A STOCHASTIC APPROXIMATION APPROACH TO LOAD SHEDDING IN POWER NETWORKS

Nikolaos Gatsis

The University of Texas at San Antonio
Dept. of Electrical and Computer Engineering
Email: nikolaos.gatsis@utsa.edu

Antonio G. Marques

King Juan Carlos University
Dept. of Signal Theory and Communications
Email: antonio.garcia.marques@urjc.es

ABSTRACT

A system comprising a utility company serving a set of electricity end-users is considered. The utility company can purchase energy from the wholesale market. It is also connected to a renewable energy production facility, from which it can harvest energy at no cost, and also to a battery for energy storage. Ahead of a scheduling horizon, the utility purchases energy based on forecasted demand and renewable energy production. During online operation, if the renewable energy is not adequate, real-time decisions with respect to user load shedding, energy procurement, and battery charging or discharging need to be made. The problem is cast in a stochastic approximation framework, and is solved online via a dual stochastic subgradient method with low per-slot complexity.

Index Terms— Dual stochastic subgradients, load shedding, smart grid, stochastic approximation.

1. INTRODUCTION

Environmental concerns and requirements on reduced dependency on non-replenishable resources such as oil and coal are driving the high integration of renewable energy into modern power systems across the world. The main challenge with such integration is the stochastic nature of renewable energy, which makes it largely unpredictable in time scales greater than a few hours.

This paper aspires to tackle this challenge for a system comprising a utility company serving a set of residential users. The utility company relies on energy procurement from the wholesale market to serve its users, but also has a renewable energy production facility that incurs zero cost to the utility, as well as a battery to store energy. The research issue addressed pertains to the situation where the produced renewable energy is not enough to cover the total load, and decisions have to be made with respect to real-time energy procurement at unpredictable real-time prices versus shedding users' load, which both incur costs to the utility company. The problem is formulated in a stochastic approximation framework. Relying on computationally efficient dual stochastic subgradient iterations, a novel online algorithm for load shedding, energy procurement, and battery charging or discharging is developed. The decisions are made in a fashion adaptive to the renewable energy production at every slot.

1.1. Prior art and paper contributions

Load shedding has been considered in [1, Ch. 11] and [2] in a static manner that does not adapt to the renewable energy uncertainty. More recent works place the interplay between renewable energy production, storage, and user demand in a stochastic setting [3–

8]. Incorporating battery dynamics naturally leads to finite-horizon stochastic dynamic programming (DP) approaches. To render the problem tractable in such setups, unrealistic probabilistic modeling assumptions for the renewable energy are adopted [3, 4]. For instance, renewable energy production is modeled as Gaussian, which is not quite realistic; see e.g., [9, Ch. 3], [10], and [11] for more pertinent models. In general, intricate approximations may be employed to solve related stochastic DP problems as in [5], which have high complexity and do not lend themselves to easily implementable online algorithms. The algorithm developed here in contrast bypasses the DP limitations, has minimal per step complexity, and is not tied to specific probabilistic models for the uncertainty.

Infinite-horizon stochastic DP is employed in [6] to coordinate renewable energy production and demand, where users have power requests whose time durations are i.i.d. and exponentially distributed. In the present framework, no explicit modeling of the users' power demands is needed. Lyapunov-based stochastic network optimization techniques are used in [7] and [8]. In the former, the focus is on users whose power consumption can be deferred for later if the utility company cannot cover the demand, while user consumption is curtailed instead in the present work. An algorithm for user consumption adjustments is also pursued in [8], where deterministic bounds need to be known for all uncertainties in advance in order to run the algorithm, while these bounds pose explicit constraints on the storage size. No such requirements are present in the stochastic approximation framework of this paper.

The remainder of the paper is organized as follows. Section 2 presents the problem formulation, while Section 3 gives characterizations of the optimal solution. The online solver is developed in Section 4, followed by numerical tests in Section 5.

2. PROBLEM FORMULATION

Consider a set of electricity end-users, denoted as $\{1, \ldots, K\}$, all served by the same utility company (also called load-serving entity, LSE). The LSE has three sources of energy. Specifically, it is connected to a renewable energy production facility, from which it can draw power to provide to the end-users. This renewable energy source does not incur any cost to the LSE. The LSE may also buy energy from a wholesale supplier (or the market). The third source is a battery used by the LSE to store the excess of energy.

The LSE makes a forecast of both the user electricity demand and the renewable energy production ahead of a system operation horizon. Suppose the horizon comprises a set of slots denoted as $\{1, \ldots, T\}$. For instance, the horizon may be a day, consisting of $T = 144$ 10-minute intervals. The LSE procures from the wholesale supplier energy equal to the forecasted demand offset by the

forecasted renewable energy production ahead of the horizon.

Now turning attention to the system operation within the horizon, if the actual renewable energy production is less than its forecasted value, the LSE has three options, namely, 1) to shed user load, 2) to buy energy in real time from the wholesale supplier, and 3) to use energy stored in the battery. Specifically, let π^t and w^t respectively denote the actual demand minus the procured energy at slot t and the produced renewable energy at slot t. The LSE can purchase energy in real time with price a^t in slot t. The real-time prices are not known in advance, and therefore, they are modeled as random.

The decision variables for LSE are denoted as s_k^t, b^t, and e_{out}^t, which respectively stand for the amount of load shedded for user k at slot t, the amount of energy bought at slot t, and the amount of energy drawn from the battery at slot t. These must satisfy per slot t

$$\pi^t - w^t \leq \sum_{k=1}^{K} s_k^t + b^t + \eta_{\text{dis}} e_{\text{out}}^t \qquad (1)$$

where $0 < \eta_{\text{dis}} < 1$ is the battery discharging efficiency, indicating the fraction of the extracted energy that can actually be provided to the load connected to the battery. The LSE can purchase up to $b^{\max}$ amount of energy per slot, which is dictated by the transfer capability of the LSE's connection with the main grid, that is, it must hold that $b^t \leq b^{\max}$ for all t. In addition, there is a maximum limit on the amount of load shedding per user, expressed as $s_k^t \leq s_k^{\max}$.

The aim is to determine s_k^t, b^t, and e_{out}^t per user and slot. These will be selected based on the actual renewable energy production w^t and the real time price a^t, which are modeled as random. Note that the actual demand may also differ from the forecasted one, in which case, π^t may also be modeled as random. This assumption is not made here, because in general, load forecasting methods are more accurate than renewable energy forecasting methods, and also because the ensuing stochastic approximation framework can easily handle random π^t with minor modifications.

The battery dynamics are described next. Specifically, let r^t be the amount of energy stored in the battery at the end of slot t. Let e_{in}^t be the amount of energy that is stored into the battery at slot t. The quantities r^t, e_{in}^t, and e_{out}^t are related as follows:

$$r^t = r^{t-1} + e_{\text{in}}^t - e_{\text{out}}^t, \ t = 1, \ldots, T \qquad (2a)$$

$$0 \leq r^t \leq R, \ t = 1, \ldots, T \qquad (2b)$$

where r^0 is given, and R is the battery capacity.

Suppose that the battery is charged only when there is excess of renewable energy (in particular, the LSE does not procure energy in real time to charge the battery, but only to avoid user load shedding). The following must then hold for e_{in}^t and e_{out}^t for all t:

$$e_{\text{in}}^t = \eta_{\text{ch}} \min\{e_{\text{in}}^{\max}, [\pi^t - w^t]^-\}; \ 0 \leq e_{\text{out}}^t \leq e_{\text{out}}^{\max} \qquad (3a)$$

where η_{ch} is the charging efficiency, which is the fraction of the energy supplied to the battery that can actually be stored in the battery, and $[x]^- := \max\{-x, 0\}$.

The customer will typically be compensated by the LSE for the inconvenience caused due to load shedding. In general, this compensation will depend on the load shedded during the entire horizon for each user. In what follows, a cost function is set up to guide the load shedding decisions per user and slot. Upon defining $\bar{s}_k = \frac{1}{T}\sum_{t=1}^{T} s_k^t$, the shedding decisions can be made by penalizing $\bar{s}_k$, which is proportional to the total load shedded, through a cost function $J_k(\bar{s}_k)$. The cost function serves two purposes, as described next.

1. *User compensation:* The user is compensated for the inconvenience due to load shedding, and this compensation is proportional to the total load shedded.

2. *User fairness:* A proper choice for the shape of $J_k(\bar{s}_k)$ can ensure that no user suffers from excessive load shedding relative to other users. This situation is analogous to rate fairness in communication networks—see e.g., [12].

The cost functions $J_k(\bar{s}_k)$ are selected to be strictly increasing and convex (increasing marginal cost).

All the pieces are in place now in order to present the optimization formulation. The problem can be stated with finite horizon T, which would necessitate the use of dynamic DP. This approach is computationally challenging. On the other hand, the infinite-horizon version will afford very efficient online solvers based on stochastic approximation techniques. The formulation is presented next. The random processes w^t and a^t are assumed to be stationary and ergodic. This is supported by different works in the literature, see e.g., [9, Ch. 3], [10], [11] for the case of wind energy.

$$\min_{\{s_k^t\}, \{b^t\}, \{e_{\text{out}}^t\}, \{\bar{s}_k\}} \ \sum_{k=1}^{K} J_k(\bar{s}_k) + \lim_{T\to\infty} \frac{1}{T}\sum_{t=1}^{T} a^t b^t \qquad (4a)$$

$$\text{subj. to} \quad \lim_{T\to\infty} \frac{1}{T}\sum_{t=1}^{T} s_k^t \leq \bar{s}_k, k = 1, \ldots, K \qquad (4b)$$

$$\lim_{T\to\infty} \frac{1}{T}\sum_{t=1}^{T} e_{\text{out}}^t = \lim_{T\to\infty} \frac{1}{T}\sum_{t=1}^{T} e_{\text{in}}^t \qquad (4c)$$

$$(1), (3), 0 \leq b^t \leq b^{\max}, \ t = 1, 2, \ldots \qquad (4d)$$

$$0 \leq s_k^t \leq s_k^{\max}, \ t = 1, \ldots, k = 1, \ldots, K \qquad (4e)$$

The objective consists of two parts, namely, the penalty related to load shedding and the energy procurement cost. The variables are the instantaneous decisions as well as the auxiliary variables $\{\bar{s}_k\}$. Note that (4b) will hold as equality at optimality, because J_k is strictly increasing.

Constraint (4c) guarantees that, in the long term, the energy stored in the battery is the same than that taken from the battery. Clearly, (4c) represents a relaxation of the battery dynamical equation (2) that is only tight if the limits of the battery are never reached. To better motivate the constraint, note that (2a) and $r^t \geq 0$ are equivalent to the condition $0 \leq \sum_{t=1}^{t'}(e_{\text{in}}^t - e_{\text{out}}^t) + r^0$ for $t' = 1, \ldots, T$. Likewise, (2a) and $r^t \leq R$ are equivalent to the condition $\sum_{t=1}^{t'}(e_{\text{in}}^t - e_{\text{out}}^t) + r^0 \leq R$ for $t' = 1, \ldots, T$. Now, substituting $t' = T$ into the previous equations, dividing by T, and taking $T \to \infty$, (4c) follows. A similar constraint has been previously used for batteries in energy-harvesting networks [13].

The optimal solution of (4) is characterized in the next section, while the online solver is developed in Section 4.

3. OPTIMAL SOLUTION

The objective in (4) is convex while the constraints are linear, so the problem is convex. The solution approach is to dualize the long-term constraints and leverage the problem separability. To be specific, let σ_k denote the Lagrange multiplier associated with the kth constraint in (4b) and ρ the multiplier associated with the long-term power conservation constraint in (4c). The Lagrangian is then formed by augmenting the objective with the dualized constraints. Provided that

the value of $\boldsymbol{\lambda} := [\sigma_1, ..., \sigma_K, \rho]$ is known, the following proposition describes the solution that minimizes the Lagrangian.

Proposition 1.a): The optimum solution for the average variables satisfies $\tilde{s}_k^*(\boldsymbol{\lambda}) = \arg\min\{J_k(s) - \sigma_k s\}, \forall k$.

Proposition 1.b): The optimal instantaneous allocation at time t follows a greedy strategy:

b.1) If $\pi^t - w^t \leq 0$, there is an instantaneous energy surplus, which is stored in the battery, that is, $e_{\text{in}}^t = \eta_{\text{dis}} \min\{e_{\text{in}}^{\max}, w^t - \pi^t\}$, while $e_{\text{out}}^{t*}(\boldsymbol{\lambda}) = 0$, $b^{t*}(\boldsymbol{\lambda}) = 0$ and $s_k^{t*}(\boldsymbol{\lambda}) = 0\ \forall k$.

b.2) If $\pi^t - w^t > 0$, there is an instantaneous energy deficit, which is covered by activating at least one of the three resources available, while $e_{\text{in}}^t = 0$. Specifically, the optimal $e_{\text{out}}^{t*}(\boldsymbol{\lambda})$, $b^{t*}(\boldsymbol{\lambda})$, and $s_k^{t*}(\boldsymbol{\lambda})$ are found as the solution of the following linear program (LP):

$$\min_{\substack{\{0 \leq s_k^t \leq s_k^{\max}, 0 \leq b^t \leq b^{\max}, \\ 0 \leq e_{\text{out}}^t \leq e_{\text{out}}^{\max}\}}} a^t b^t + \rho e_{\text{out}}^t + \sum_{k=1}^{K} \sigma_k s_k^t \quad (5a)$$

$$\text{subj. to:} \quad \pi^t - w^t = b^t + \eta_{\text{dis}} e_{\text{out}}^t + \sum_{k=1}^{K} s_k^t. \quad (5b)$$

The first observation is that, thanks to the decomposability of the Lagrangian, average and instantaneous variables can be found separately. Note that the variables $\tilde{s}_k^{t*}$ in 1.a) are found separately for each user k too. To gain some intuition on the solution, we will interpret σ_k as the price of shedding user k and ρ as the price of extracting energy from the battery. Part a) of the proposition states that the optimal $\tilde{s}_k^*$ is found so that, in the long term, the marginal cost of shedding user k coincides with the equilibrium price σ_k. In fact, if $J_k(\cdot)$ is differentiable, one can readily see that the first-order optimality condition yields $J_k'(\tilde{s}_k^*) = \sigma_k$. Regarding the instantaneous variables, part 1.b.2) of the proposition dictates that, when a deficit of instantaneous energy exists, a greedy approach is optimal. More specifically, while the deficit can be covered by any combination of $\{e_{\text{out}}^{t*}, b^{t*}, s_1^{t*}, ..., s_K^{t*}\}$, the LP in 1.b.2) implies that, if feasible, only the resource with minimum price shall be used to cover it, while all others must be set to zero. When the deficit is so high that it exceeds the maximum value of the resource with minimum price, the resource with the second smallest price has to be activated to cover the remaining deficit, and so on.

When more than one resource achieve the minimum price (e.g., when $\rho = a_t$ and $\sigma_k > \rho$ for all k), any feasible combination of the resources with minimum price (e_{out}^t and b^t in our example) covering the deficit minimizes the Lagrangian. However, not all such combinations will necessarily be actual (feasible) solutions of the original primal problem (4). This is a well-known issue with the dual methods for not strictly convex problems. In fact, it is not difficult to show that, except in trivial cases, the optimal solution to our problem requires $\sigma_k^* = \sigma_{k'}^* = \rho^*$ for any two users k and k', so the aforementioned issue will show up. Different approaches exist in the literature to solve this problem; see e.g., [14–17]. Leveraging the results in [15], it is possible to construct a solution of (5) that guarantees feasibility in the primal domain.[1] However, it is worth stressing that the focus of the paper is on the design of stochastic schemes that overcome this issue, as will be explained in the ensuing section.

---

[1] The specific combination is $s_k^{t*} = (\pi^t - w^t)\tilde{s}_k^*/\Omega_0$, $e_{\text{out}}^{t*} = (\pi^t - w^t)\bar{e}_{\text{out}}/(\eta_{\text{dis}}\Omega_0)$, and $b^{t*} = (\pi^t - w^t)\bar{a}_{eq}/\Omega_0$, where $\bar{e}_{\text{out}} = \eta_{\text{ch}}\eta_{\text{dis}}\mathsf{E}[(w^t - \pi^t)\mathbf{1}_{\{w^t - \pi^t \geq 0\}}]$, $\bar{a}_{eq} = \mathsf{E}[(\pi^t - w^t)\mathbf{1}_{\{(\pi^t - w^t) > 0\}}] - \mathsf{E}[(\pi^t - w^t)\mathbf{1}_{\{a^t < \lambda^*\}}] - \bar{e}_{\text{out}} - \sum_k \tilde{s}_k^*$ and $\Omega_0 := \mathsf{E}[(\pi^t - w^t)\mathbf{1}_{\{(\pi^t - w^t) > 0\}}\mathbf{1}_{\{a_t \geq \rho^*\}}] = \bar{e}_{\text{out}} + \bar{a}_{eq} + \sum_k \tilde{s}_k^*$. One can easily prove that this allocation is optimal and feasible.

4. ALGORITHM

4.1. Stochastic approximation

The main obstacle in implementing the solution in the previous section is to find the optimal value for the multipliers $\{\sigma_k^*, \rho^*\}$. Although a classical offline iterative subgradient dual method can be used to accomplish such a task, averaging over all time instants is required at each iteration; see e.g., [16]. The latter is challenging because 1) knowledge of the joint distribution of $\{w_t, a_t\}$, or, for finite horizons, the entire path $\{w_t, a_t\}_{t=1}^T$, would be required and 2) the algorithm will be computationally expensive. An effective alternative to bypass such problems entails stochastic approximation iterations [18–20], whose goal is to obtain samples $\{\sigma_k^t, \rho^t\}, t = 1, 2, \ldots$ that are sufficiently close to the optimal dual variables. The merits of stochastic approximation techniques are threefold: 1) the joint distribution of $\{w_t, a_t\}$ is not required; 2) the computational complexity of stochastic approximation schemes is significantly lower than that of their off-line counterparts; and 3) stochastic schemes can cope with non-stationary environments. In fact, stochastic schemes are causal, meaning that at time t, only the history of the system up to time t, i.e., $\{w_\tau, a_\tau\}_{\tau=1}^t$, is required.

With $\mu_\sigma > 0$ and $\mu_\rho > 0$ denoting constant stepsizes, the following stochastic iterations yield the desired multipliers for all t:

$$\rho^{t+1} = \left[\rho^t - \mu_\rho\left(\eta_{\text{ch}} e_{\text{in}}^{t*}(\boldsymbol{\lambda}^t) - e_{\text{out}}^{t*}(\boldsymbol{\lambda}^t)\right)\right]^+ \quad (6)$$

$$\sigma_k^{t+1} = \left[\sigma_k^t + \mu_\sigma\left(\tilde{s}_k(\sigma_k^t) - s_k^{t*}(\boldsymbol{\lambda}^t)\right)\right]^+ \quad (7)$$

where $[x]^+ = \max\{x, 0\}$. The update terms in the right-hand sides of (6) and (7) form an *unbiased* stochastic subgradient of the dual function of (4), and they are *bounded* almost surely; see e.g., [14, 20]. It can be shown that the sample average of the solution of the LP using the iterates $\rho^t, \{\sigma_k^t\}_{k=1}^K$ instead of $\boldsymbol{\lambda}$ satisfies the long-term constraints in (4), and incurs minimal performance loss relative to the optimal (off-line) solution of (4). Rigorously stated, define $\mu := \max\{\mu_\sigma, \mu_\rho\}$; $\bar{s}_k^t := \frac{1}{t}\sum_{\tau=1}^t s_k^{\tau*}(\boldsymbol{\lambda}^\tau)$; $\bar{J}^t := \sum_k J_k(\bar{s}_k^t) + \frac{1}{t}\sum_{\tau=1}^t a^\tau b^{\tau*}(\boldsymbol{\lambda}^\tau)$ and J^* as the optimal value of (4). Then, it holds with probability one that as $t \to \infty$, 1) the solution is feasible; and 2) $\bar{J}^t \leq J^* + \delta(\mu)$, where $\delta(\mu) \to 0$ as $\mu \to 0$. A proof of this result is not presented here due to space limitations, but it relies on the convergence of stochastic (epsilon) subgradient methods and can be derived following the lines of [20, 21].

In the context of energy-harvesting sensor networks, it has been recently shown that an alternative way to estimate ρ^* stochastically is to replace (6) with (cf. [13])

$$\rho^t = [C - \mu_\rho r^t]^+ \quad (8)$$

where C is a design constant to be selected based on the battery capacity and r_t is updated recursively according to (2). In words, a scaled version of the battery level r^t can be used as an estimate of ρ^*. Since ρ^* can be viewed as the price of the energy stored in the battery, (8) states that the estimate of the stored energy price decreases as the amount of energy actually stored in the battery increases.

4.2. Resource allocation algorithm

With all previous considerations in mind, Algorithm 1 is the on-line solver for (4). The algorithm implements the solution given in Proposition 1 after replacing σ_k and ρ with the stochastic estimates σ_k^t and ρ^t in (7) and (8), respectively, with two modifications described next.

Algorithm 1 Load shedding with procurement and storage

For each k, initialize σ_k^1 to a small random number.

for $t = 1, \ldots, T$ **do**

    [s1] Find $\tilde{s}_k^*(\sigma_k^t)$ for all k using Prop. 1.a.

    [s2.0] Set the value of ρ^t using (8).

    [s2.1] If $\pi^t - w^t \leq 0$, then find e_{in}^t, e_{out}^{t*}, b^{t*} and s_k^{t*} using Prop. 1.b.1 with $\sigma_k = \sigma_k^t$ and $\rho = \rho^t$. Go to **[s3]**.

    [s2.2.1] If $\pi^t - w^t > 0$, then set $e_{\text{in}}^t = 0$ and solve the LP in Prop. 1.b.2 with $\sigma_k = \sigma_k^t$ and $\rho = \rho^t$ and replacing $e_{\text{out}}^{\max}$ with $\tilde{e}_{\text{out}}^{\max} := \min\{e_{\text{out}}^{\max}, r^t\}$. If multiple solutions exist, pick one at random. Go to **[s3]**.

    [s3] Using the outputs of steps **[s1]** and **[s2]**, update the battery r^t via (2) and σ_k^t via (7).

end for
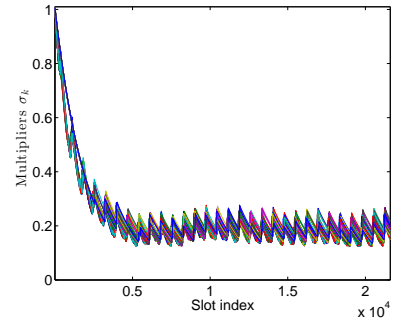

Fig. 1. Convergence of Lagrange multipliers.

First, to account for the battery limits, the actual value of r^t is used in the algorithm (cf. Step [s2.0]) and the upper bound $e_{\text{out}}^{\max}$ is replaced by $\tilde{e}_{\text{out}}^{\max}$ (cf. Step [s2.2.1]). This way, we guarantee that the solution is feasible not only for (4), but also over a finite horizon. Note however that by properly choosing the constants in (8), re-defining the bound $\tilde{e}_{\text{out}}^{\max}$ is not necessary. To illustrate this, suppose that C in (8) is set to a very high value. Then, the price of e_{out}^t when $r^t \to 0$ will be very high, so that it is very unlikely for the the solution to set $e_{\text{out}}^t > 0$ and, hence, to deplete the battery.

The second modification deals with the event of two resources corresponding to the minimum price (i.e., multiple solutions to the LP in Prop. 1.b.2). In such a hypothetical event, the algorithm just picks one of them at random. Intuitively, such an event will not affect the long-term feasibility, because small deviations from the optimal solution that could potentially render the long-term constraints infeasible, are immediately compensated via the update of the multipliers associated with those constraints. In addition, since the values of the multipliers are random, the event that two or more of them taking the exact same value is highly unlikely—especially when the number of values that the random variables can take grows large.

5. NUMERICAL TESTS

A system with $K = 50$ residential users is considered. The horizon consists of $T = 21600$ 2-minute intervals over 30 days. There are two classes of users with cost functions $J_k(s_k) = 0.5s_k^2$ for $1 \leq k \leq 25$, and $J_k(s_k) = s_k^2$ for $26 \leq k \leq 50$. The quantity π^t follows a daily pattern peaking in the early evening at 300 kW, and having lowest values in late night hours—see [17, Fig. 4] and [22, Ch. 2] for examples. The real-time prices are selected i.i.d. uniform over the interval $(0, 5)$ cents/kW. The system also has installed wind capacity of 130kW. The samples w^t are generated i.i.d., using the sampling method and parameters detailed in [23]. The battery capacity is $R = 50$kWh. Instantaneous bounds on the decision variables were selected so that they are not binding, in order to simplify the presentation. The online algorithm of Table 1 is used with stepsizes $\mu_\rho = 0.001$ and $\mu_\sigma = 0.5$, and battery parameter $C = 10$.

The convergence of the Lagrange multipliers σ_k is depicted in Fig. 1, where it is seen that they all converge to the same value, as it is expected for this system. The energy shortage and the load shedding, real-time energy procurement, and battery charging/discharging decisions are illustrated in Fig. 2. It is interesting to note that the peaks of load shedding coincide with the shortage peaks. But not all shortage is covered by load shedding; in fact the battery is also discharged, while buying energy happens less frequently. The rea-

son for less frequent energy procurement here is that the prices a^t are generally higher than the user marginal costs—that is, the values of the derivatives $J_k'(s)$ at the optimal point.

The latter are shown in the bottom plot of Fig. 3, while the top part illustrates the running averages $\bar{s}_k$ of the load shedding decisions. Note that the derivatives $J_k'(\bar{s}_k)$ have roughly the same values as the Lagrange multipliers (cf. Fig. 1), which is another manifestation of the proper algorithm convergence. It is finally interesting to observe how the different choice of user cost functions leads to different load shedding decisions. Specifically, Class 1 users have double load shedding in the long run than Class 2 users. This is explained by the fact that user cost derivatives must be equal across all users, and these derivatives are linear in the case of quadratic cost functions.
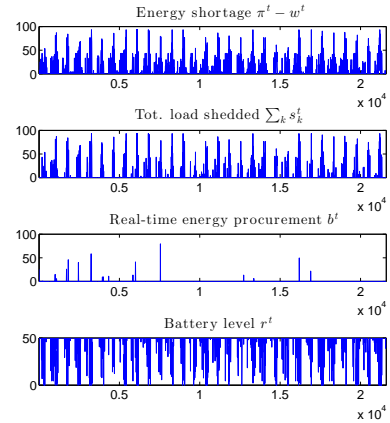
Fig. 2. Real-time decisions. The x-axis is slot index.
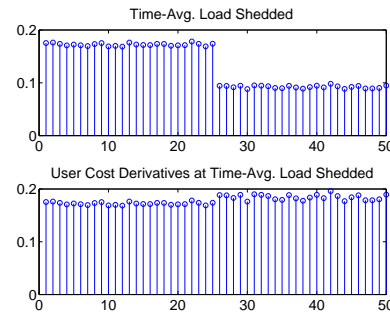
Fig. 3. Time-average load shedded $\bar{s}_k = \frac{1}{T}\sum_{t=1}^{T} s_k^t$ and user cost derivatives $J_k'(s)$ evaluated at $\bar{s}_k$. The x-axis is user index.

6. REFERENCES

[1] J. Zhu, *Optimization of power system operation*, Wiley-IEEE, Hoboken, NJ, 2009.

[2] D. Xu and A.A. Girgis, "Optimal load shedding strategy in power systems with distributed generation," in *Proc. IEEE Power Engineering Society Winter Meeting*, Columbus, OH, Jan. 2001, vol. 2, pp. 788–793.

[3] L. Jiang and S. H. Low, "Real-time demand response with uncertain renewable energy in smart grid," in *Proc. 49th Allerton Conf. Communications, Control, and Computing*, Monticello, IL, Sept. 2011, pp. 1334–1341.

[4] M. He, S. Murugesan, and J. Zhang, "Multiple timescale dispatch and scheduling for stochastic reliability in smart grids with wind generation integration," in *Proc. IEEE INFOCOM*, Shanghai, China, Apr. 2011, pp. 461–465.

[5] A. Papavasiliou and S. S. Oren, "Supplying renewable energy to deferrable loads: Algorithms and economic analysis," in *Proc. IEEE Power and Energy Society General Meeting*, Minneapolis, MN, July 2010, pp. 1–8.

[6] I. Koutsopoulos, V. Hatzi, and L. Tassiulas, "Optimal energy and storage control policies for the smart power grid," in *Proc. 2nd IEEE Int. Conf. Smart Grid Communications*, Brussels, Belgium, Oct. 2011, pp. 475–480.

[7] M. J. Neely, A. S. Tehrani, and A. G. Dimakis, "Efficient algorithms for renewable energy allocation to delay tolerant consumers," in *Proc. 1st Int. Conf. Smart Grid Communications*, Gaithersburg, MD, Oct. 2010, pp. 549–554.

[8] L. Huang, J. Walran, and K. Ramchandran, "Optimal demand response with energy storage management," May 2012, http://arxiv.org/abs/1205.4297/.

[9] A. Conejo, M. Carrion, and J. M. Morales, *Decision Making Under Uncertainty in Electricity Markets*, Springer, New York, NY, 2010.

[10] S. V. Dhople and A. D. Dominguez-Garcia, "A framework to determine the probability density function for the output power of wind farms," in *Proc. North American Power Symp.*, Urbana, IL, Sept. 2012.

[11] J. A. Carta, P. Ramírez, and S. Velázquez, "A review of wind speed probability distributions used in wind energy analysis: Case studies in the Canary islands," *Renew. Sust. Energ. Rev.*, vol. 13, pp. 933–955, 2009.

[12] M. Chiang, S. H. Low, A. R. Calderbank, and J. C. Doyle, "Layering as optimization decomposition: A mathematical theory of network architectures," *Proc. of the IEEE*, vol. 95, no. 1, pp. 255–312, Jan. 2007.

[13] J. Fernandez-Bes, A. G. Marques, and J. Cid-Sueiro, "Battery-aware selective communications in energy-harvesting sensor networks: Optimal solution and stochastic dual approximation," in *Proc. 10th Int. Symp. Wireless Communication Systems*, Aug. 2013.

[14] D. Bertsekas, A. Nedic, and A. E. Ozdaglar, *Convex Analysis and Optimization*, Athena Scientific, 2003.

[15] A. G. Marques, G. B. Giannakis, and J. Ramos, "Optimizing orthogonal multiple access based on quantized channel state information," *IEEE Trans. Sig. Proc.*, vol. 59, no. 10, pp. 5023–5038, Oct. 2011.

[16] N. Gatsis, A. Ribeiro, and G. B. Giannakis, "A class of convergent algorithms for resource allocation in wireless fading networks," *IEEE Trans. Wireless Commun.*, vol. 9, no. 5, pp. 1808–1823, May 2010.

[17] N. Gatsis and G. B. Giannakis, "Residential load control: Distributed scheduling and convergence with lost AMI messages," *IEEE Trans. Smart Grid*, vol. 3, no. 2, pp. 770–786, June 2012.

[18] L. Georgiadis, M. J. Neely, and L. Tassiulas, "Resource allocation and cross-layer control in wireless networks," *Foundations and Trends in Networking*, vol. 1, no. 1, pp. 1–144, 2006.

[19] A. G. Marques, X. Wang, and G. B. Giannakis, "Dynamic resource management for cognitive radios using limited-rate feedback," *IEEE Trans. Sig. Proc.*, vol. 57, no. 9, pp. 3651–3666, Sep. 2009.

[20] A. Ribeiro, "Ergodic stochastic optimization algorithms for wireless communication and networking," *IEEE Trans. Sig. Proc.*, vol. 58, no. 12, pp. 6369–6386, Dec. 2010.

[21] A. G. Marques, L. M. Lopez-Ramos, G. B. Giannakis, J. Ramos, and A. Caamano, "Optimal cross-layer resource allocation in cellular networks using channel and queue state information," *IEEE Trans. Veh. Technol.*, vol. 61, no. 6, pp. 2789 – 2807, Jul. 2012.

[22] H. Lee Willis, *Power Distribution Planning Reference Book*, CRC Press, New York, NY, 2nd edition, 2004.

[23] Y. Zhang, N. Gatsis, and G. B. Giannakis, "Risk-constrained energy management with multiple wind farms," in *Proc. IEEE Innovative Smart Grid Technologies Conf.*, Washington, D.C., Feb. 2013.