1}\sum_{i=1}^{I}\frac{1}{I}\left(c^+_{t+k}\epsilon^{\mathrm{s},i}_{t+k|t} - c^-_{t+k}\epsilon^{\mathrm{l},i}_{t+k|t}\right) \tag{1a}$$

s.t.

$$\begin{aligned} x^{\mathrm{SoC},i}_{t+k+1|t} &= x^{\mathrm{SoC},i}_{t+k|t} + (P^{\mathrm{DR},i}_{t+k|t} + \xi_{t+k|t})\Delta t \\ &\qquad \forall k \in \mathbb{N}_0^{N-1}, \forall i \in \mathbb{N}_1^I \end{aligned} \tag{1b}$$

$$\begin{aligned} 0 &\le x^{\mathrm{SoC},i}_{t+k|t} \le C \\ 0 &\le P^{\mathrm{DR},i}_{t+k|t} \le P^{\mathrm{N}} \\ &\qquad \forall k \in \mathbb{N}_0^{N-1}, \forall i \in \mathbb{N}_1^I \end{aligned} \tag{1c}$$

$$r = P^{\mathrm{DR},i}_{t|t} \qquad \forall i \in \mathbb{N}_1^I \tag{1d}$$

$$\begin{aligned} \epsilon^{\mathrm{s},i}_{t+k|t} - \epsilon^{\mathrm{l},i}_{t+k|t} &= P^{\mathrm{BG},i}_{t+k|t} + P^{\mathrm{DR},i}_{t+k|t} - \overline{u}_{t+k} \\ 0 &\le \epsilon^{\mathrm{s},i}_{t+k|t} \\ 0 &\le \epsilon^{\mathrm{l},i}_{t+k|t} \qquad \forall k \in \mathbb{N}_0^{N-1}, \forall i \in \mathbb{N}_1^I \end{aligned} \tag{1e}$$

$$x^{\mathrm{SoC},i}_t = \hat{x}^{\mathrm{SoC}}_{t|t} \tag{1f}$$

$$0.5C \le x^{\mathrm{SoC},i}_{t+N} \tag{1g}$$

The horizon N depends on the execution time. The schedule is known until mid-night the next day. The horizon is therefore reduced every iteration, until a new schedule is approved, which again extends the horizon by 24 hours.

Feasibility is inherent to the problem formulation under certain assumptions. The initial SoC has to be within the constraints and mean consumption has to be lower than the nominal power of the boiler aggregation, implying that the installed EWHs are capable of providing the requested amount of hot water. But mainly, for each time step positive BE has to be more costly than spot energy, which in turn needs to be more expensive than negative BE

$$c^+_k > c^{\mathrm{spot}}_k > c^-_k \qquad \forall k \qquad . \tag{2}$$

Table 1. Aggregated parameters of flexible loads for the base simulation case

Variable	Description	Value
C	Capacity	54 MW h
$\bar{\xi}$	Mean demand	1.5 MW
P^{N}	Nominal power	9 MW

If this condition is violated, the problem is unbounded as it would be possible to buy, e.g., positive balancing energy and sell it at the spot market, making an infinite profit. There have been occurrences in the past when capped BE prices were lower than exceptionally high spot prices, motivating some BGs to "buy" energy via balancing markets. To prevent this kind of abusive behavior, the Swiss market design was changed to guarantee (2).

4. SIMULATION SETUP AND CASE STUDY

4.1 Description of Case Study

The base case for our case study is modelled around a typical Swiss 40 MW substation, relying on data provided by the utility of the Kanton Zurich (EKZ). The substation is serving mainly residential customers. Around one sixth of the electric energy consumption is used for electric water heating. Most EWHs are resistive heaters with a storage tank and can be switched using ripple control. The EKZ, which serves nearly 300 000 customers, has around 100 000 EWHs with a nominal power of more than 440 MW in a basic load-shifting scheme. EWHs are activated during off-peak (night) hours via ripple-control, and the resulting cost-savings are handed to customers in form of lower night-tariffs. The data about number of EWHs and statistics about consumption can be used to estimate the available number and power of EWHs at a single substation and the amount of energy used for water heating as follows: The average demand at the substation under discussion is slightly below 50 %. If furthermore half of the demand is attributed to residential customers, and of this one sixth to electric water heating, the mean hot-water consumption $\bar{\xi}$ is around 1.5 MW, giving a daily energy consumption of 36 MW h. The EWHs are allowed to run for 4 hours per day, and thus need a nominal power P^{N} of 9 MW to provide sufficient hot water. This is equivalent to 2000 EWHs installed and seems reasonable for a city of 25 000 inhabitants, considering much of the electric water heating is also done using gas or oil. The capacity is estimated to be 6 hours times the nominal heating power, giving $C = 54$ MW h in total. A similar estimate for $\bar{\xi}$ and P^{N} is found by dividing the 440 MW over the number of substations in the EKZ grid, and also by analyzing measured load patterns in which the peaks from the EWH switching at night times are clearly visible.

Other loads with load-shifting potential are not regarded in the base case. Some load types simple have a low penetration in Switzerland, e.g., air-conditioning. Others are difficult to control, e.g., refrigerators. Space-heating with heat pumps is similar to air-conditioning, but penetration is only recently increasing in Switzerland. However, the load characteristics – such as storage capacity and duty cycle – can have a significant impact on the load behavior in a DR scheme. This will be investigated in more detail in Section 5.3.

Prices. We use prices from the swissIX, the Swiss spot market, as well as balancing energy prices published by the Swiss Transmission System Operator (TSO) (Swissgrid, 2013). Spot prices are changing every hour, while balancing energy cost is computed for each 15-minute period. Spot prices are dominated by low prices at night due to abundant nuclear production, and high prices during day and especially peak hours. Balancing energy cost is rather predictable, as much of Switzerland's energy is produced by hydro storage plants, which are very flexible. In that respect, the Swiss market behavior is considerably different from other markets, such as Nordpool, where prices are much more flat during the day, or the Belgian power market, where balancing energy costs are much more volatile and unpredictable. We will argue later why this matters much for the problem at hand.

Load Forecast and Scenario Generation. Our approach for load forecasting and scenario generation is described in more detail in our previous work, see (Borsche et al., 2013a). To forecast the load, a neural auto-regressive network with exogenous inputs is used. These inputs are month, day-of-week, work- or holiday, and temperature. The MAPE for the 40 MW substation under discussion is around 4.9 %, which seems reasonable. Scenarios are generated using a method proposed by (Pinson et al., 2009). First, historic forecast-realization pairs are used to identify quantiles and auto-correlation of forecast errors. Then, a new forecast for the next day is generated. The knowledge about past errors is used to generate scenarios for the load deviation from that forecast. All load curves and forecasts are in 15-minute resolution.

4.2 Benchmarks

To assess the potential cost savings and performance of the **scenario-based intra-day DR (sb-idDR)**, it is compared to two benchmarks

day-ahead DR (daDR) the traditional approach, where flexibility is only used in the day-ahead scheduling to minimize energy acquisition cost,

Scenario-based day-ahead DR (sb-daDR) similar to the traditional approach, but the day-ahead schedule is now made using scenarios, thus taking into account uncertainty in the forecast already during the planning phase,

performance bound on intra-day DR (pb-idDR) DR is used to minimize both energy cost and schedule deviations. While the day-ahead scheduling is done under uncertainty, the intra-day decisions about the DR activation are done with perfect knowledge. This is the performance bound for the scenario-based MPC.

All costs given in this paper are the additional cost relative to the cost which would be achieved in a deterministic acquisition with perfect knowledge of future demand. For a day-ahead DR, this is the cost for BE. If the intra-day DR approach is used, it is the combination of opportunity cost and cost for BE.

5. RESULTS AND DISCUSSION

The following section gives results on the base case described in the previous section. To get a wider comprehension of the behaviour of the intra-day DR scheme, sensitivity analyses were run. These analyses cover the number of scenarios used in the optimization, the sensitivity to load parameters capacity, duty cycle and nominal power, and finally the effect of the price structure on the load behaviour. The section is concluded by a discussion of the findings.

5.1 Base Case Results

A simulation covering half a year of simulation time was run for the base case. Figure 2 shows results for a typical day. The blue line gives the behaviour of the day-ahead DR which activates loads at low spot prices. The green line is the activation when using intra-day DR. Here, some loads are activated during more expensive hours around mid-day in order to reduce BE cost. However, even with the intra-day DR, most of the energy is consumed during low price hours. Over the half-year period, on average 22.6 % of the energy was used for schedule compliance.

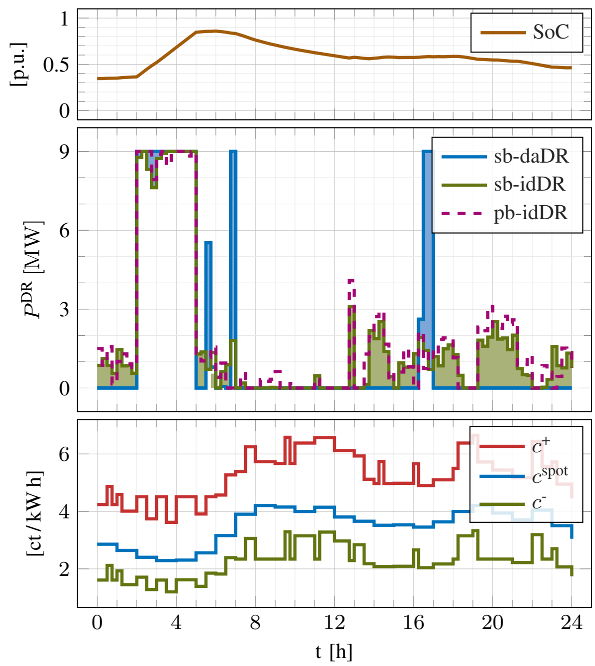

Fig. 2. Sample Day. Top shows SoC evolution, middle shows DR activation in different schemes and bottom shows the prices driving the activation. The scenario-based day-ahead DR (sb-daDR, middle, blue) activates flexible loads at low cost hours. The intra-day DR (sb-idDR) also activates loads during more expensive afternoon hours to reduce schedule deviations. The performance bound computed with no uncertainty during the day is given in the dashed line (pb-idDR).

Accordingly, cost savings are small. The cost given in Table 2 are the additional cost on top of an energy acquisition with perfect knowledge and no uncertainty. The traditional day-ahead DR approach has the highest associated cost of around EUR 485 per day. Using scenarios to account for the uncertainty in the load forecast, this cost can already be reduced to EUR 401. If the flexible loads are used to balance schedule deviations, cost can be further reduced to EUR 249 per day – assuming a controller with perfect knowledge. The proposed scenario-based MPC does not quite achieve this, but with associated BE cost of EUR 353 is still offers a cost saving of EUR 130 or 26 % over the traditional approach. While savings scale with the number of substation of a utility, the savings per participant are rather low. Assuming 2000 EWHs in the DR scheme, each EWH accounts for a saving of EUR 24 per year.

Table 2. Results for the base case simulation, comparing cost incurred with day-ahead DR (daDR) and intraday DR (idDR).

DR scheme	daily BE cost [EUR]	daily BE cost %	flexible demand used for BE reduction [%]
daDR	485.15	2.14	-
sb-daDR	401.40	1.77	-
pb-idDR	248.86	1.10	24.44
sb-idDR	353.03	1.56	22.27

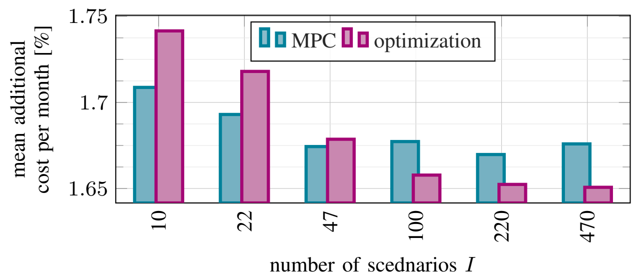

Fig. 3. Sensitivity of cost to number of scenarios used in the MPC and the optimization stage.

5.2 Sensitivity to Number of Scenarios

To assess the effect of the number of scenarios on the cost savings, the same simulation was run repeatedly with different numbers of scenarios for both the day-ahead optimization, I^{opt}, and the intra-day MPC, I^{MPC}. Both are chosen from $\{10, 22, 47, 100, 220, 470\}$. All combinations between these two sets are used, giving 36 parameter pairs. Each simulation is covering a period of one month.

The average BE and opportunity cost is shown in Figure 3, plotted over the scenarios used for both the optimization phase and the MPC. For example, the first purple bar shows the average additional cost for all simulations where $I^{\text{opt}} = 10$, while the first blue bar is the mean additional cost for all simulations with $I^{\text{MPC}} = 10$. It can be clearly seen that additional scenarios have no noticeable effect when $I^{\text{MPC}} \geq 50$ and $I^{\text{opt}} \geq 100$. It is also noteworthy, that the optimization phase seems to depend generally much more on the numbers of scenarios chosen.

5.3 Sensitivity to Load Parameters

Load parameters might seriously affect the usage pattern of the intra-day DR system. Consider a load which can only be shifted by half an hour: reducing BE cost might

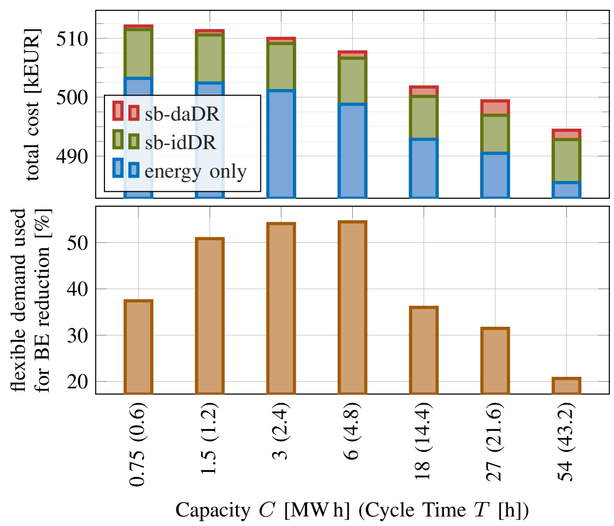

Fig. 4. *Sensitivity to capacity C.* Top: increasing the capacity reduces overall cost, while cost for BE first decreases, then increases again. Bottom: in an intra-day DR scheme, the amount of energy used to reduce schedule deviations is highest for a capacity of 3 MW h, which translates into a cycle time of $T \approx$ 2.4 h.

be much more interesting as low energy prices are out of reach. To understand the effect of load type better, we run a sensitivity analysis with regard to the nominal power P^{N}, nominal capacity C and duty cycle D. These parameters are linked to the mean demand $\bar{\xi}$ and the natural cycle period T as follows

$$\bar{\xi} = DP^{\mathrm{N}} \quad , \tag{3}$$

$$T = \frac{C}{P^{\mathrm{N}} - \bar{\xi}} + \frac{C}{\bar{\xi}} \quad . \tag{4}$$

Increasing the duty cycle with fixed P^{N} implies increasing demand, while decreasing the storage capacity linearly affects the natural cycle period.

First, the effect of storage capacity was investigated. Results are given in Figure 4. Figure 4 top shows total cost over storage capacity. The blue bar gives the energy cost that would be incurred in a deterministic acquisition. The cost decreases, as a bigger storage allows to make more use of low cost hours. The red bar refers to the scenario-based day-ahead DR, while the green bar is the cost incurred when using the intra-day DR. The visible red area is thus the amount that can be saved by using DR for schedule compliance. These savings generally increase with storage capacity. However, this is not true for the biggest storage investigated. This result, counter-intuitive at first sight, is due to the co-optimization of BE and energy cost. Figure 4 bottom depicts the percentage of flexible energy available that is shifted from low cost hours to more expensive hours to reduce BE cost. In other words, it shows how the intra-day DR differs form the day-ahead DR in terms of demand allocation. It can be seen that most energy is shifted for a storage capacity around 5000 kW h or $T \approx 3$ h. Shorter cycle times do not allow for effective shifting, while for longer cycle times buying at low cost is more effective than reducing BE cost.

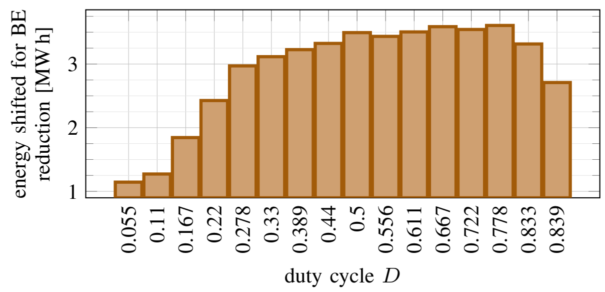

Fig. 5. *Sensitivity to demand.* Shifted energy increases with the duty cycle.

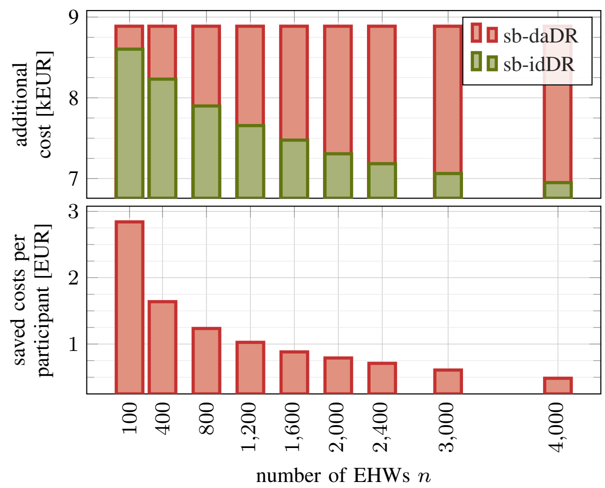

Fig. 6. *Sensitivity to number of participants n.* Top: Saved cost for BE grow sub-linear. Bottom: Participants therefore earn most when the savings are split between few.

Next, the power rating was adjusted between 3000 kW and 12 000 kW, while ξ and C were kept constant. Results are inconclusive. As expected, cost is reduced when the nominal power is increased – demand is now even more concentrated at the lowest price hours. However, savings in BE cost is only slightly affected by the power rating.

Third, increasing the duty cycle increases the incentive to shift energy. The duty cycle was changed between $D = 0.056$ and 0.889, which is equal to an demand between 500 kW and 8000 kW. Figure 5 shows the absolute amount of energy shifted with respect to the duty cycle. Long duty cycles "block" low cost hours, leading to more energy shifted in the remaining time.

Finally, the parameters $\bar{\xi}$, C and P^{N} are scaled to represent different numbers of participants in the intra-day DR program. The base case assumes 2000 EWHs. Figure 6 top shows BE cost for a day-ahead DR in red and intra-day DR in green. While cost reduction increase with the number of flexible loads, the growth seems to saturate at a certain point. In turn, Figure 6 bottom shows how the reward per EWH is reduced as savings are split between more participants.

5.4 Sensitivity to Price Structure

As argued earlier, prices in Switzerland are different from other markets with respect to spread in the spot market

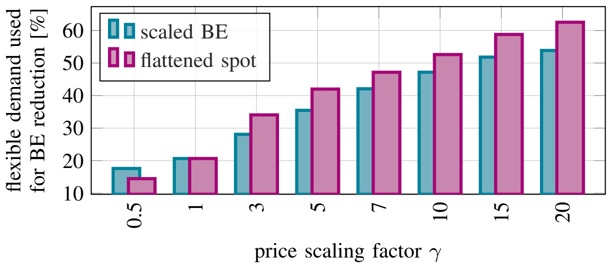

Fig. 7. Both a higher spread between BE price and spot price as well as more flat spot prices significantly increase usage of the intra-day DR for schedule compliance.

and spread between spot and BE prices. To assess the price-sensitivity of the intra-day DR scheme, the spread between spot price and price for BE was scaled by a factor γ

$$c^{+\prime} = c^{\text{spot}} + \gamma(c^{+} - c^{\text{spot}}) \quad , \tag{5}$$

$$c^{-\prime} = c^{\text{spot}} + \gamma(c^{-} - c^{\text{spot}}) \quad . \tag{6}$$

In a second simulation, the spot prices were flattened according to

$$c^{\text{spot}\prime} = \frac{1}{\gamma}\left(c^{\text{spot}} + (\gamma - 1)\bar{c}^{\text{spot}}\right) \quad , \tag{7}$$

$$c^{+\prime} = c^{\text{spot}\prime} + (c^{+} - c^{\text{spot}}) \quad , \tag{8}$$

$$c^{-\prime} = c^{\text{spot}\prime} + (c^{-} - c^{\text{spot}}) \quad , \tag{9}$$

where $\bar{c}^{\text{spot}}$ refers to the mean of the spot price. This keeps the spread between BE and spot as in the original data, but reduces the spread between high and low price hours at the spot market. Figure 7 shows how the flexibility is used more and more to decrease BE energy, if either the BE price is increased or the spread at the spot market reduced. While an increase in the BE cost by a factor of five or more seems unlikely, a flattening of the spot price is feasible – and might actually be caused by an increasing number of flexible consumers trying to take advantage of low price hours.

5.5 *Discussion*

The low earnings observed in the base case might be prohibitive considering the infrastructure needed to implement an intra-day DR scheme. Considering the sensitivity analysis, several more statements can be made. Increasing the storage capacity of flexible demand reduces overall cost, but mainly due to reduced costs at the spot market. Increasing the power rating does not affect the profitability. However, when loads have a duty cycle of around one half, much of the flexibility can be shifted economically to avoid schedule deviations. Loads with such a behaviour might be the most promising candidates for reducing schedule deviations. Next, if such a scheme is implemented, initial cost will be high compared to incremental costs for an additional participant. However, our results indicate that too many participants will also be detrimental, as profits scale sub-linear.

6. CONCLUSION AND OUTLOOK

In this paper, a control scheme allowing a balance group to significantly reduce its schedule deviations was presented. The control handles the trade-off between the cost for balancing energy and opportunity cost of not utilizing low spot market prices. Estimating the future demand is essential to making good real-time decisions about the activation of flexible loads. This uncertainty is handled using scenarios based on the most recent forecast.

Results show a certain potential to reduce overall costs incurred by the balance group. Under the current price regime, it is not profitable to use such an elaborate scheme with electric water heaters in Switzerland. However, loads with a higher duty cycle might be of interest. Also, the price structure heavily influences the profitability of the scheme. Finally, there is an optimal number of participants, as potential savings saturate. Ideally, such a control would rely on existing infrastructure. Future work should focus on the implementation, as the main unknowns in this setup are investment and operating costs.

REFERENCES

Borsche, T., Oldewurtel, F., and Andersson, G. (2013a). Balance Group Optimization by Scenario-based Day-ahead Scheduling and Demand Response. In *4th IEEE PES Innovative Smart Grid Technologies Europe (ISGT Europe)*. Copenhagen.

Borsche, T., Oldewurtel, F., and Andersson, G. (2013b). Minimizing Communication Cost for Demand Response using State Estimation. In *IEEE PowerTech*. Grenoble.

Callaway, D.S. (2011). Can smaller loads be profitably engaged in power system services? In *2011 IEEE Power and Energy Society General Meeting*, 1–3. San Diego. doi:10.1109/PES.2011.6039890.

Koch, S., Mathieu, J.L., and Callaway, D.S. (2011). Modeling and control of aggregated heterogeneous thermostatically controlled loads for ancillary services. In *17th Power Systems Computation Conference*. Stockholm.

Kondoh, J., Lu, N., and Hammerstrom, D.J. (2011). An evaluation of the water heater load potential for providing regulation service. *IEEE Transactions on Power Systems*, 26(3), 1309–1316. doi:10.1109/TPWRS.2010.2090909.

Mathieu, J.L., Koch, S., and Callaway, D.S. (2013). State Estimation and Control of Electric Loads to Manage Real-Time Energy Imbalance. *IEEE Transactions on Power Systems*, 28(1), 430–440. doi:10.1109/TPWRS.2012.2204074.

Pinson, P., Madsen, H., Nielsen, H.A., Papaefthymiou, G., and Klöckl, B. (2009). From probabilistic forecasts to statistical scenarios of short-term wind power production. *Wind Energy*, 12(1), 51–62. doi:10.1002/we.284.

Swissgrid (2013). swissgrid balancing energy prices. `http://www.swissgrid.ch/swissgrid/de/home/experts/topics/bgm/balance_energy.html`. Accessed: 2013-10-30.

Vrettos, E., Oldewurtel, F., Vasirani, M., and Andersson, G. (2013). Centralized and decentralized balance group optimization in electricity markets with demand response. In *IEEE PES PowerTech 2013*. Ieee, Grenoble. doi:10.1109/PTC.2013.6652519.