

Operation and Configuration of a Storage Portfolio via Convex Optimization

Matt Kraning, Yang Wang, Ekine Akuiyibo, Stephen Boyd*

* All authors are members of the Information System Laboratory,
Department of Electrical Engineering, Stanford University; email:
{mkraning, yw224, ekine, boyd}@stanford.edu

Abstract: We consider a portfolio of storage devices which is used to modify a commodity flow so as to minimize an average cost function. The individual storage devices have different parameters that characterize attributes such as capacity, maximum charging rates, and losses in charging and storage. We address two problems related to such a system. The first is the problem of *operating* a portfolio of storage devices in real-time, *i.e.*, making real-time decisions as to how to charge or discharge each of the storage devices in response to the fluctuating commodity flow and cost function. The second is the problem of *configuring* the portfolio of storage devices, *i.e.*, choosing a single portfolio from a set of candidate portfolios. Here we are given the cost of each candidate portfolio as a function of its parameters, and seek to minimize a combination of initial configuration cost and average operating cost. In this paper, we show how both problems can be approximately solved using convex optimization.

Keywords: Convex optimization; Predictive control; Energy management systems

1. INTRODUCTION

We propose the use of multiple storage devices, operated in a coordinated fashion to serve as a single storage portfolio, which is used to modify a commodity flow so as to minimize an average cost function. We consider two problems: The (real-time) operation of such a storage portfolio, and the configuration of such a portfolio (which is a planning or design decision) in order to trade off the running cost and the capital construction cost of the storage portfolio. Although our method is general and can include any commodity, the primary application is energy storage.

In power systems, energy storage devices are used to modify a given input energy flow to help synthesize a desired output energy flow. The uses for energy storage devices are extremely broad and include frequency regulation (Oudalov et al. [2007]), system stability (Mercier et al. [2009]), peak shaving (Even et al. [1993]), and spinning reserve (Kottick and Blau [1993]). Due to the varied uses, scales, and requirements for energy storage applications, the number of different types of energy storage devices is equally broad, and includes pumped hydro, compressed air energy storage (CAES), battery energy storage systems (BESS), supercapacitors, and deferrable loads, as just some examples.

A recent application for energy storage devices is buffering the output power flow of intermittent, renewable energy sources including wind (Zeng et al. [2006], Arulampalam et al. [2006]) and solar (Teleke et al. [2010]) farms. These systems are typically built with an accompanying battery bank that is able to maintain a desired output power flow

in the presence of either cloud cover or decreased wind speed by being charged during periods of increased power generation. Because both the capital costs of the storage systems and the potential energy savings can range from the tens of thousands (Barton and Infield [2004]) into the tens of millions (Alt et al. [1997]) of dollars per year for large-scale systems, operation and configuration of energy storage portfolios is of paramount importance.

There is an extensive literature on the operation and sizing of batteries (Lee and Chen [1993], Yoshimoto et al. [2006], Banos et al. [2006]), with optimal operation and sizing being considered under many different operating conditions (Oudalov et al. [2006]). Different battery sizes and types are considered under multiple usage scenarios, and the best performing type and size of battery is subsequently picked as the (single device) storage portfolio. However, single energy storage systems are typically specialized for specific scenarios. We will see that a properly constructed portfolio of different energy storage devices can better meet a wide range of requirements.

We describe a method called receding horizon control (RHC), which we will use for operating the storage portfolio. In RHC, we solve an optimization problem at each time step to determine a plan of action over a finite time horizon. The first step of the plan is executed, and the process is repeated at the next time step, incorporating new measurements and external data that have become available (Maciejowski [2002], Bemporad [2006]). RHC has been successfully applied in a wide range of settings, including chemical process control (Qin and Badgwell [2003]), supply chain management (Cho et al.

[2003]), stochastic control in economics and finance (Herzog [2005], Talluri and Ryzin [2004]), and many others.

RHC is an ideal controller for energy storage portfolios for several reasons. First, it is highly versatile, and handles a variety of objectives and constraints naturally and directly. In classical control systems, objectives and constraints are instead handled indirectly, by adjusting controller coefficients via a cumbersome trial and error process. Another advantage is that RHC does not require a formal stochastic model of the uncertainty, which is not easy to obtain in practical settings. The controller only requires predictions of future quantities, which can be based on historical data, stochastic models, weather forecasts, or analyst predictions. In many problems, even when the predictions are poor, the controller often performs exceptionally well (Wang and Boyd [2009]).

Recent advances in convex optimization allow for RHC problems to now be solved at millisecond and microsecond time-scales (Wang and Boyd [2008], Mattingley et al. [2010], Bemporad et al. [2002], Bemporad and Filippi [2004]). This allows the control policy to be implemented online at kilohertz rates, which covers a wide range of applications, and essentially all energy storage problems. Our ability to solve the required optimization problems at very high speed is useful even when the actual application does not require high speed execution; for example, offline Monte Carlo simulations to evaluate the system performance can be carried out quickly. In this way, RHC can be extensively tested offline via simulation, before online implementation.

Another important problem we consider is how to choose the configuration of the storage devices in our portfolio. Larger sizes lead to better performance, but larger batteries/reservoirs/pumps are also more expensive. The configuration problem is to choose the parameters of the storage system in order to minimize a trade-off between the average per-period operation cost plus the initial cost of constructing the storage portfolio, amortized over the portfolio lifetime. In this paper we propose leveraging the ability to rapidly solve optimization problems to solve the the configuration problem via Monte Carlo simulations.

The rest of the paper is organized as follows. In section 2 we give the system model, objective, constraints, and we define the operation and configuration problems. In section 3 we describe methods for solving these problems, and in section 4 we present a numerical example.

2. MODEL DESCRIPTION

2.1 System dynamics and constraints

We consider a portfolio of n storage devices (shown in Figure 1) with dynamics

$$q_{t+1} = \eta^l \circ q_t + \eta^c \circ u_t^+ - (1/\eta^d) \circ u_t^- + w_t,$$

for $t = 1, 2, \dots$, where $q_t \in \mathbf{R}^n$ is the vector of charge levels, $u_t^+ \in \mathbf{R}^n$ is the vector of charging rates or inflows into the storage devices, and $u_t^- \in \mathbf{R}^n$ is the vector of discharging rates or outflows from the storage

devices, at time t . The parameters $\eta^l, \eta^c, \eta^d \in (0, 1]^n$ are vectors of storage leakage, and charging and discharging efficiencies, respectively; $1/\eta^d$ is interpreted elementwise; and \circ denotes the Hadamard (elementwise) product. Lastly, $w_t \in \mathbf{R}^n$ is an exogenous input at time t . In the simplest case with lossless storage devices we have $\eta^l = \eta^c = \eta^d = \mathbf{1}$, where $\mathbf{1}$ is the vector with all entries equal to one.




Fig. 1. Storage model with n storage devices. Dotted box encloses the storage system.

We require that q_t , u_t^+ , and u_t^- lie in the ranges

$$0 \leq q_t \leq Q, \quad 0 \leq u_t^+ \leq C, \quad 0 \leq u_t^- \leq D, \quad t = 1, 2, \dots,$$

where $Q \in \mathbf{R}^n$ is the vector of storage capacities, and $C \in \mathbf{R}^n$ and $D \in \mathbf{R}^n$ are vectors of maximum charging and discharging rates. (The inequalities are interpreted elementwise.) The total inflow into our storage system at time t is $\mathbf{1}^T u_t^+$ and the total outflow is $\mathbf{1}^T u_t^-$. The storage portfolio is completely characterized by the parameters $(Q, C, D, \eta^l, \eta^c, \eta^d)$.

The amount of energy pulled from the source at time t is denoted by $s_t \in \mathbf{R}$, and the amount of energy delivered is $d_t \in \mathbf{R}$. Energy balance can be expressed as

$$d_t - s_t + \mathbf{1}^T u_t^+ - \mathbf{1}^T u_t^- = 0.$$

To simplify our notation, we let

$$v_t = (d_t, s_t, u_t^+, u_t^-)$$

be the vector of control variables at time t .

2.2 Objective

The cost incurred in each time period is denoted $\ell_t(v_t, q_t)$, where $\ell_t : \mathbf{R}^{2n+2} \times \mathbf{R}^n \rightarrow \mathbf{R}$ is a convex stage cost function. The cost depends on the amount of charging and discharging, the amount of energy pulled from the source, the amount of energy delivered, and the charge levels. For the problems we consider, the stage cost function has the separable form

$$\ell_t(v_t, q_t) = \phi_t^{\text{sr}}(s_t) + \phi_t^{\text{de}}(d_t) + \phi_t^{\text{ch}}(u_t^+, u_t^-) + \phi_t^{\text{st}}(q_t).$$

We can interpret each term of this decomposition. The first term, $\phi_t^{\text{sr}}(s_t)$, is a cost for pulling energy from the source. In many applications we simply have $\phi_t^{\text{sr}}(s_t) =$

$p_t s_t$, where p_t is the energy price at time t . Another choice for ϕ_t^{sr} is

$$\phi_t^{\text{sr}}(s_t) = \begin{cases} p_t s_t, & 0 \leq s_t \leq S^{\text{max}} \\ \infty, & \text{otherwise,} \end{cases}$$

which additionally constrains s_t to lie in the interval $[0, S^{\text{max}}]$. The second term, $\phi_t^{\text{de}}(d_t)$, is a cost that depends on the amount of energy delivered. For example, we can have $\phi_t^{\text{de}}(d_t) = \alpha(r_t - d_t)_+$, where r_t is the amount of energy requested at time t , α is a penalty for not meeting the requested demand, and $(z)_+ = \max\{z, 0\}$. In other applications we might penalize excessive charging and discharging via ϕ_t^{ch} , or add an energy storage cost, ϕ_t^{st} .

We do not assume that ℓ_t is known in advance. Indeed, for many of the examples we consider, uncertainties such as time-varying price or demand are precisely captured by the uncertain stage cost function. Our overall objective is the average operation cost

$$J^{\text{op}} = \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} \ell_t(v_t, q_t), \quad (1)$$

where we assume the limit exists.

2.3 Control policy

The control policy is a function that chooses the vector of control variables at time t , v_t , based on information available at time t . The information available can include measured quantities (such as a history of charge levels), but it can also include estimates of quantities that are not known (such as future stage costs) based on known information. We denote these estimates by $\hat{w}_{\tau|t}$ and $\hat{\ell}_{\tau|t}$, where the subscript $\tau|t$ denotes an estimate of the quantity at time τ , based on information available at time t . All other quantities, such as the parameters associated with the storage portfolio, $Q, C, D, \eta^l, \eta^c, \eta^d$, are known and fixed during operation.

These estimates can be obtained by many methods. When we have a statistical model of the uncertain quantity, the estimates can be the conditional expectations. For example when $w_t, t = 1, 2, \dots$, are independent and identically distributed, we can take $w_{\tau|t} = \mathbf{E} w_t$. However, estimates need not come from statistical models; indeed, in many applications such models do not even exist. Instead, external data such as weather forecasts, historical patterns, or analyst predictions can also be used to obtain estimates.

2.4 Operation

The operation problem is to find a policy that chooses the control variables v_t , as a function of the information available at time t , so that the average operation cost, J^{op} , is minimized and all constraints are satisfied. We should point out that our description of the optimal operation problem is informal, since we have not specified statistical models for the unknown quantities. We could state the problem formally as a stochastic control problem; however, for practical applications accurate statistical models

are not usually available, so such a formal statement would not be particularly meaningful.

2.5 Configuration

We consider a set of N candidate portfolios, which are constructed from combinations and configurations of individual storage devices, such as batteries, reservoirs, and capacitors. Each portfolio has an associated capital cost J^{cap} , which represents the cost to build or acquire that candidate portfolio. Given a particular control policy, each portfolio also incurs an average operating cost J^{op} . The configuration problem is to find portfolios that are Pareto optimal with respect to the costs $(J^{\text{cap}}, J^{\text{op}})$.

3. METHODS

3.1 Operation

We will use receding horizon control as the control policy for operating storage portfolios. At time t , we consider a fixed time interval extending T steps into the future: $t, t+1, \dots, t+T-1$. We first form predictions of exogenous inputs, $\hat{w}_{t|t}, \dots, \hat{w}_{t+T-1|t}$, and stage costs $\hat{\ell}_{t|t}, \dots, \hat{\ell}_{t+T-1|t}$ over this time interval. Next, we solve

$$\begin{aligned} & \text{minimize} \quad \frac{1}{T} \sum_{\tau=t}^{t+T-1} \hat{\ell}_{\tau|t}(\hat{v}_{\tau}, \hat{q}_{\tau}) \\ & \text{subject to} \quad \hat{q}_{\tau+1} = \eta^l \circ \hat{q}_{\tau} + \eta^c \circ \hat{u}_{\tau}^+ - (1/\eta^d) \circ \hat{u}_{\tau}^- + \hat{w}_{\tau|t}, \\ & \quad \hat{d}_{\tau} - \hat{s}_{\tau} + \mathbf{1}^T u_{\tau}^+ - \mathbf{1}^T u_{\tau}^- = 0, \\ & \quad 0 \leq \hat{q}_{\tau} \leq Q, \quad 0 \leq u_{\tau}^+ \leq C, \\ & \quad 0 \leq u_{\tau}^- \leq D, \quad \tau = t, \dots, t+T-1 \\ & \quad \hat{q}_t = q_t, \quad \hat{q}_{t+T} = q_{\text{final}}, \end{aligned}$$

with variables $\hat{q}_t, \dots, \hat{q}_{t+T}$ and $\hat{v}_t, \dots, \hat{v}_{t+T-1}$. At time t , let $\hat{q}_t^*, \dots, \hat{q}_{t+T}^*, \hat{v}_t^*, \dots, \hat{v}_{t+T-1}^*$ be an optimal solution to the RHC problem. The RHC policy takes $v_t = \hat{v}_t^*$. The process is repeated at the next time step, with new data and predictions.

We make a few comments about this control policy. First, when $\hat{\ell}_{t|t}, \dots, \hat{\ell}_{t+T-1|t}$ are convex, this is a convex optimization problem and can be solved efficiently (Boyd and Vandenberghe [2004], Nocedal and Wright [1999], Grant and Boyd [2008]). This will be the case for the example we present in this paper. Problems with nonconvex objective and constraints can also be handled. In these instances, we solve the nonconvex problem by solving a sequence of convex optimization problems to obtain good local solutions (Diehl et al. [2002], Houska and Ferreau [2008]). We have also added a terminal constraint $\hat{q}_{t+T} = q_{\text{final}}$, to ensure that the storage system is not depleted at the end of the time horizon.

3.2 Configuration

We identify Pareto optimal portfolios by evaluating J^{cap} and J^{op} for each candidate portfolio. We assume J^{cap} is given for each candidate portfolio. The operation cost J^{op} is evaluated by simulating the operation of the portfolio with the RHC policy, and averaging the stage costs

incurred over a large number of time periods. If N is too large to efficiently perform exhaustive simulations, we can grid or find representative samples of regions of the set of candidate portfolios and then use successive refinement to exhaustively evaluate all candidate portfolios in good performing areas.

One way to choose a particular portfolio is to amortize J^{cap} over the operational lifetime of the constituent storage devices in each candidate portfolio. We then pick the candidate portfolio that achieves the smallest sum of average operation cost and amortized capital configuration cost.

4. NUMERICAL EXAMPLE

Our storage portfolio is comprised of $n = 3$ different types of storage devices with parameters listed in Table 1. We reference these devices by their storage capacity as large (L), medium (M), and small (S). We refer to the portfolio containing one large, one medium, and one small device as the basic candidate storage portfolio. A candidate portfolio is constructed by selecting 0, 1, 2, or 3 units of each storage device, giving $N = 4^3$ candidate portfolios.

| device | Q | C | D | η | η_c | η_d | $J^{\text{cap}}/\text{unit}$ |
|--------|-----|------|------|--------|----------|----------|------------------------------|
| L | 5 | 0.75 | 0.75 | 0.98 | 0.8 | 0.8 | 5 |
| M | 2 | 0.5 | 0.5 | 0.99 | 0.9 | 0.9 | 3 |
| S | 1 | 0.5 | 0.5 | 0.995 | 1 | 1 | 2 |

Table 1. Storage device parameters.

We discretize time into 30 minute intervals and use the two term stage cost

$$\ell_t(v_t, q_t) = p_t s_t + \alpha(r_t - d_t)_+,$$

where p_t is the price we pay for pulling energy from the source at time t , and $\alpha > 0$ is a penalty for not delivering the requested energy r_t ; typically we have $\alpha \gg p_t$.

Uncertainty in ℓ_t follows from uncertainty in r_t and p_t . We model r_t and p_t as stochastic processes with diurnal components, where the initial time $t = 0$ corresponds to midnight, and the expected peak demand request and price occur at 3PM and 6PM each day, respectively. These processes are given by

$$\begin{aligned} \log r_t &= a_r + b_r \cos(2\pi t/48 - 5\pi/4) + u_t + x_t, \\ \log p_t &= a_p + b_p \cos(2\pi t/48 - 3\pi/2) + u_t + y_t, \\ u_t &= hu_{t-1} + z_t, \end{aligned}$$

for $t = 1, 2, \dots$, where x_t, y_t, z_t are each IID gaussian processes, with $x_t \sim \mathcal{N}(0, \sigma_x^2)$, $y_t \sim \mathcal{N}(0, \sigma_y^2)$, $z_t \sim \mathcal{N}(0, \sigma_z^2)$, and u_t is a zero mean, first order autoregressive process — common to both prices and requests — with autocorrelation

$$\mathbf{E} u_t u_{t+\tau} = \begin{cases} \sigma_z^2/(1-h^2), & \tau = 0 \\ h^\tau \sigma_z^2/(1-h^{2\tau}), & \tau \neq 0. \end{cases}$$

At time t , we assume that the controller has access to the charge q_t , the current and previous 48 (24 hours) prices and energy requests $p_t, p_{t-1}, \dots, p_{t-48}, r_t, r_{t-1}, \dots, r_{t-48}$,

as well as all model parameters. We take the planning horizon of our RHC control policy to be the next 24 hours ($T = 48$). We note that $\log r_\tau$ and $\log p_\tau$ are gaussian processes, and their conditional distributions are given by

$$\begin{aligned} \log r_\tau | r_t, \dots, r_{t-48} &\sim \mathcal{N}(\mu_{r_\tau|t}, \sigma_{r_\tau|t}^2), \\ \log p_\tau | p_t, \dots, p_{t-48} &\sim \mathcal{N}(\mu_{p_\tau|t}, \sigma_{p_\tau|t}^2), \end{aligned}$$

for $\tau = t, t+1, \dots, t+47$, where $\mu_{r_\tau|t}, \sigma_{r_\tau|t}^2, \mu_{p_\tau|t}$, and $\sigma_{p_\tau|t}^2$ are easily computed as affine functions of $\log r_t, \dots, \log r_{t-48}, \log p_t, \dots, \log p_{t-48}$. We predict the future demands and prices over the RHC horizon using their conditional expectations

$$\begin{aligned} \hat{r}_{\tau|t} &= \exp(\mu_{r_\tau|t} + \sigma_{r_\tau|t}^2/2), \\ \hat{p}_{\tau|t} &= \exp(\mu_{p_\tau|t} + \sigma_{p_\tau|t}^2/2), \end{aligned}$$

for $\tau = t, t+1, \dots, t+47$. Thus our stage cost predictions are

$$\hat{\ell}_{\tau|t}(v_t, q_t) = \hat{p}_{\tau|t} s_t + \alpha(\hat{r}_{\tau|t} - d_t)_+,$$

for $\tau = t, t+1, \dots, t+47$. For this example, we used the parameters $a_r = 0.2, a_p = 0.15, b_r = 0.4, b_p = 0.4, h = 0.9, \sigma_x^2 = 0.01, \sigma_y^2 = 0.01$, and $\sigma_z^2 = 0.01$. In addition, $\alpha = 20$, and the maximum energy purchase is restricted to $S^{\text{max}} = 1.5$, for all t .

We should mention that our model for price and demand is chosen to be simple to describe but also have some characteristics of real electrical demand and price. These include a large, predictable baseline diurnal variation, common and coupled variations from this baseline (given by u_t), and small, short-term, unpredictable fluctuations (x_t and y_t). Of course, we could easily use more complex models for predicting future price and demand for the RHC control policy.

4.1 Operation

We used CVXGEN (Mattingley and Boyd [2010a,b]) to solve for the average operating cost of the RHC control policy over a 1 year period (365 days) for each candidate portfolio. CVXGEN transformed the original RHC optimization problem into a standard form quadratic program with 576 variables and 1296 constraints. On a 3.2 Ghz Intel Core i3, the CVXGEN code took an average of 6.56 ms (on a single core) to solve at each time step, so that a year long simulation of a candidate portfolio could be carried out in under 2 minutes. The same problem took an average of 3.23 seconds to solve using the generic optimization solver SDPT3 (Toh et al. [1999], Grant and Boyd [2008]), so that a year long simulation of a single candidate portfolio would take in excess of 15 hours.

Figures 2 and 3 show typical 5-day price and demand, and charge trajectories, respectively, for the basic portfolio. The top plot in Figure 2 shows demand requested (blue) and price (red). The bottom plot compares the demand requested (blue) to the power delivered by RHC (green) and the power delivered with no storage (black). With energy storage, we can match the demand much more closely, only failing to meet the demand for a fraction of one of the five days.

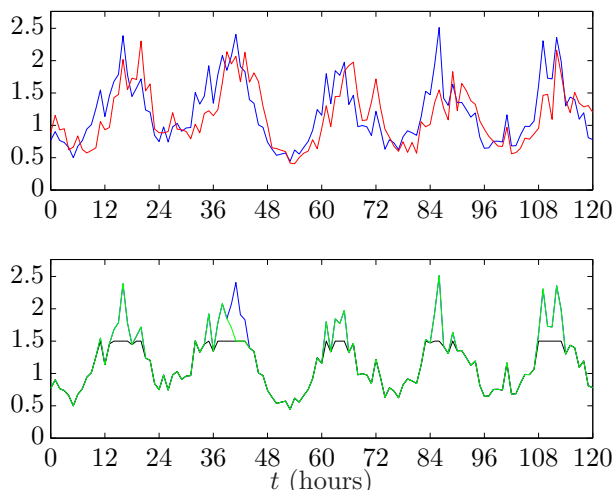


Fig. 2. Top: demand (blue), price (red). Bottom: demand (blue), energy delivered using RHC (green), energy delivered with no storage (black).




Fig. 3. Storage device charge profiles for basic portfolio. Top: $(q_t)_L$, Middle: $(q_t)_M$. Bottom: $(q_t)_S$.

4.2 Configuration

Figure 4 shows the capital cost versus average cost of operating the system for 1 year for all 64 candidate portfolios. The Pareto optimal configurations are indicated in red while the basic portfolio is indicated in black. Of all candidate portfolios, 13 are Pareto optimal with respect to the costs $(J^{\text{cap}}, J^{\text{op}})$.

We make a few statements about particular Pareto optimal portfolios. The four leftmost (smallest J^{cap}) Pareto optimal portfolios correspond to, in increasing order of J^{cap} , no storage, one small, one medium, and one large storage device, and have average operating costs of 4.16, 4.07, 4.04, and 3.60, respectively. These four portfolios all operate in a ‘capacity limited’ regime, where they are each unable to fully meet daily energy requests. Consequently, the portfolios with larger energy storage capacity can meet more of the larger demand, thereby reducing their average operating costs by incurring a smaller demand shortfall penalty. The three rightmost Pareto optimal

configurations correspond to configurations with 3 units of both the small and medium storage devices, and, in increasing order of J^{cap} , 1, 2, and 3 units of the large storage device, with average operating costs 2.74, 2.722, and 2.720 respectively. In this ‘capacity satisfied’ regime, all three storage portfolios can (nearly) match energy requests on every day. Thus, the addition of large storage devices adds substantially to capital costs, while minimally improving operational performance.




Fig. 4. Capital cost versus average operating cost for all configurations.

5. CONCLUSIONS

In this paper we have presented methods, based on convex optimization, for operating and configuring a portfolio of energy storage devices. A major advantage of our methods is that they are highly versatile, and apply to a wide range of energy storage applications, without requiring a specific model of future uncertainty.

ACKNOWLEDGEMENTS

This material is based on work supported by Precourt 1140458-1-WPIAE, JPL contract 1400723, and AFOSR grant FA9550-09-0130.

REFERENCES

- J. Alt, M. Anderson, and R. Jungst. Assessment of utility side cost savings from battery energy storage. *IEEE Transactions on Power Systems*, 12(3):1112–1120, August 1997.
- A. Arulampalam, M. Barnes, N. Jenkins, and J. Ekanayake. Power quality and stability improvement of a wind farm using STATCOM supported with hybrid battery energy storage. *IEE Proceedings — Generation, Transmission and Distribution*, 153(6): 701–710, November 2006.
- C. Banos, M. Aten, P. Cartwright, and T. Green. Benefits and control of STATCOM with energy storage in wind power generation. In *The 8th IEE International Conference on AC and DC Power Transmission, 2006*, pages 230–235, March 2006.

- J. Barton and D. Infield. Energy storage and its use with intermittent renewable energy. *IEEE Transactions on Energy Conversion*, 19(2):441–448, June 2004.
- A. Bemporad. Model predictive control design: New trends and tools. In *Proceedings of 45th IEEE Conference on Decision and Control*, pages 6678–6683, 2006.
- A. Bemporad and C. Filippi. Suboptimal explicit receding horizon control via approximate multiparametric quadratic programming. *Journal of Optimization Theory and Applications*, 117(1):9–38, November 2004.
- A. Bemporad, M. Morari, V. Dua, and E. N. Pistikopoulos. The explicit linear quadratic regulator for constrained systems. *Automatica*, 38(1):3–20, 2002.
- S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
- E. G. Cho, K. A. Thoney, T. J. Hodgson, and R. E. King. Supply chain planning: Rolling horizon scheduling of multi-factory supply chains. In *Proceedings of the 35th conference on Winter simulation: driving innovation*, pages 1409–1416, 2003.
- M. Diehl, R. Findeisen, S. Schwarzkopf, I. Uslu, F. Allgöwer, H. G. Bock, E. D. Gilles, and J. P. Schlöder. An efficient algorithm for nonlinear model predictive control of large-scale systems. Part I: Description of the method. *Automatisierungstechnik*, 50(12):557–567, 2002.
- A. Even, J. Neyens, and A. Demousselle. Peak shaving with batteries. In *12th International Conference on Electricity Distribution. CIRED.*, pages 5.17/1–5.17/6 vol.5, May 1993.
- M. Grant and S. Boyd. CVX: Matlab software for disciplined convex programming (web page and software). <http://www.stanford.edu/boyd/cvx/>, July 2008.
- F. Herzog. *Strategic Portfolio Management for Long-Term Investments: An Optimal Control Approach*. PhD thesis, ETH, Zurich, 2005.
- B. Houska and H. J. Ferreau. ACADO toolkit: Automatic control and dynamic optimization (web page and software). <http://www.acadotoolkit.org/>, August 2008.
- D. Kottick and M. Blau. Operational and economic benefits of battery energy storage plants. *International Journal of Electrical Power & Energy Systems*, 15(6):345–349, 1993.
- T. Lee and N. Chen. Optimal capacity of the battery energy storage system in a power system. *IEEE Transactions on Energy Conversion*, 8(4):667–673, December 1993.
- J. Maciejowski. *Predictive Control with Constraints*. Prentice Hall, 2002.
- J. Mattingley and S. Boyd. CVXGEN: Automatic convex optimization code generation (web page and software). <http://cvxgen.com/>, April 2010a.
- J. Mattingley and S. Boyd. Automatic code generation for real-time convex optimization. In D. P. Palomar and Y. C. Eldar, editors, *Convex optimization in signal processing and communications*, pages 1–41. Cambridge University Press, 2010b.
- J. Mattingley, Y. Wang, and S. Boyd. Code generation for receding horizon control. In *Proceedings of the IEEE Multi-Conference on Systems and Control*, pages 985–992, September 2010.
- P. Mercier, R. Cherkaoui, and A. Oudalov. Optimizing a battery energy storage system for frequency control application in an isolated power system. *IEEE Transactions on Power Systems*, 24(3):1469–1477, August 2009.
- J. Nocedal and S. J. Wright. *Numerical Optimization*. Springer, 1999.
- A. Oudalov, D. Chartouni, C. Ohler, and G. Linhofer. Value analysis of battery energy storage applications in power systems. In *Power Systems Conference and Exposition, 2006. PSCE '06. 2006 IEEE PES*, pages 2206–2211, October 2006.
- A. Oudalov, D. Chartouni, and C. Ohler. Optimizing a battery energy storage system for primary frequency control. *IEEE Transactions on Power Systems*, 22(3):1259–1266, August 2007.
- S. J. Qin and T. A. Badgwell. A survey of industrial model predictive control technology. *Control Engineering Practice*, 11(7):733–764, 2003.
- K. T. Talluri and G. J. Van Ryzin. *The Theory and Practice of Revenue Management*. Springer, 2004.
- S. Teleke, M. Baran, S. Bhattacharya, and A. Huang. Rule-based control of battery energy storage for dispatching intermittent renewable sources. *IEEE Transactions on Sustainable Energy*, 1(3):117–124, October 2010.
- K. C. Toh, M. J. Todd, and R. H. Tütüncü. SDPT3—a Matlab software package for semidefinite programming, version 1.3. *Optimization Methods and Software*, 11(1):545–581, 1999.
- Y. Wang and S. Boyd. Fast model predictive control using online optimization. In *Proceedings IFAC World Congress*, pages 6974–6997, July 2008.
- Y. Wang and S. Boyd. Performance bounds for linear stochastic control. *System and Control Letters*, 53(3):178–182, March 2009.
- K. Yoshimoto, T. Nanahara, G. Koshimizu, and Y. Uchida. New control method for regulating state-of-charge of a battery in hybrid wind power/battery energy storage system. In *Power Systems Conference and Exposition, 2006. PSCE '06. 2006 IEEE PES*, pages 1244–1251, October 2006.
- J. Zeng, B. Zhang, C. Mao, and Y. Wang. Use of battery energy storage system to improve the power quality and stability of wind farms. In *International Conference on Power System Technology, 2006. PowerCon 2006*, pages 1–6, October 2006.