Scheduling Distributed Energy Storage Units to Provide Multiple Services

Olivier Mégel
Power Systems Laboratory
ETH Zürich, Switzerland
megel@eeh.ee.ethz.ch

Johanna L. Mathieu
Electrical and Computer Engineering
University of Michigan, USA
jlmath@umich.edu

Göran Andersson
Power Systems Laboratory
ETH Zürich, Switzerland
andersson@eeh.ee.ethz.ch

***Abstract*—When energy storage units (*ESUs*) within the distribution grid, e.g. batteries, provide local services such as PV integration support, peak shaving, and infrastructure upgrade deferral, they are inactive or only partially used most of time. Moreover, they are often not profitable because of their high investment costs. Their unused capacities could be used to provide power system services, such as frequency control, allowing them to generate additional revenues. However, individual units might not be available to provide system services over the entire contract duration, since they must also provide their local services. This paper shows how an aggregation of distributed ESUs can simultaneously provide local services individually and system services in aggregate. Using a model predictive control approach, a central scheduler dynamically allocates parts of the energy and power capacities of each ESU to either the local or grid service with the objective of maximizing the profit of the aggregation. A key contribution of this paper is the development of an algorithm that handles both resource aggregation and optimal provision of multiple services. We find that multitasking can almost double an ESU's profits as compared with a single-service approach, and that the benefits from aggregation depend on the grid service market structure and how often the local service is required.**

Keywords– distributed storage; battery; scheduling

I. Introduction

The number of energy storage units (*ESUs*) within the distribution grid is likely to increase since they can be used for a variety of local services including photovoltaic (*PV*) integration support, peak shaving, infrastructure upgrade deferral, and powering electric vehicles. However, the purchase cost of distributed ESUs, especially batteries, is expected to remain high in the near- to middle-term future [1]. A way to improve the economics of an ESU was described in [2]: when not fully used for its local service, an ESU could provide other services to power systems, such as frequency control. This so-called multitasking approach has been the subject of several recent publications, for example, [3] which analyzes storage capacity allocation under grid constraints, considering spot and intraday markets, and frequency control markets simultaneously, and [4] where the focus is on peak shaving, electricity trading, frequency control, and uninterruptible power supply service.

A key challenge to power system service provision with ESUs is that individual units might not be available to provide the contracted service over the entire contract duration because they must also provide their local service. Therefore, there is a benefit to resource aggregation. Many papers have investigated the use of aggregations of distributed resources with limited energy capacities to provide both local and system services, for example, [5] develops methods to schedule/control electric vehicle charging to minimize charging costs and provide frequency regulation while minimizing negative network impacts, [6] evaluates the ability of electric water heaters to provide frequency control, and [7] develops methods to schedule/control thermostatically controlled loads to provide frequency regulation in addition to actively managing the distribution network to increase PV energy absorption. However, these papers do not co-optimize the allocation of resources to the local and system services. Instead, they either allocate a predefined power capacity to frequency control, or they optimize the schedule/control for the local services and then use the remaining flexibility for frequency control. Also, they focus on aggregations of resources providing the same local service, rather than resources providing diverse local services.

The contributions of this paper are threefold. First, we develop an algorithm that dynamically co-optimize the allocation of diverse ESUs' energy and power capacities over local and power system services with the aim of maximizing the profit of the aggregation. Our algorithm provides day- and hour-ahead allocation schedules that tell each ESU how much power and energy capacity to reserve for each service at each time step; our focus is not on real-time control. Second, through a case study, we investigate the benefits of combining aggregation and multitasking. Third, an additional case study highlights the impact of the power system service contract duration.

We restrict our analysis to batteries as the sole energy storage technology; however, we consider batteries providing different local services. Since batteries have significant degradation costs (per kWh cycled), they are only attractive for local services when the alternative would lead to a higher cost per kWh. We therefore restrict ourselves to local services for which batteries would perform well:

- *PV-l*: minimization of PV curtailment when subject to line export limitation. If no battery were installed at this location, the alternative would be to either curtail the PV generator or to upgrade the line;
- *PV-t*: minimization of PV curtailment when the PV generator is connected to the grid through a transformer with limited thermal capacity. If no battery were installed at this location,

the alternatives would be to either curtail the PV generator or to upgrade the transformer;

- *Load-l*: customer load smoothing when the load profile sometimes exceeds the line import power. If no battery were installed at this location, the alternative would be to either curtail the load or upgrade the line;
- *Load-t*: customer load smoothing when the customer is connected to the grid through a transformer with limited thermal capacity. If no battery were installed at this location, the alternative would be to either curtail the load or upgrade the transformer.

The only power system service we consider is primary frequency control (*PFC*), since we found in [8] that it might soon become cost effective in Europe to provide this service with batteries. If secondary frequency control revenues increase or battery costs decrease, it could be considered as well, using the same methodology as for PFC.

Section II describes the models and scheduling algorithm. Section III details the case studies, Section IV discusses the results, and Section V provides concluding remarks.

II. Models and Scheduling Algorithm

To model an ESU aggregation, we define a Local Area Control (*LAC*) as a building block. Each LAC, designated by the subscript i, represents one ESU and its local environment, and contains at least:

- One ESU, characterized by its energy and power capacities (E_i^{cap}, P_i^{cap}), its charge and discharge efficiencies (η_i^{c}, η_i^{d}), its self-discharge per time step (1-η_i^{sd}), and its linear and quadratic degradation costs (d_i^{l}, d_i^{q});
- One (aggregated) load profile;
- One electricity tariff profile (purchases) and one electricity feed-in tariff profile (sales).

Depending on the local service provided, the following options can also be present:

- One PV generation profile (local service: PV-l or PV-t);
- One line with limited export/import capacity, which creates a bottleneck between the grid and the LAC (local service: PV-l or Load-l);
- One transformer with limited thermal capacity that creates a bottleneck between the grid and the LAC (local service: PV-t or Load-t),

and the following optimization variables:

- One PV curtailment profile (local service: PV-l or PV-t);
- One load curtailment profile (local service: Load-l or Load-t).

Since we focus on grid constraints associated with ESUs connected to the rest of the network through a bottleneck, we do not explicitly model power flow, as in [7]. In the future, our algorithm could be extended to explicitly include network constraints.

The goal of the scheduling algorithm is to maximize the profit of an existing ESU aggregation. We do not consider investment costs, but instead we do consider ESU battery degradation costs as operational costs. The problem of how to attribute the benefits to all possible stakeholders is not considered here, nor do we address regulatory barriers to multitasking [3].

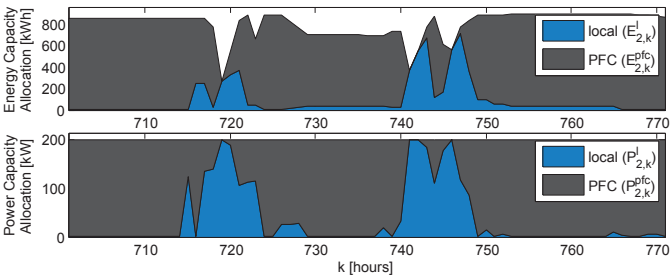

Fig. 1. Example of energy and power capacity allocation schedule.

We use a Model Predictive Control (*MPC*) [9] approach to compute the allocation schedule. This receding horizon approach allows us to handle plant-model mismatch and could easily be extended in the future to handle forecast error and stochasticity as well. In this paper, we are primarily interested in understanding upper bounds on ESU aggregation revenues achievable via multitasking and so we do not investigate the effect of forecast error and stochasticity. However, we use MPC to handle limited-horizon, but perfect forecasts (detailed in Section III) and the mismatch between the transformer model used to represent the real system and that used within our controller (detailed in Section II-C). We used YALMIP [10] to represent our set of equations and constraints, and to build our MPC controller. Since we aim to control large numbers of ESUs, we use linear and quadratic models, ensuring that the system model is computationally tractable.

In the following four sections we describe our modeling approaches. Section II-A presents our ESU energy and power capacity allocation model, while Section II-B describes our methods of modeling the cost/profit associated with buying/selling electricity from/to the grid, the cost of battery degradation, and the load curtailment penalty. Section II-C describes our transformer model and overheating penalty, and Section II-D details how we model the profit realized through PFC provision. Finally, Section II-E gives the full mathematical description of our MPC controller by bringing together the results of the previous sections.

A. Allocation of Capacities

Our algorithm dynamically allocates fractions of the energy and power capacities of each ESU to either its local service or to frequency control, as shown in Fig. 1. For each timestep k and each LAC, E_i^{cap} and P_i^{cap} are divided into a part that serves the local service ($E_{i,k}^{\text{l}}$, respectively $P_{i,k}^{\text{l}}$) and a part that serves PFC ($E_{i,k}^{\text{pfc}}$, respectively $P_{i,k}^{\text{pfc}}$):

$$SoC_i^{\text{l}} \cdot E_i^{\text{cap}} \leq E_{i,k}^{\text{l}} + E_{i,k}^{\text{pfc}} \leq SoC_i^{\text{u}} \cdot E_i^{\text{cap}} \quad (1)$$

$$0 \leq P_{i,k}^{\text{l}} + P_{i,k}^{\text{pfc}} \leq P_i^{\text{cap}} \quad (2)$$

where SoC_i^{l} and SoC_i^{u} are lower and upper State-of-Charge (SoC) limits (enforced to avoid operating areas associated with

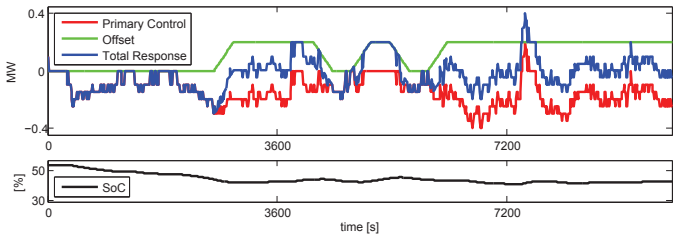

Fig. 2. Example of frequency response offsetting. The opposite of the offset (green line) should be provided by a different resource.

excessive degradation) and $P^{\mathrm{l}}_{i,k}$ and $P^{\mathrm{pfc}}_{i,k}$ represent the absolute value of the power that can be extracted or injected for the local service and PFC, respectively (we assume symmetric charge and discharge power capacities).

As it is impossible to predict the instantaneous energy and power capacities requested by PFC, we use a statistical approach to estimate values of $E^{\mathrm{pfc}}_{i,k}$ and $P^{\mathrm{pfc}}_{i,k}$ needed to deliver a specific amount of PFC, leading to a reformulation of (1). For this, we assume that an ESU's SoC is managed via an "offsetting mechanism" as described in [8]. In the next few paragraphs, we provide a short summary of this work.

When providing frequency control, ESUs might run out of energy or become fully charged if the frequency signal is not zero-mean over a certain time period, and hence no longer be able to contribute to PFC. A way to solve this problem is to break the frequency control signal into two parts: i) a fast, zero-mean signal and ii) a time-dependent offset, as shown in Fig. 2. The ESU follows the fast, zero-mean signal, while the opposite of the offset is contracted from other, less-flexible, but energy-unconstrained resources such as conventional power plants. Different methods of offset computation can be found in [8], [11]–[14]. These methods have in common that a part of $P^{\mathrm{pfc}}_{i,k}$ has to be reserved for the offset ($P^{\mathrm{pfc,off}}_{i,k} > 0$):

$$P^{\mathrm{pfc}}_{i,k} = P^{\text{e-pfc}}_{i,k} + P^{\mathrm{pfc,off}}_{i,k} \tag{3}$$

where $P^{\text{e-pfc}}_{i,k}$ is the power capacity which is effectively available for PFC (the value that would be communicated to the system operator). The size of $P^{\text{e-pfc}}_{i,k}$ compared to $P^{\mathrm{pfc}}_{i,k}$ depends on $E^{\mathrm{pfc}}_{i,k}/P^{\mathrm{pfc}}_{i,k}$, the ESU's characteristics, and the choice of the offset mechanism.

Ref. [8] analyzed three years of historical frequency data, and found that an ESU providing PFC with an appropriate offset mechanism has a narrow SoC distribution centered around $SoC^{\mathrm{pfc}}_{\mathrm{avg}}$, as shown in Fig. 3. Note that it is very rare for the SOC to reach $SoC^{\mathrm{pfc}}_{\mathrm{min}}$ or $SoC^{\mathrm{pfc}}_{\mathrm{max}}$. For simplicity we assume that the energy content of the portion of the ESU dedicated to PFC is $SoC^{\mathrm{pfc}}_{\mathrm{avg}} \cdot E^{\mathrm{pfc}}_{i,k}$ at all hourly timesteps (yellow line in Fig. 4), but that the SoC^{pfc} may vary between $SoC^{\mathrm{pfc}}_{\mathrm{min}}$ and $SoC^{\mathrm{pfc}}_{\mathrm{max}}$ within a timestep, due to the frequency profile. To include these assumptions, we replace (1) with:

$$SoC^{\mathrm{l}}_i \cdot E^{\mathrm{cap}}_i \leq E^{\mathrm{l}}_{i,k} + SoC^{\mathrm{pfc}}_{\mathrm{min}} \cdot E^{\mathrm{pfc}}_{i,k} \tag{4}$$

$$E^{\mathrm{l}}_{i,k} + SoC^{\mathrm{pfc}}_{\mathrm{max}} \cdot E^{\mathrm{pfc}}_{i,k} \leq SoC^{\mathrm{u}}_i \cdot E^{\mathrm{cap}}_i \tag{5}$$

where (4) corresponds in Fig. 4 to the distance between the green line and the bottom dashed line, and (5) to the distance

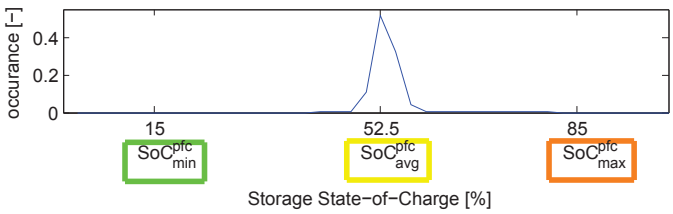

Fig. 3. SoC distribution of an ESU providing PFC with appropriate offsetting mechanism, from [8].

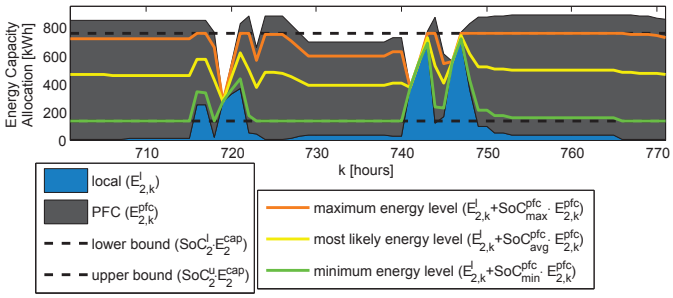

Fig. 4. Example of energy and power allocation schedule including SoC limits, $SoC^{\mathrm{pfc}}_{\mathrm{min}}$, $SoC^{\mathrm{pfc}}_{\mathrm{avg}}$, and $SoC^{\mathrm{pfc}}_{\mathrm{max}}$. The allocation is the same as in Fig. 1. The color coding corresponds to that of Fig. 3. When no PFC is provided, the three color lines converge since $E^{\mathrm{pfc}}_{i,k} = 0$.

between the orange line and the top dashed line. In this paper, we use values of $SoC^{\mathrm{pfc}}_{\mathrm{min}}$ and $SoC^{\mathrm{pfc}}_{\mathrm{max}}$ from [8], which are displayed in Fig. 3, and we set $SoC^{\mathrm{l}}_i = SoC^{\mathrm{pfc}}_{\mathrm{min}}$ and $SoC^{\mathrm{u}}_i = SoC^{\mathrm{pfc}}_{\mathrm{max}}$.

For the local service, $E^{\mathrm{l}}_{i,k}$ and $P^{\mathrm{l}}_{i,k}$ represent both how much capacity (in terms of kWh and kW, respectively) is used by the local service and how much is reserved for it. In other words, the optimizer reserves only as much energy or power capacities that it predicts it will need, which is valid because of our assumption of perfect forecasts. Therefore, $P^{\mathrm{l}}_{i,k}$ can be expressed as:

$$P^{\mathrm{l}}_{i,k} = \max(p^{\mathrm{s,i}}_{i,k}, p^{\mathrm{s,o}}_{i,k}) \tag{6}$$

where $p^{\mathrm{s,i}}_{i,k}$ and $p^{\mathrm{s,o}}_{i,k}$ (both positive) represent the power sent to and requested from the ESU i by its local service. In a real implementation including forecast error, one would need to build in robustness margins.

The partitioning of energy and power capacities can change at any timestep, as shown in Fig. 1. To represent this, we include the variable $\Delta E_{i,k}$ which represents a virtual energy transfer from one service to the other, i.e.

$$SoC^{\mathrm{pfc}}_{\mathrm{avg}} \cdot E^{\mathrm{pfc}}_{i,k+1} = \eta^{\mathrm{sd}}_i \left(SoC^{\mathrm{pfc}}_{\mathrm{avg}} \cdot E^{\mathrm{pfc}}_{i,k} + \Delta E_{i,k} \right) \tag{7}$$

$$E^{\mathrm{l}}_{i,k+1} = \eta^{\mathrm{sd}}_i \left(E^{\mathrm{l}}_{i,k} + \left(\eta^{\mathrm{c}}_i p^{\mathrm{s,i}}_{i,k} - \frac{p^{\mathrm{s,o}}_{i,k}}{\eta^{\mathrm{d}}_i} \right) \Delta t - \Delta E_{i,k} \right) \tag{8}$$

where Δt represents the timestep duration (one hour). Note that these energy conservation equations use $SoC^{\mathrm{pfc}}_{\mathrm{avg}} \cdot E^{\mathrm{pfc}}_{i,k}$ (the yellow line in Fig. 4), not the min/max values.

B. Physical Constraints, Cost, & Profits

For each LAC, the power balance must be satisfied at all timesteps:

$$L_{i,k} - L^{\mathrm{X}}_{i,k} = p^{\mathrm{g,i}}_{i,k} + p^{\mathrm{s,o}}_{i,k} + PV_{i,k} - p^{\mathrm{g,o}}_{i,k} - p^{\mathrm{s,i}}_{i,k} - PV^{\mathrm{X}}_{i,k} \tag{9}$$

where $L_{i,k}$ and $L^{\mathrm{X}}_{i,k}$ represent the load profile and load curtailment profile, $p^{\mathrm{g,i}}_{i,k}$ and $p^{\mathrm{g,o}}_{i,k}$ the power consumed from and exported to the grid, $PV_{i,k}$ and $PV^{\mathrm{X}}_{i,k}$ the PV generation profile and PV curtailment profile. PV curtailment is not explicitly penalized, but there is an opportunity cost associated with the loss of revenues from the feed-in tariff. The load curtailment penalty $C^{\mathrm{X}}_{i,k}$ is calculated as:

$$C^{\mathrm{X}}_{i,k} = x \cdot L^{\mathrm{X}}_{i,k} \cdot \Delta t \tag{10}$$

with x being the cost of load curtailment (set to 10€/kWh, which is in line with [15]). The cost/profit associated with buying/selling electricity from/to the grid $C^{\mathrm{e}}_{i,k}$ is computed assuming that the same tariffs apply to the ESU as to the load and PV:

$$C^{\mathrm{e}}_{i,k} = \left(p^{\mathrm{g,i}}_{i,k} \cdot g^{\mathrm{in}}_{i,k} - p^{\mathrm{g,o}}_{i,k} \cdot g^{\mathrm{out}}_{i,k}\right) \Delta t \tag{11}$$

where $g^{\mathrm{in}}_{i,k}$ and $g^{\mathrm{out}}_{i,k}$ are the purchase and feed-in tariffs. We assume that the battery degradation cost function $C^{\mathrm{d}}_{i,k}$ is the sum of linear and quadratic functions on the battery power profile:

$$C^{\mathrm{d}}_{i,k} = \left(d^{\mathrm{l}}_i \cdot (p^{\mathrm{s,i}}_{i,k} + p^{\mathrm{s,o}}_{i,k}) + d^{\mathrm{q}}_i \frac{(p^{\mathrm{s,i}}_{i,k} + p^{\mathrm{s,o}}_{i,k})^2}{E^{\mathrm{cap}}_i}\right) \Delta t \tag{12}$$

The line constraints are modeled as:

$$p^{\mathrm{g,i}}_{i,k} \leq p^{\mathrm{g,i,max}}_{i,k} \tag{13}$$

$$p^{\mathrm{g,o}}_{i,k} \leq p^{\mathrm{g,o,max}}_{i,k} \tag{14}$$

where $p^{\mathrm{g,i,max}}_{i,k}$ and $p^{\mathrm{g,o,max}}_{i,k}$ are the line limitations.

C. Transformer overheating

The power rating of a transformer can cause a bottleneck in the grid. The limiting factor is not usually the power rating itself, but rather the transformer hot spot temperature, which is related to the load factor profile through a set of thermal relations considering thermal inertia. An elevated hot spot temperature shortens the transformer lifetime, and can also lead to immediate damage if the temperature is above a certain threshold [16].

To ensure that our MPC algorithm is computationally tractable, we developed a simplified transformer overheating model based upon the higher order models of [16] and [17], and extending the initial work of [18]. Using the numerical values of distribution transformer model parameters given in [16], we first found that $\Delta\theta_{\mathrm{h2},k}$ from [17], a system state of the transformer model, can be neglected. Since we also assume the transformer load factor $K_{i,k}$ to be constant over one hourly timestep, the equation (3) from [17] can be approximated to a form that does not include dependency on the previous timestep, and the equation (5) can be simplified as well. These three simplifications allow us to use an hourly timestep (as opposed to minute duration time steps in [17]) and lead to the following equation set:

$$\theta^{\mathrm{h}}_{i,k} = \theta^{\mathrm{o}}_{i,k} + \Delta\theta h^{\mathrm{r}} K^{1.6}_{i,k} \tag{15}$$

$$\theta^{\mathrm{o}}_{i,k} = \frac{\theta^{\mathrm{o}}_{i,k\text{-}1}}{1.3957} + 0.2835\left(\left(\frac{1+K^2_{i,k}R}{1+R}\right)^{0.8} \Delta\theta^{\mathrm{r}}_o + \theta^{\alpha}_{i,k}\right) \tag{16}$$

$$K_{i,k} = \frac{\max(p^{\mathrm{g,i}}_{i,k}, p^{\mathrm{g,o}}_{i,k})}{S_i} \tag{17}$$

where $\theta^{\mathrm{h}}_{i,k}$ is the hot spot temperature of the transformer; $\theta^{\mathrm{o}}_{i,k}$ its top-oil temperature; $\theta^{\alpha}_i(k)$ the ambient temperature; $\Delta\theta h^{\mathrm{r}}$, $\Delta\theta^{\mathrm{r}}_o$, and R are thermal parameters; and S_i is the transformer rated power.

Equations (15) and (16) are convex but not linear or quadratic, and so we approximate them with sets of linear equations that provide piecewise linear lower bounds on $K_{i,k}$. We then penalize $K_{i,k}$ in the objective function so that it stays on the lower bound. Tests show that the output of our model and the model in [17] match well; the outputs do not differ by more than a few temperature degrees.

Since we want to focus on extreme overloading events, we penalized the rise of the hot spot temperature only above $160^{\circ}\mathrm{C}$, and since the risk of damage increases rapidly in high temperature ranges [16], we use a quadratic penalty function:

$$C^{\mathrm{t}}_{i,k} = h^{\mathrm{t}}_i \cdot \left(\max(\theta^{\mathrm{h}}_{i,k} - 160^{\circ}\mathrm{C}, 0)\right)^2 \tag{18}$$

where $C^{\mathrm{t}}_{i,k}$ is the transformer overheating penalty and h^{t}_i the cost of transformer overheating.

This transformer model is used within the MPC controller; however, to represent the real transformer we use the original formulation from [17] which provides $\theta^{\mathrm{o2}}_{i,k}$, a more accurate transformer top-oil temperature estimate. This introduces plant-model mismatch. At each iteration of the controller, we set the initial value of $\theta^{\mathrm{o}}_{i,k}$ to the corresponding timestep value of the $\theta^{\mathrm{o2}}_{i,k}$ profile.

D. Primary Frequency Control Revenues & Costs

Using the results from [8], it is possible to extract the average hourly operating costs $C^{\mathrm{pfc}}_{i,k}$ of an ESU providing PFC. These operating costs consist of the cost of offsetting (i.e. the resource providing the offset compensation has to be paid) and the battery usage degradation. These values depend on the ESU characteristics, $E^{\mathrm{pfc}}_{i,k}$, and $P^{\mathrm{pfc}}_{i,k}$. We computed ESU operating costs over discretized values of $E^{\mathrm{pfc}}_{i,k}$ and $P^{\mathrm{pfc}}_{i,k}$ and found that these costs are almost a convex function in $E^{\mathrm{pfc}}_{i,k}$ and $P^{\mathrm{pfc}}_{i,k}$. We therefore use the set of linear constraints describing the convex hull to form a piecewise linear lower bound on PFC operating costs:

$$C^{\mathrm{pfc}}_{i,k} \geq a^{\mathrm{pfc},1}_{i,j} \cdot P^{\mathrm{pfc}}_{i,k} + b^{\mathrm{pfc},1}_{i,j} \cdot E^{\mathrm{pfc}}_{i,k} \;\; \forall\, i,j,k \tag{19}$$

where j indexes the planes constituting the piecewise affine approximation, and $a^{\mathrm{pfc},1}_{i,j}, b^{\mathrm{pfc},1}_{i,j} \geq 0 \,\forall\, i,j$. For a given $P^{\mathrm{pfc}}_{i,k}$, $C^{\mathrm{pfc}}_{i,k}$ decreases when $E^{\mathrm{pfc}}_{i,k}$ increases as the offset mechanism needs to be called less often, leading to lower degradation and smaller offsetting costs. We can describe the dependency of $P^{\text{e-pfc}}_{i,k}$ on $P^{\mathrm{pfc}}_{i,k}$ and $E^{\mathrm{pfc}}_{i,k}$ using the same approach:

$$P^{\text{e-pfc}}_{i,k} \leq a^{\mathrm{pfc},2}_{i,j} \cdot P^{\mathrm{pfc}}_{i,k} + b^{\mathrm{pfc},2}_{i,j} \cdot E^{\mathrm{pfc}}_{i,k} \;\; \forall\, i,j,k \tag{20}$$

where $a^{\mathrm{pfc},2}_{i,j}, b^{\mathrm{pfc},2}_{i,j} \geq 0 \,\forall\, i,j$. The ratio $P^{\text{e-pfc}}_{i,k}/P^{\mathrm{pfc}}_{i,k}$ is an increasing function of $E^{\mathrm{pfc}}_{i,k}/P^{\mathrm{pfc}}_{i,k}$ which tends to one when the later ratio tends to infinity (i.e. no offset mechanism needed if $E^{\mathrm{pfc}}_{i,k}$ is infinitely large).

We define the PFC profit $B^{\text{pfc}}_{i,k}$ as

$$B^{\text{pfc}}_{i,k} = P^{\text{e-pfc}}_{i,k} \cdot r^{\text{pfc}} - C^{\text{pfc}}_{i,k} \tag{21}$$

where r^{pfc} is the PFC hourly revenue per kW of effective PFC power capacity.

PFC markets in Europe require participants to provide constant PFC capacity over the duration of a contract (d^{pfc} timesteps). Values of d^{pfc} vary from 4h to a week. Therefore, the ESU aggregator should ensure that the sum of effective PFC capacities over all LACs remains constant over a PFC contract duration:

$$\sum_i P^{\text{e-pfc}}_{i,k^*} = \sum_i P^{\text{e-pfc}}_{i,k^*+k_2} \; \forall \, k_2 \in [1, d^{\text{pfc}} - 1] \tag{22}$$

where k^* represents the beginning of a PFC market interval. To avoid integer variables in the formulation, we assume that $\sum_i P^{\text{e-pfc}}_{i,k}$ can take any positive value and we do not include limits on the minimum power capacity provided.

E. MPC Controller Formulation

The MPC controller solves the following optimization problem:

$$\min_{\mathbf{u}_k} J(\mathbf{u}_k, x_k) = \sum_{j=0}^{N-1} \sum_i \Big(C^{\text{X}}_{i,k+j} + C^{\text{e}}_{i,k+j} + C^{\text{d}}_{i,k+j} + C^{\text{t}}_{i,k+j} - B^{\text{pfc}}_{i,k+j} \Big) + T_k \tag{23}$$

subject to (2)-(14), (17)-(22), and approximated versions of (15) and (16). The control input and system states are

$$\begin{aligned} \mathbf{u}_k &= [u_k, u_{k+1}, ..., u_{k+N\text{-}1}] \geq 0 \\ u_k &= [E^{\text{l}}_k, E^{\text{pfc}}_k, p^{\text{s,i}}_k, p^{\text{s,o}}_k, P^{\text{pfc}}_k, L^{\text{X}}_k, PV^{\text{X}}_k]^{\text{T}} \\ x_k &= [E^{\text{l}}_{k\text{-}1}, E^{\text{pfc}}_{k\text{-}1}, \theta^{\text{o2}}_{k\text{-}1}]^{\text{T}} \geq 0 \, . \end{aligned}$$

The variable N is the MPC optimization horizon and $T(k)$ is the terminal cost function representing the value of the energy within the ESUs at the end of the horizon:

$$T_k = \sum_i \bar{g}^{\text{in}}_{i,k} \cdot \Big(E^{\text{l}}_{i,k+N\text{-}1} + SoC^{\text{pfc}}_{\text{avg}} E^{\text{pfc}}_{i,k+N\text{-}1} \Big) . \tag{24}$$

The input data required by (23) are the load, PV, and temperature forecasts, and electricity tariffs. The optimal allocation schedule over the next k^{u} timesteps (control update rate) is applied to the ESUs and the process is repeated every k^{u} timesteps until the end of the simulation time.

III. Case Study Definitions

To demonstrate the behavior of our algorithm and highlight the importance of aggregation and PFC contract duration, we conduct two case studies. Both use four LACs, described in Table I, and the parameter values from Table II. We set the values in Table I to represent a situation where, in the absence of ESUs, each LAC faces a grid bottleneck during periods of either high load or high PV generation. Specifically, without ESUs, our simulations lead to 0.2% of load consumption curtailment for LAC 3, 18% of PV production curtailment for LACs 1 and 2, and to the hot spot temperature of the LAC 4 transformer reaching almost $240°\text{C}$. The ratings of the ESUs were set large enough to alleviate these bottlenecks, but we did not optimize over the ratings.

We use load and PV profiles covering 18 weeks – July to November 2012. LAC 1's load profile and all PV profiles come from measurements taken on a Swiss low voltage grid. LAC 2's load profile is a downscaled Switzerland-wide consumption profile (ENTSO-E data). The load profiles of LACs 3 and 4 come from synthetic residential and commercial profiles from Bavaria [19]. We added Brownian motion and white Gausian noise to the load profiles of LACs 2-4 to represent the increased variance of a small number of customers compared to that of the aggregate/average. We scaled the noise profiles so that they constitute up to about one fourth of the load profiles.

We use the electricity tariffs of residential and small commercial buildings from four Swiss providers. All the tariffs display the same general time-of-use structure: on-peak tariff from the morning to the evening Monday through Friday (sometimes as well during some hours on Saturday), and off-peak tariff the rest of the time. We modeled the feed-in tariff as half the retail rate for the corresponding time and LAC. The ratio $g^{\text{in}}_{i,k}/g^{\text{out}}_{i,k}$ approximately corresponds to the situation in Germany for new PV installations. To highlight the behavior of the algorithm, we set the battery degradation costs to lower values than what would be realistic today. Similarly, the transformer overheating cost was set to dissuade overheating.

We assume that we have perfect load and PV forecasts over N timesteps. Since our forecasts have limited horizons, we primarily use MPC to handle the new (perfect) forecasts that become available throughout the simulation period (in addition to transformer plant-model mismatch as described in Section II-C). Therefore, the results presented here constitute approximate upper bounds on the revenue achievable via aggregation and multitasking given the parameter values in Table I. More generally, the value of our case studies lies in the qualitative comparisons, not the quantitative results.

The first case study focuses on the additional profits generated by multitasking. We compare the total optimization objective value of four scenarios:

- *Scenario O (base case)*: We assume no ESUs are installed. The objective function is therefore penalized by PV and load curtailment as well as transformer overheating.
- *Scenario L (local services only)*: The ESUs are installed but only provide their local services.
- *Scenario D (local services and PFC)*: We assume the ESUs provide PFC (with a four-hour contract duration) in addition to their local services. This case corresponds to the framework detailed in the previous sections.
- *Scenario D2 (local services and PFC – different batteries)*: We assume two identical sets of ESUs (each identical to the set used in scenarios L and D). One set provides *only* local services to the four LACs, while the other set is installed in a location not facing a bottleneck and provides *only* PFC.

Comparing the results of the last two scenarios tells us whether

TABLE I
LAC PARAMETERS

LAC	1: PV-l	2: PV-t	3: Load-l	4: Load-t
Energy capacity E_i^{cap} [kWh]	800	900	600	600
Power capacity P_i^{cap} [kW]	150	200	180	150
Linear ESU degradation costs d_i^{l} [€/kWh]	0.03	0.05	0.10	0.10
Quadratic ESU degradation costs d_i^{q} [€/kW]	0.10	0.07	0.20	0.20
Min; Max; Mean load [kW]	144; 459; 288	150; 385; 264	53; 619; 269	72; 716; 300
PV rated power [kW]	1640	1760	-	-
On-; Off-peak rate [€/kWh]	0.161; 0.116	0.188; 0.105	0.185; 0.107	0.196; 0.123
Transformer rated power S_i [kW]	-	350	-	300
Transformer overheating cost h_i^{t} [€/kW^2]	-	0.05	-	0.05
Line limitations $p_i^{\text{g,i,max}}$; $p_i^{\text{g,o,max}}$ [kW]	-;450	-;-	470;-	-;-

TABLE II
SIMULATION PARAMETERS

Variables	Values
Timestep duration Δt [hours]	1
Control update rate k^{u} [timesteps]	168
MPC optimization horizon N [timesteps]	240
PFC hourly revenue r^{pfc} [€/(kW·h)]	0.016*

*current Swiss and German market price.

the simultaneous provision of local and grid services leads to a significant reduction in the quality of either service.

The second case study focuses on the importance of aggregation and establishes the relationship between aggregation and PFC contract duration. The following four scenarios are analyzed:

- *Scenario A*: This scenario is the most conservative. We assume a weekly PFC contract duration (which is the current regulation in Germany and Switzerland) and aggregation of ESUs is not allowed. Hence, each unit is operated independently and must provide a constant effective PFC power capacity. In this case, the summation signs of (22) disappear, and this equation is enforced separately for each LAC.
- *Scenario B*: We still assume a weekly PFC contract duration, but aggregation of ESUs is allowed, i.e. (22) applies. Hence, only the sum of the $P_{i,k}^{\text{e-pfc}}$ profiles is a commitment, and not the PFC schedule of each ESU.
- *Scenario C*: This scenario represents a much more flexible PFC market, with four-hour contract durations (which is the current regulation in Denmark), but aggregation is not allowed, as in Scenario A.
- *Scenario D*: This scenario is the same as the scenario D from the first case study. It represents the most innovative scenario, combining the improvements of Scenarios B and C: four-hour PFC contract duration and aggregation of ESUs.

TABLE III
RESULTS OF FIRST CASE STUDY. ALL VALUES ARE IN K€.

Scenario	O	L	D	D2
C^{e}	351	340	343	340 + 0
C^{d}	0	11	8	11 + 0
C^{X}	13	0	0	0 + 0
C^{t}	15	2	2	2 + 0
$-B^{\text{pfc}}$	0	0	-18	0 + -21
Objective Value	380	352	335	331

IV. RESULTS

Table III shows the results of the first case study. The values in the first four rows are the sums of the objective value components over the LACs and the simulation period. For D2, the values on each side of the + sign correspond to the different ESU sets. The table shows that scenario D leads to a gain of 17k€ compared to scenario L. This gain comes from the provision of PFC in addition to the local services. It also shows that the difference between scenarios D and D2 is only 4k€, which means that there is little conflict between PFC and the local services. Since the battery investment cost in scenario D2 is exactly twice the one in scenario D, this result is a strong argument for multitasking.

Table IV shows the results of the second case study. The additional benefit of aggregation is bigger when considering weekly PFC contract durations (2.0k€ from scenario A to B) than when considering four-hour contract durations (0.3k€ from C to D). This is because the shorter the PFC contract, the lesser the impact of the needs of the local service on the optimal PFC schedule. By extrapolation, using one-hour PFC contract duration would result in no difference between aggregation and no aggregation, since (22) would be trivially satisfied. It also shows that a four-hour PFC market, whose implementation is the system operator's responsibility, leads to better results than an aggregation scheme, whose implementation is the responsibility of the aggregator, under a weekly PFC market. The importance of a short PFC contract duration for multitasking was also highlighted in [3].

The results show that even the most conservative scenario (A) already achieves a significant portion of the improvement of the most advanced scenario (D), i.e. the difference between scenarios D and A (3.7k€) is more than three times smaller than the difference between scenarios A and L (13.6k€). The small difference between scenarios D and A comes from the fact that we selected local services that are not often requested (but that provide a high revenue per kWh cycled). Performing the same analysis with ESUs whose local applications run more often (such as electric vehicle batteries or thermostatically controlled loads) would likely lead to a different result. Specifically, aggregation would be critical to ensure a constant availability of PFC power capacity over some hours.

When analyzing the results of Tables III and IV, one should keep in mind that the biggest component of the operating costs is the cost of energy exchange with the grid (C^{e}), on which the ESUs have only limited influence. This explains why the difference between scenarios might seem small when expressed

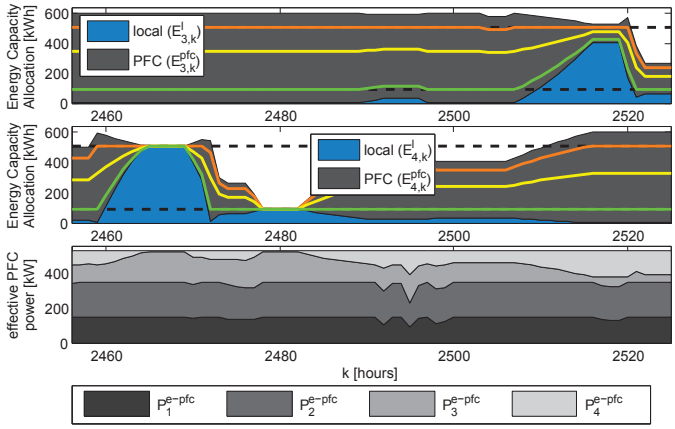

Fig. 5. Illustration of how the central controller allocates capacities in order to provide constant PFC capacity over a PFC contract duration, Scenario B. The top and middle plots display energy capacity allocation for LACs 3 and 4. The energy allocation of LAC 1 and 2 is omitted here since it does not bring additional information because their local services were barely requested. The bottom plot shows the individual contribution of each LAC to the cumulative PFC power capacity.

TABLE IV
RESULTS OF THE SECOND CASE STUDY. ALL VALUES ARE IN K€.

Scenario	A	B	C	D
Objective Value	338.4	336.4	334.9	334.7

in % of the operating costs of scenario O.

Figure 5 shows how the central scheduler dispatches the four LACs in scenario B to ensure constant PFC power capacity over a week. Since LACs 3 and 4 face demand for their local services at a specific time during this week, the optimal cumulative effective PFC power capacity is only 530kW for the whole week, whereas it climbs to 610kW for weeks with less request from local services.

V. CONCLUSION

We have presented an algorithm that co-optimizes simultaneous provision of local and power system services by an aggregation of energy storage units. We showed that multitasking can increase the profitability of batteries in power systems, especially if their local services are related to overloading and are requested rarely. In such cases, we find that the profits generated by a battery set providing both local and system services are almost equal to the sum of the profits from two identical battery sets, where one provides only the local service and the other only the system service. We also found that the importance of aggregation decreases when the PFC contract duration decreases. Finally, we showed that improvements generated by multitasking are significantly greater than improvements resulting from shorter PFC contract duration or from ESUs aggregation for PFC provision.

Future research includes extending our framework to consider other kinds of resources as well as the impact of prediction uncertainty and stochasticity. The present work was performed assuming perfect forecasts over the optimization horizon, and we expect the importance of aggregation will increase if we consider uncertainty, as aggregation will reduce the noise on the prediction. Furthermore, we will conduct sensitivity analyses to assess the importance of specific parameters, especially battery degradation costs.

ACKNOWLEDGMENT

The authors would like to thank the Swiss Commission for Technology and Innovation (project no. 14478).

REFERENCES

[1] McKinsey & Company, "A portfolio of power-trains for Europe: a fact-based analysis," 2010.
[2] J. Eyer and G. Corey, "Energy storage for the electricity grid: Benefits and market potential assessment guide," SAND2010-0815, Sandia National Laboratories, Tech. Rep., 2010.
[3] B. Wasowicz, S. Koopmann, T. Dederichs, A. Schnettler, and U. Spaetling, "Evaluating regulatory and market frameworks for energy storage deployment in electricity grids with high renewable energy penetration," in *9th European Energy Market International Conference*, may 2012.
[4] M. Braun and T. Stetz, "Multifunctional photovoltaic inverters–economic potential of grid-connected multifunctional PV-battery-systems in industrial environments," in *23rd European Photovoltaic Solar Energy Conference, Valencia, Spain*, 2008.
[5] M. G. Vayá and G. Andersson, "Combined smart-charging and frequency regulation for fleets of plug-in electric vehicles," in *IEEE Power and Energy Society General Meeting*, 2013.
[6] J. Kondoh, N. Lu, and D. Hammerstrom, "An evaluation of the water heater load potential for providing regulation service," in *Power and Energy Society General Meeting, IEEE*, 2011.
[7] E. Vrettos and G. Andersson, "Combined load frequency control and active distribution network management with thermostatically controlled loads," in *IEEE International Conference on Smart Grid Communications, Vancouver, Canada*, 2013.
[8] O. Mégel, J. L. Mathieu, and G. Andersson, "Maximizing the potential of energy storage to provide fast frequency control," *4th IEEE Innovative Smart Grid Technologies Europe, Copenhagen*, October 2013.
[9] E. F. Camacho and C. Bordons, *Model predictive control*, 2nd ed. Springer London, 2007, vol. 2.
[10] J. Löfberg, "YALMIP : A toolbox for modeling and optimization in MATLAB," in *Proceedings of the CACSD Conference*, Taipei, Taiwan, 2004. [Online]. Available: http://users.isy.liu.se/johanl/yalmip
[11] D. Tretheway, "Regulation energy management draft final proposal," CAISO, Tech. Rep., 2011.
[12] T. Borsche, A. Ulbig, M. Koller, and G. Andersson, "Power and energy capacity requirements of storages providing frequency control reserves," in *Proceedings of the IEEE PES General Meeting*, 2013.
[13] C. Jin, N. Lu, S. Lu, Y. Makarov, and R. Dougal, "Coordinated control algorithm for hybrid energy storage systems," in *Proceedings of the IEEE PES General Meeting*, 2011, pp. 1–7.
[14] A. Oudalov, D. Chartouni, and C. Ohler, "Optimizing a battery energy storage system for primary frequency control," *Power Systems, IEEE Transactions on*, vol. 22, no. 3, pp. 1259 –1266, 2007.
[15] A. van der Welle and B. van der Zwaan, "An overview of selected studies on the value of lost load (VOLL)," *Energy Research Centre of the Netherlands (ECN)*, 2007.
[16] International Electrotechnical Commission, "Power transformers - part 7: Loading guide for oil-immersed power transformers," 2005, iEC 60076-7:2005.
[17] M. C. Kuss, A. J. Markel, and W. E. Kramer, "Application of distribution transformer thermal life models to electrified vehicle charging loads using monte-carlo method," *25th World Battery, Hybrid and Fuel Cell Electric Vehicle Symposium & Exhibition, Shenzhen, China*, 2010.
[18] C. Tzanetopoulou, "Impacts of high penetration of pv in distribution grids and mitigation strategies," Master Thesis, ETH Zurich, 2013.
[19] E.ON Mitte AG. (2012). [Online]. Available: http://www.eon-mitte.com/de/netz/veroeffentlichungen/strom_/standardlastprofil_verfahren/standardlastprofile_2012