Decentralized Large-Scale Power Balancing

Rasmus Halvgaard, John B. Jørgensen,
Niels K. Poulsen, and Henrik Madsen
Department of Applied Mathematics and Computer Science
Technical University of Denmark, DK-2800 Kgs. Lyngby, Denmark
Email: {rhal,jbjo,nkpo,hmad}@dtu.dk

Lieven Vandenberghe
Department of Electrical Engineering
UCLA, 66147L Engineering IV
Los Angeles, CA 90095 USA
Email: lieven.vandenberghe@ucla.edu

***Abstract*—A power balancing strategy based on Douglas-Rachford splitting is proposed as a control method for large-scale integration of flexible consumers in a Smart Grid. The total power consumption is controlled through a negotiation procedure between all units and a coordinating system level. The balancing problem is formulated as a centralized large-scale optimization problem but is then decomposed into smaller subproblems that are solved locally by each unit connected to an aggregator. For large-scale systems the method is faster than solving the full problem and can be distributed to include an arbitrary number of units.**

I. Introduction

A large number of units with flexible consumption, e.g. electrical heat pumps for heating in buildings and Electric Vehicles (EVs) with batteries that can be charged and discharged, are expected to be part of the future Danish power system. These units could potentially provide large-scale flexible consumption by aggregating or pooling the units together. If the aggregated units are controlled and coordinated well, they could help to partially balance the fluctuations in the power system caused by renewable energy sources such as wind or solar power. In real-time electricity markets the aggregated units could help to minimize the imbalances caused by forecast errors and in general provide ancillary services. Controlling a large number of units in real-time requires fast evaluation of the control algorithm that coordinates the power consumption. Thus methods for solving this large-scale optimization problem in real-time must be developed.

In this paper we consider the problem of real-time large-scale power balancing. We apply Douglas-Rachford splitting [1] to decompose the general problem into smaller subproblems to be distributed and solved locally by each unit. Each unit has its own model, constraints, and variables, and can even make decisions based on its own local control objective. All units must communicate their predicted consumption plan to an aggregator that coordinates the need for system level flexibility and minimizes imbalances. The aggregator continuously communicates control signals and synchronizes the global negotiation. This negotiation procedure is required to converge in every time step and thus requires fast evaluation of the unit subproblems that can be cast as convex optimal control problems. Convergence is achieved by coordinating all units' consumption through a negotiation procedure with the dual variables that are updated in the coordinating system level referred to as an aggregator.

With a tracking control objective, the aggregator is able to control and deliver the requested combined power consumption of all units in real-time and can provide running forecasts of the consumption by applying a receding horizon control principle. This principle, referred to as Model Predictive Control (MPC), solves the optimization problem online, but only implements the first part of the solution. At the next sample time, the procedure is repeated by using the new measurements and by moving the prediction window one step. The sampling time in real-time markets could be around five minutes or lower. The precise sampling time is dictated by the settlement requirements of the regulating power market [2]. This short sampling time motivates computational efficient optimization algorithms for the MPC that balances the power. For large-scale systems such as power balancing problems, decomposition methods are one computationally attractive option.

The power reference to be tracked by the aggreagtor could be a fluctuating production from a local wind power. However in this paper we use the Nordic power market framework, where the goal of the aggregator would be to minimize the deviation from a consumption plan already negotiated in advance at the day-ahead market. Any deviations from this day-ahead plan cause power imbalances that must be settled on the regulating power market. These imbalances can be minimized by the aggregator by tracking the day-ahead plan as accurate as possible, while trying to eliminate forecast errors and unforeseen disturbances along the way. If the regulating power prices could be forecasted, even just a few hours ahead, the aggregator could adopt an economic control objective to reflect the actual imbalance costs. Many studies do not consider the imbalances they introduce when only optimizing over the day-ahead market prices [3]. We therefore provide an example of taking regulating power prices into account. Grid capacity constraints on the active power can also be applied as an aggregator constraint.

Compared to dual decomposition that uses the mandatory subgradient projection with rather slow convergence [4], the Douglas-Rachford splitting used in this paper is often faster.

Furthermore the primal is not always easily recoverable from the dual in dual decomposition. Another advantage of the presented method is that its prox-operator easily evaluates complicated and even non-linear expressions. Compared to e.g. Dantzig-Wolfe decomposition [5], which only handles linear programming problems, this method easily handles non-linear convex functions, and converges under very mild conditions. In our case, the subproblems even reduce to simple optimal control problems that can be solved very efficiently using the Riccati recursions [6], [7]. This would not be the case for a similar ADMM formulation [8], where the subproblems are more complicated. By splitting the problem and solving the local subproblems in parallel by each unit, we can control a large number of units in real-time. Even when solving the Douglas-Rachford subproblems sequentially, a significant speedup is observed compared to just solving the centralized problem with standard solvers.

In this paper we illustrate the advantage of the splitting method and show how an aggregator can use the method for power balancing based on flexible consumption units, e.g. heat pumps in buildings. The method can be completely decentralized down to communication with neighbors only [9].

This paper is organized as follows. In Section II we formulate the centralized large-scale optimization problem for the aggregator with linear systems. Section III describes the Douglas-Rachford splitting and arrives at the suggested algorithm that controls the aggregated units. Different aggregator objectives are derived in Section IV. In Section V the control method is demonstrated through simulation and its convergence are discussed. Finally, Section VI provides conclusions.

II. Problem formulation

It is assumed that a certain power consumption profile $q(t)$ is given, e.g. it is settled in advance on a day-ahead market. This profile must be followed by the aggregator, such that the combined power consumption from all units sum to this at every time instant t. The centralized large-scale problem that includes all variables and local unit models is

$$\begin{array}{ll} \text{minimize} & g(p(t)) \\ \text{subject to} & p(t) = \displaystyle\sum_{k=1}^{n} u_k(t) \\ & u_k \in F_k \,. \end{array} \tag{1}$$

The total consumption $p(t)$ with $t = 1, ..., N$ over a period of length N, is a sum of the predicted consumption profiles $u_k(t), k = 1, ..., n$ for each of the n units. A power capacity limitation can also be included by adding the inequality constraint $p(t) \le p^{\max}(t)$. We define F_k as a closed convex set containing the following kth linear state-space system and its constraints

$$F_k \left\{ \begin{array}{l} x_k(t+1) = A_k x_k(t) + B_k u_k(t) \\ y_k(t) = C_k x_k(t) \\ y_k^{\min} \le y_k(t) \le y_k^{\max} \\ u_k^{\min} \le u_k(t) \le u_k^{\max} \,. \end{array} \right. \tag{2}$$

$x_k(t)$ is the state vector with discrete-time state-space system defined by the matrices (A_k, B_k, C_k). $y_k(t)$ is the output, e.g. a temperature, of a linear system with input consumption $u_k(t)$. The input and output limits are superscripted with max and min. The variables are $p(t)$, $x_k(t)$, and $u_k(t)$, while the aggregator objective function is $g(p(t))$. To simplify notation further, the time argument will be omitted from here on. The constraints in F_k can be moved to the objective by putting them in an indicator function

$$f_k(u_k) = \left\{ \begin{array}{cl} 0 & \text{if } u_k \in F_k \\ +\infty & \text{otherwise} \end{array} \right.$$

The optimization problem to be solved at every time instant is

$$\min_{u_k{}_{k=0}^{n}} \quad \sum_{k=1}^{n} f_k(u_k) + g\left(\sum_{k=1}^{n} u_k\right) \tag{3}$$

The functions $f_k : \mathbf{R}^n \to \mathbf{R}$ and $g : \mathbf{R}^n \to \mathbf{R}$ are closed and convex, with nonempty domains. The variable is $u_k \in \mathbf{R}^n$, $k = 1, \ldots, n$ and u_k represents the power consumption of unit k over a given time period N. f_k is its cost function. The sum $p = \sum_k u_k$ is the total power consumption of all units and g is the cost function for the aggregator. The functions f_k and g may include indicator functions that represent constraints on the variables u_k or their sum.

III. Douglas-Rachford splitting

A distributed method such as Douglas-Rachford splitting can solve the problem (3) iteratively. But the dual problem must first be formulated. Define

$$u = [u_1^T, u_2^T, \ldots, u_n^T]^T \in \mathbf{R}^{Nn}$$

and stack the indicator functions

$$f(u) = \sum_{k=1}^{n} f_k(u_k)$$

then (3) can be expressed as

$$\min_u \quad f(u) + g(Cu)$$

where the individual power consumption from each unit has been summed through the matrix C

$$C = \left[\begin{array}{cccc} I & I & \cdots & I \end{array} \right]$$

With the variable p the problem can be reformulated as

$$\begin{array}{ll} \text{minimize} & f(u) + g(p) \\ \text{subject to} & Cu = p \,. \end{array}$$

The Lagrange dual of this problem is

$$\text{maximize} \quad -f^*(-C^T z) - g^*(z)$$

where z is the Lagrange multiplier for the constraint $Cu = p$, and f^* and g^* are the conjugate functions of f and g. The optimality conditions are

$$0 \in \begin{bmatrix} 0 & C^T \\ -C & 0 \end{bmatrix} \begin{bmatrix} u \\ z \end{bmatrix} + \begin{bmatrix} \partial f(u) \\ \partial g^*(z) \end{bmatrix}. \tag{4}$$

Here $\partial f(u)$ and $\partial g^*(z)$ are the subdifferentials of f and g^*, respectively.

The Douglas-Rachford splitting algorithm [10] finds a solution of the optimality conditions via the simple iteration

$$u^+ = \mathrm{prox}_{tf}(v) \tag{5a}$$

$$z^+ = \mathrm{prox}_{tg^*}(s) \tag{5b}$$

$$\begin{bmatrix} w^+ \\ m^+ \end{bmatrix} = \begin{bmatrix} I & tC^T \\ -tC & I \end{bmatrix}^{-1} \begin{bmatrix} 2u^+ - v \\ 2z^+ - s \end{bmatrix} \tag{5c}$$

$$v^+ = v + w^+ - u^+ \tag{5d}$$

$$s^+ = s + m^+ - z^+. \tag{5e}$$

The superscript $^+$ indicates the next iterate, e.g. $v^+ = v_{j+1}$ if j is the iteration number. The aggregator broadcasts a vector v to the units and the units respond with a consumption profile u_k. u_k is determined by evaluating the prox-operator with their local model, constraints, and variables. The remaining steps (5b)-(5e) are computed by the aggregator alone. The (w, m) update gathers the unit consumption profiles and involves multiplications with C and C^T, which can be simplified to

$$\begin{bmatrix} I & tC^T \\ -tC & I \end{bmatrix}^{-1} = \begin{bmatrix} I & 0 \\ 0 & 0 \end{bmatrix} + \frac{1}{b} \begin{bmatrix} -tC^T \\ I \end{bmatrix} \begin{bmatrix} tC & I \end{bmatrix}$$

where $b = 1 + nt^2$. The other steps in the algorithms are local to the aggregator and simply update the coordination variables (v,s). The prox-operator in the first step (5a) that solves the kth unit subproblem and is defined as

$$\mathrm{prox}_{tf}(v) = \arg\min_{\tilde{v}} \left(f(\tilde{v}) + \frac{1}{2t} \|\tilde{v} - v\|_2^2 \right).$$

In our case we must solve this subproblem for all units and stack their solutions in u^+ . Due to the decoupled units, separability of $f(u) = \sum_k f_k(u_k)$ implies that

$$\mathrm{prox}_{tf}(v) = \left(\mathrm{prox}_{tf_1}(v_1), \ldots, \mathrm{prox}_{tf_n}(v_n)\right)$$

if $v = [v_1^T, v_2^T, \ldots, v_n^T]^T$ and each of these prox-operators involve only one unit. We thus have the unit subproblem

$$\mathrm{prox}_{tf_k}(v_k) = \arg\min_{u_k} \left(f_k(u_k) + \frac{1}{2t} \|u_k - v_k\|_2^2 \right)$$

with the standard QP formulation

$$\begin{array}{ll} \text{minimize} & \frac{1}{2}\frac{1}{t} u_k^T u_k - \frac{1}{t} v_k^T u_k \\ \text{subject to} & u_k \in F_k\,. \end{array} \tag{6}$$

In our case, all of these subproblems reduce to a finite horizon constrained LQR problems that can be solved efficiently by methods based on the Riccati recursion [11], [12]. Note also that v_k can be interpreted as an individual linear coefficient for each unit.

IV. Aggregator objective

The dual update of z in (5b) must be evaluated through the prox-operator and depends on the choice of aggregator objective $g(p)$. In this paper we consider the objective of tracking the power consumption profile q. We assume that q is given, e.g. has been negotiated on a day-ahead market a day in advance of the actual consumption. We present two aggregator objectives that do this. First the traditional tracking MPC with quadratic penalty on the residual. Secondly, we also set up an economic MPC that includes the actual costs for imbalances. Both objectives are easily evaluated by the prox-operator that yields simple analytic expressions in both cases.

A. Tracking objective

If q represents a desired total power consumption profile to be followed by the aggregated units, we can choose the differentiable quadratic aggregator objective

$$g(p) = \frac{1}{2}||p - q||_2^2 \tag{7}$$

that is easily evaluated by the prox-operator required in the Douglas-Rachford z^+ update (5b). The conjugate of $g(p)$ is

$$g^*(y) = \sup_p (y^T p - g(p)) = \frac{1}{2}||y||_2^2 + q^T y \tag{8}$$

and the prox-operator simplifies to the analytic expression

$$\mathrm{prox}_{tg^*}(s) = \frac{1}{t+1} s - \frac{t}{t+1} q \tag{9}$$

to be substituted for (5b).

B. Economic objective

The prox-operator of the conjugate g^* is related to the prox-operator of g via

$$\mathrm{prox}_{tg^*}(s) = s - t\mathrm{prox}_{g/t}(s/t)\,. \tag{10}$$

We can exploit this relation when the conjugate function can not be found analytically, e.g. when the aggregator objective is non-differentiable. This is the case when incorporating the regulating power prices in the aggregator objective. When the aggregated power consumption p deviates from the reference plan q, this power imbalance, denoted $d = p - q$, has a penalty equal to the regulating power prices. The price depends on the sign of the imbalance and the overall system imbalance. These regulating power prices are difficult to forecast, but for the Nordic power market the day-ahead price is a known lower or upper bound depending on imbalance direction, i.e. the sign of d. The total cost J of the imbalances within the prediction horizon N sum to

$$J = \sum_t^N \left(\pi_t p_t + \max\left(-\pi_t^+ d_t, \pi_t^- d_t \right) \right) \tag{11}$$

where the up and down-regulating power market prices are forecasted as (π^+, π^-), respectively. π is the fixed day-ahead price traded in advance. In the unlikely case of no imbalances, where d is zero for all time steps, then the price paid by the aggregator is simply π. Note that we introduced the time argument as a subscript, not to be confused with the prox-operator scaling constant t in (10). When integrating this function into the aggregator objective, we get the following piecewise linear cost function to be minimized

$$g(d) = \sum_t \max\left(-\pi_t^+ d_t, \pi_t^- d_t\right) \tag{12}$$

where $d_t = p_t - q_t$ is the tracking error. Since this objective is both separable in units and in time, the prox-operator from (10) is

$$\text{prox}_{g/t}(s/t) = \arg\min_{\tilde{s}_t} \sum_t \left(g_t(\tilde{s}_t) + \frac{t}{2} ||\tilde{s}_t - s_t/t||_2^2 \right)$$

The prox-operator evaluation can be divided into several cases for the unconstrained minimum to be found, since it is a sum of a piecewise linear function $g(d)$ and a quadratic function. Analytically, this leaves the minimum to be found in the following three cases

$$\left(\text{prox}_{g/t}(s_t/t)\right)_t = \begin{cases} (s_t - \pi_t^-)/t & \text{if } s_t \le tq_t - \pi_t^- \\ (s_t + \pi_t^+)/t & \text{if } s_t \ge tq_t + \pi_t^+ \\ q_t & \text{otherwise} \end{cases}$$

Insert this expression into (10) to obtain the final analytical expression

$$\left(\text{prox}_{tg^*}(s_t)\right)_t = \begin{cases} \pi_t^- & \text{if } s_t \le tq_t - \pi_t^- \\ -\pi_t^+ & \text{if } s_t \ge tq_t + \pi_t^+ \\ s_t - tq_t & \text{otherwise} \end{cases} \tag{13}$$

V. Numerical Example

To illustrate the method we provide a simulation with only $n = 2$ different first order thermal storage systems. The models have unity gain, time constants $\tau_1 = 5$ and $\tau_2 = 10$ such that the kth unit has the transfer function

$$G_k(s) = \frac{y_k}{u_k} = \frac{1}{\tau_k s + 1}.$$

u_k is the consumption and y_k is the output temperature. The unit models could also be changed to represent EVs or any other flexible Smart Grid unit. However, on/off control signals, e.g. charge or no charge of an EV, will lead to a non-convex problem, that will not converge. The results for $t = 0.5$ after 50 iterations is shown in Fig. 1. Both the temperatures and the consumption are kept within their operating intervals. Their combined consumption p is seen to match the reference q very well. The dual variable z can be interpreted as a price and expresses the aggregator's need for flexibility at every time instant. It could be used to bill units.

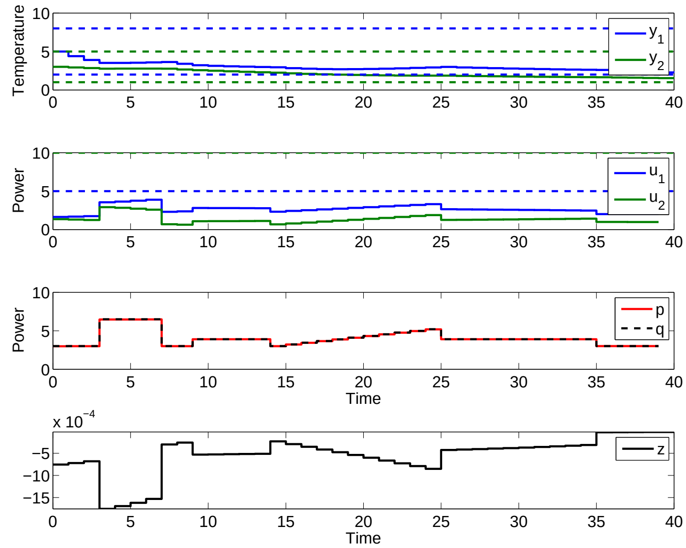

Fig. 1. *Simulation of power balancing with two first order systems. The two input/output consumption/temperature pairs (blue/red) with constraints (dotted) are shown as the first two plots. Right below is the resulting power p and tracking profile q. The lower plot shows the converged dual variable z after 50 iterations.*

A. Large-scale example

To demonstrate that the algorithm works for a larger number of units, we chose $n = 100$ units with uniform randomly generated parameters $\tau_k \in U(5, 10)$, $(u_k^{\min}, u_k^{\max}) \in (0, U(5, 10))$, $(y_k^{\min}, y_k^{\max}) \in (U(1, 5), y_k^{\min} + U(0, 5))$. $U(a, b)$ indicates uniform distribution in the interval from a to b. In Fig. 2 a consumption profile q to be tracked is given and the open-loop profile is plotted. The two upper plots show the normalized temperature and power consumption for the different units. The normalization was done for plotting purposes and the variables were scaled to the interval $[0; 1]$ for each unit and its feasible temperature interval and power consumption. For the lower tracking plot the highest and lowest possible power consumption is plotted as the dotted black lines. The upper bound is the largest possible consumption physically limited by the hard input constraints. (2). The lowest possible consumption is also limited in the same way to zero. The steady state power consumption for the minimum and maximum temperature was also plotted as the dotted red line. These lines define the power operating area, where any q within this region can be followed, assuming no disturbances at the units. Tracking a power reference outside this interval is not always possible and depends on the state of the system, i.e. how far the temperature is from its constraints and also the time constant. In the case of disturbances, i.e. an outdoor temperature affecting the temperature outputs, then the operating interval is time varying and will not be constant over time. The power tracking residual is very small since no uncertainties are present and only deviates when q is outside the feasible operating area.

B. Convergence

From theory, it is known that the step size t in the algorithm must remain constant. However, various heuristics provide adaptive strategies, see for instance the references in [8]. In the numerical example provided in this paper, t was found experimentally, based on the observed convergence behavior. We use the primal and dual optimality conditions from (4) to

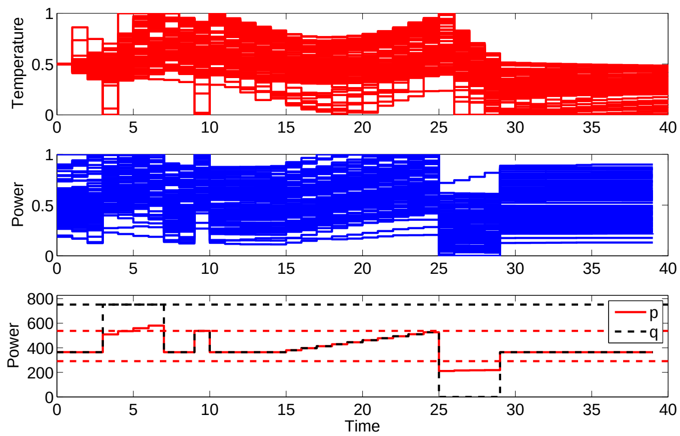

Fig. 2. *Simulation of power balancing with $n = 100$ randomly generated first order systems, converged in 20 iterations with $t = 0.1$. The output temperatures (red) with worst-case constraints (dotted) are shown above the input consumption profiles (blue). The lower plot shows the total power consumption (red), and the target consumption (black). The dotted red lines are the estimated maximum and minimum consumption.*

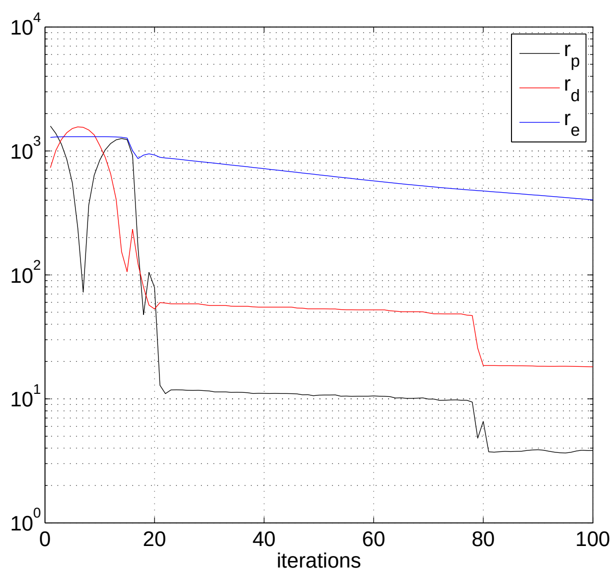

Fig. 3. *Convergence of Douglas-Rachford splitting for scenario with $n = 100$ units, $t = 0.5$. The two-norm was taking for all residuals.*

provide a measure for convergence, i.e.

$$r_p = \frac{s - z^+}{t} - Cu^+ \qquad r_d = \frac{v - u^+}{t} + C^T z^+. \tag{14}$$

In Fig. 3 the convergence rates for the simulation are shown. The convergence measures (14) similar to the tracking error, r_e, and the change in dual variable z, r_z,

$$r_e = q - Cu^+ \qquad\qquad r_z = z - z^+ \tag{15}$$

has also been plotted. We use (14) as stopping criteria for the algorithm. For the simulation with $n = 100$ units, the algorithm is seen to converge within 20 iterations.

VI. Conclusion

In this paper we solve the power balancing problem using a constrained model predictive controller with a least squares tracking error criterion. This is an example of a large-scale optimization problem that must be solved reliably and in real-time. We demonstrated how Douglas-Rachford splitting can be applied in solving this problem. By decomposing the original optimization problem thousands of units can be controlled in real-time by computing the problem in a distributed (parallel) manner. We considered a large-scale power balancing problem with flexible thermal storage units. A given power consumption profile can be followed by controlling the total power consumption of all flexible units through a negotiation procedure with the dual variables introduced in the method. An economic aggregator objective that takes the regulating power prices into account was derived. The solution obtained converges towards the original problem solution and requires fast two-way communication between units and the coordinating level. The resulting power balancing performance is very accurate while the local constraints and objectives for each unit are satisfied and aggregator operation costs are reduced.

References

[1] J. Eckstein and D. Bertsekas, "On the Douglas-Rachford splitting method and the proximal point algorithm for maximal monotone operators," *Mathematical Programming*, vol. 55, pp. 293–318, 1992.

[2] Y. Ding, S. Pineda, P. Nyeng, J. Ostergaard, E. M. Larsen, and Q. Wu, "Real-Time Market Concept Architecture for EcoGrid EU - A Prototype for European Smart Grids," *IEEE Transactions on Smart Grid*, vol. 99, pp. 1–11, 2013.

[3] R. Halvgaard, N. K. Poulsen, H. Madsen, and J. B. Jørgensen, "Economic Model Predictive Control for Building Climate Control in a Smart Grid," in *Proceedings of IEEE PES Innovative Smart Grid Technologies (ISGT)*, 2012.

[4] B. Biegel, P. Andersen, J. Stoustrup, and J. Bendtsen, "Congestion management in a smart grid via shadow prices," in *IFAC Power Plant and Power Systems Control, Toulouse, France*, 2012.

[5] K. Edlund, J. D. Bendtsen, and J. B. Jørgensen, "Hierarchical model-based predictive control of a power plant portfolio," *Control Engineering Practice*, vol. 19, no. 10, pp. 1126–1136, 2011.

[6] J. B. Jørgensen, J. Rawlings, and S. B. Jørgensen, "Numerical Methods for Large Scale Moving Horizon Estimation and Control," in *7th International Symposium on Dynamics and Control of Process Systems (DYCOPS), Boston, USA*, 2004.

[7] J. B. Jørgensen, G. Frison, N. F. Gade-Nielsen, and B. Damman, "Numerical Methods for Solution of the Extended Linear Quadratic Control Problem," 2012.

[8] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers," *Foundations and Trends in Machine Learning*, vol. 3, no. 1, pp. 1–122, 2011.

[9] M. Kraning, E. Chu, J. Lavaei, and S. Boyd, "Message Passing for Dynamic Network Energy Management," *[To appear in] Foundations and Trends in Optimization*, 2013.

[10] L. Vandenberghe, "Lecture notes on operator splitting methods," 2011. [Online]. Available: www.ee.ucla.edu/ vandenbe/ee236c.html

[11] J. B. Jørgensen, "Moving Horizon Estimation and Control," Ph.D. dissertation, Technical University of Denmark, 2005.

[12] C. V. Rao, S. J. Wright, and J. B. Rawlings, "Application of interior-point methods to model predictive control," *Journal of optimization theory and applications*, vol. 99, no. 3, pp. 723–757, 1998.