Model predictive control of thermal storage for demand response

Invited Paper, 2015 American Control Conference, Chicago, IL

Kevin J. Kircher and K. Max Zhang

***Abstract*— Buildings with thermal storage use it mainly to shift cooling loads. Ice or chilled water is produced when electricity prices are low and stored to provide cooling when prices are high. While this price-based load shifting has value for power system operators, buildings with thermal storage could provide more direct grid services by reacting to demand charges and demand response calls. In this paper, we consider the problem of cooling a building under these incentives. The context is a New York City office building with passive and active thermal storage, subject to Consolidated Edison's (ConEd's) default rate plan for large commercial buildings. This rate plan includes a three-tiered demand charge and hourly energy prices determined by the system operator's day-ahead dispatch. We also model a ConEd demand response program, and consider the thermal comfort of building occupants. The problem is formulated in the language of stochastic optimal control and solved approximately using model predictive control (MPC). Extending previous work on MPC of thermal storage, which has focused on dynamic energy prices, we include the full set of economic incentives directly in the stage and terminal costs. Simulations of the hottest day of 2013 demonstrate the value of realistic economic modeling. They also highlight an interesting tension between the various incentives, which all compete for shiftable load.**

I. Introduction

In electric power systems, hours of peak demand carry disproportionately high social costs. Because power plants are dispatched in order of ascending marginal cost, energy from the plants that come online to meet system peaks is the most expensive. Additionally, in systems with capacity markets, grid operators must procure enough capacity each year to meet the forecasted peak demand. This often means paying for the ongoing availability of peaking plants that are only dispatched a few hundred hours per year. Peaking plants also raise public health concerns, as they tend to emit pollutants at high rates and at times of already poor air quality.

Recently, system operators have turned to demand-side resources to mitigate the costs of system peaks. Demand response programs, which encourage consumers to curtail load when the grid is stressed, allow system operators to procure less generation capacity, and to dispatch peaking plants less often. A recent FERC study estimates that U.S. system operators have 28.3 GW of registered demand response capacity, equivalent to about 750, 75-MW peaking plants. [1]

System peaks typically occur during the hottest hours of the year, driven by residential and commercial air conditioning. Therefore, cooling loads are valuable demand response resources. To building operators, however, curtailing cooling systems during hot hours risks compromising occupants' thermal comfort. By reducing this risk, thermal storage could play an important role in expanding the pool of demand response resources.

Thermal storage is a mature technology that has been used for decades to shift cooling loads in commercial buildings. A typical thermal storage system consists of an insulated tank and a chiller that fills it with ice or cold water. The capital costs of thermal storage are estimated at 14-20 $/kWh$_{\text{th}}$ [2], significantly lower than those of electrochemical storage[1], which (for lithium-ion batteries) are in the 500-600 $/kWh range [3].

Most building operators with thermal storage systems control them with heuristics such as storage-priority, chiller-priority or constant-proportion control. [4] Although these heuristics have been shown to perform well under time-of-use rate plans, they are becoming outdated as building operators face increasingly complex economic environments. Optimization-based control methods are better equipped to negotiate the trade-offs between competing incentives such as dynamic energy prices, multi-tiered demand charges, demand response revenues, and occupant thermal comfort.

MPC is one such optimization-based method that is widely used in building control problems. [5] MPC was first applied to thermal storage by Henze *et al.* in 1997. [6] Modeling a chiller and tank as a SISO LTI system, the authors showed that under dynamic energy prices, MPC significantly outperforms the previously mentioned heuristics. In the last eighteen years, this work has been extended to make use of a building's passive thermal mass [7], [8]; to incorporate building simulation tools such as EnergyPlus and TRNSYS [8]; to actuate temperature setpoints and mass flow rates [9]; and to handle on/off chiller constraints [10].

While extensive, this body of work has, with a few exceptions, paid little attention to several key objectives that arise in real-world building operations. Perhaps the most important of these is the demand charge that many large commercial buildings face. Demand charges penalize a building's peak power consumption over a monthly billing period. These

[1]The subscript "th" in kWh$_{\text{th}}$ is included to emphasize that thermal energy, rather than electrical energy (kWh), is measured. We adopt this convention throughout the paper. The thermal energy evacuated by a chiller and the electrical energy required to do so are related by the chiller's coefficient of performance. For example, a typical ice-making chiller has a coefficient of performance of 2.5-3, so a thermal storage cost of 14-20 $/kWh$_{\text{th}}$ is loosely equivalent to an electric storage cost of 35-60 $/kWh.

The authors are with the Sibley School of Mechanical and Aerospace Engineering, Cornell University, Ithaca, NY 14850. {kjk82,kz33}@cornell.edu

charges couple decisions across long stretches of time, posing modeling and computational challenges.

In 2008, Henze *et al.* addressed the demand charge, in the absence of thermal storage, by applying the Nelder-Mead heuristic method to nonconvex MPC with embedded TRNSYS simulation. [11] More recently, Ma *et al.* handled the demand charge directly (again for a building without thermal storage) in a linear program. [12] Applying this method to a commercial building with dynamic energy prices, Ma *et al.* demonstrated 3-28% savings in weekly energy costs relative to several heuristics.

In this paper, we use the framework of convex optimization to include a multi-tiered demand charge and the other objectives of a challenging economic environment in an algorithm for MPC of thermal storage. We apply the algorithm to a New York City commercial building operating under ConEd's default rate plan and a ConEd demand response program. [14] For more discussion of this rate plan and demand response program, see [15].

This paper is organized as follows. In §II, we develop models of the building physics and the economic incentives. Control algorithms are presented in §III, and §IV describes the simulation parameters and results. We conclude with some discussion in §V.

II. Models

Our primary focus is on detailed economic modeling. We base the physical model on the one presented in [6], extended to include a second chiller, chiller ramping limits, time-varying coefficients of performance, and non-ideal storage and heat exchanger efficiencies. The resulting model is a perfectly observed MIMO stochastic system with linear, time-varying dynamics.

In the sequel, the variable k indexes the set $\mathcal{T} = \{0, \ldots, N-1\}$, were $N = T/\Delta t$ is the number of discrete time steps of length Δt (hours) in the control horizon T. In this paper, we take $T = 24$ h, although with minor revisions the horizon could be extended to a month or a cooling season.

A. Physics

Let $x_1(k)$ be the charge state of the ice tank (kWh$_{\text{th}}$), bounded above by the capacity $\bar{x}_1$. Let $x_2(k)$ be the building's cooling deficit (kWh$_{\text{th}}$), *i.e.*, the sum of all unmet cooling load from previous stages. Note that $x_2(k)$ may be negative, for instance if the building is pre-cooled overnight.

Let $u_1(k)$ be the power allocated to making ice (kW), let $u_2(k)$ be the cooling load met by melting ice (kW$_{\text{th}}$), and let $u_3(k)$ be the power consumed by the main chiller (kW). For each $u_i(k)$, $i \in \{1, 2, 3\}$, define the maximum operating point $\bar{u}_i \geq 0$ and the maximum ramp $\Delta \bar{u}_i$.

To enforce the ramping constraint, we append $\mathbf{u}(k-1)$ to the state, defining $\mathbf{x}(k) = [x_1(k), x_2(k), \mathbf{u}(k-1)^T]^T$. This gives the state constraint $\mathbf{x}(k) \in \mathcal{X} = [0, \bar{x}_1] \times \mathbf{R}^4$ and the control constraint $\mathbf{u}(k) \in \mathcal{U}(\mathbf{x}(k))$, where

$$\mathcal{U}(\mathbf{x}(k)) = \left\{ \mathbf{u} \in \mathbf{R}^3 \;\middle|\; u_i \in [0, \bar{u}_i], |u_i - x_{i+2}(k)| \leq \Delta \bar{u}_i \right\}.$$

Let $w_1(k)$ be the building's cooling demand (kW$_{\text{th}}$), and let $w_2(k)$ be the rest of the building's electrical demand (kW): lighting, plug loads, and so on. We assume that the disturbance $\mathbf{w}(k) = [w_1(k), w_2(k)]^T$ is jointly normal and white,

$$\begin{aligned} &\mathbf{w}(k) \sim \mathcal{N}(\bar{\mathbf{w}}(k), \Sigma(k)) \\ &\mathbf{E}\left[(\mathbf{w}(j) - \bar{\mathbf{w}}(j))(\mathbf{w}(k) - \bar{\mathbf{w}}(k))^T\right] = \delta_{jk}\Sigma(k). \end{aligned}$$

We also assume that the building operator has accurate knowledge at stage 0 of $\bar{\mathbf{w}}(k)$ and $\Sigma(k)$ for all $k \in \mathcal{T}$.

Given these definitions, the state evolves according to

$$\begin{aligned} x_1(k+1) &= \beta x_1(k) + \kappa_{\text{ice}}(k) u_1(k) \Delta t - u_2(k) \Delta t / \eta \\ x_2(k+1) &= x_2(k) + [w_1(k) - u_2(k) - \kappa_{\text{main}}(k) u_3(k)] \Delta t \\ x_i(k+1) &= u_{i-2}(k), \quad i \in \{3, 4, 5\} \end{aligned}$$

or more compactly,

$$\mathbf{x}(k+1) = A\mathbf{x}(k) + B(k)\mathbf{u}(k) + G\mathbf{w}(k)$$

where $A \in \mathbf{R}^{5 \times 5}$, $B(k) \in \mathbf{R}^{5 \times 3}$, and $G \in \mathbf{R}^{5 \times 2}$.

In the above, $\beta \in [0, 1]$ is the ice retention rate, $\eta \in [0, 1]$ is the efficiency of heat exchange between the tank and building, and $\kappa_{\text{ice}}(k)$ and $\kappa_{\text{main}}(k)$ are the coefficients of performance of the ice chiller and main chiller, respectively. The time dependence allows the chiller coefficients of performance to vary with outdoor temperature.

We assume that the controller has perfect knowledge of $\mathbf{x}(k)$ at stage k, though a Kalman filter could be added to estimate noisy or unmeasured states.

B. Economics

Letting $c_e(k)$ be the price of energy (\$/kWh) and $p(k) = u_1(k) + u_3(k) + w_2(k)$ be the net power consumption, the energy cost is

$$g_e(k, \mathbf{u}(k), \mathbf{w}(k)) = c_e(k) p(k) \Delta t.$$

Each day's energy prices $c_e(k)$ are published the previous day at 4 PM, so for a 24 hour simulation they are known deterministically.

For $i \in \{1, 2, 3\}$, let $\mathcal{T}_i \subseteq \mathcal{T}$ be the set of stages subject to the i^{th} tier of the demand charge, and let $\bar{p}_i(k)$ be the maximum power consumed in the i^{th} tier of the demand charge prior to stage k, *i.e.*, on $\mathcal{T}_i \cap \{0, \ldots, k-1\}$. The variable $\bar{p}_i(0)$ is initialized with the i^{th} peak demand in the month so far or, if $k = 0$ is the start of the month, with a target demand limit based on historical data. Each $\bar{p}_i$ obeys the dynamics

$$\bar{p}_i(k+1) = \begin{cases} \max\{\bar{p}_i(k), p(k)\}, & k \in \mathcal{T}_i \\ \bar{p}_i(k), & k \notin \mathcal{T}_i. \end{cases}$$

Let $c_d(\mathcal{T}_i) \geq 0$ be the demand price (\$/kW) over $\mathcal{T}_i$. Then the increase in the three-tiered demand charge between stages k_1 and $k_2 \geq k_1$ is

$$g_d\left(\{\mathbf{u}(k), \mathbf{w}(k)\}_{k=k_1}^{k_2}\right) = \sum_{i=1}^{3} c_d(\mathcal{T}_i) \left(-\bar{p}_i(k_1) + \max_{k \in \mathcal{T}_i \cap \{k_1, \ldots, k_2\}} p(k)\right)^+$$

where $(\cdot)^+$ denotes the positive part, $\max\{0, \cdot\}$. In the special case $k_1 = 0$ and $k_2 = N-1$, the function $g_d\left(\{\mathbf{u}(k), \mathbf{w}(k)\}_{k\in\mathcal{T}}\right)$ gives the amount by which the demand charge increases over the entire control horizon.

In order to avoid under- or over-cooling, one could impose the constraint that the cooling demand be met exactly at all stages, but this would preempt desirable strategies like pre-cooling. Instead, we impose a quadratic penalty on deviations of the delivered cooling from the desired cooling:

$$g_u(k, \mathbf{x}(k)) = c_u(k) x_2(k)^2.$$

Here $c_u(k) > 0$, the price of unmet cooling load ($/(kWh$_{\text{th}}$)^2), is modeled as proportional to the building's occupancy at stage k and inversely proportional to its thermal mass.

The demand response program modeled in this paper pays participants a flat price c_{dr} ($/kWh) for avoided energy consumption during the demand response window $\mathcal{T}_{\text{dr}}$, announced by the utility 21 hours in advance. This motivates the cost

$$g_{\text{dr}}(\{\mathbf{u}(k), \mathbf{w}(k)\}_{k\in\mathcal{T}_{\text{dr}}}) = -c_{\text{dr}}\Delta t \sum_{k\in\mathcal{T}_{\text{dr}}} [\hat{p}(k) - p(k)]$$

where $\hat{p}(k)$ is the baseline power consumption (kW) at stage k, as computed by the demand response program administrator.

To discourage tank depletion from day to day, we also impose the cost

$$g_t(\mathbf{x}(N)) = c_t\left(x_1(0) - x_1(N)\right)^+$$

where the initial state $\mathbf{x}(0)$ is given and c_t ($/kWh$_{\text{th}}$) is a tunable parameter.

The net stage cost is the sum of the costs of energy and unmet cooling load:

$$g_k(\mathbf{x}(k), \mathbf{u}(k), \mathbf{w}(k)) = g_e(k, \mathbf{u}(k), \mathbf{w}(k)) + g_u(k, \mathbf{x}(k)).$$

The terminal cost is the sum of the costs of demand, demand response, and tank depletion:

$$\begin{aligned} g_N(\mathbf{x}(N), \{\mathbf{u}(k), \mathbf{w}(k)\}_{k\in\mathcal{T}}) &= g_d\left(\{\mathbf{u}(k), \mathbf{w}(k)\}_{k\in\mathcal{T}}\right) \\ &+ g_{\text{dr}}(\{\mathbf{u}(k), \mathbf{w}(k)\}_{k\in\mathcal{T}_{\text{dr}}}) + g_t(\mathbf{x}(N)). \end{aligned}$$

C. The stochastic optimal control problem

Given an initial state $\mathbf{x}_0$, the total expected cost of a particular policy $\pi = \{\boldsymbol{\mu}_0, ..., \boldsymbol{\mu}_{N-1}\}$, where $\mathbf{u}(k) = \boldsymbol{\mu}_k(\mathbf{x}(k))$, is

$$J_\pi(\mathbf{x}_0) = \mathbf{E}\Bigg[g_N(\mathbf{x}(N), \{\mathbf{u}(k), \mathbf{w}(k)\}_{k\in\mathcal{T}}) + \sum_{k=0}^{N-1} g_k(\mathbf{x}(k), \mathbf{u}(k), \mathbf{w}(k))\Bigg]$$

where the expectation is taken over $\mathbf{w}(k)$ for all $k \in \mathcal{T}$. The stochastic optimal control problem is therefore

$$\begin{aligned} \min_{\pi} \quad & J_\pi(\mathbf{x}_0) \\ \text{s.t.} \quad & \mathbf{x}(k+1) = A\mathbf{x}(k) + B(k)\mathbf{u}(k) + G\mathbf{w}(k) \\ & \mathbf{x}(k+1) \in \mathcal{X} \\ & \mathbf{u}(k) \in \mathcal{U}(\mathbf{x}(k)) \end{aligned} \tag{1}$$

where the constraints apply for all $k \in \mathcal{T}$.

III. Algorithms

Two algorithms are used in this paper: open-loop optimal control (OLOC) and MPC.

A. Open-loop optimal control

To solve a deterministic problem by OLOC, we use the system dynamics to write all states in terms of the previous controls. The resulting problem of minimizing the future costs can be solved offline, simultaneously producing all future states and controls.

We use OLOC for three purposes. First, the demand response baseline[2] is produced at stage 0 by OLOC with the assumption of certainty equivalence, *i.e.*, that all disturbance realizations will equal their expected values. Second, OLOC is used to compute the optimal policy under the oracle information pattern, where all disturbance realizations are known perfectly at stage 0. The oracle policy is not causal, and therefore cannot be implemented in practice. However, it gives a useful lower bound on the cost of daily operations: no causal policy can outperform the oracle. Finally, OLOC is used at each MPC iteration.

B. Model predictive control

The MPC policy is computed implicitly at each stage through online convex optimization. At stage k, a certainty equivalent OLOC problem is solved for the controls $\mathbf{u}(k), \ldots, \mathbf{u}(k+L-1)$, where $L = \min\{H, N-k\}$ and $H \leq N$ is the MPC horizon. Only the first control is implemented, then the system is allowed to evolve according to the true, stochastic dynamics and the process is repeated. The truncated horizon and certainty equivalence assumption make MPC suboptimal for Problem (1).

IV. Simulation

In order to explore the control problem in a realistic economic environment, we conduct Monte Carlo simulation of a New York City building operating under the ConEd rate structures introduced in §II. The simulation day is July 18, the hottest day of 2013 in New York City. Due to the spike in air conditioning loads, demand in the New York State grid reached a record high. The wholesale energy prices in New York City rose to 29.4 ¢/kWh, about an order of magnitude higher than typical. [16] Demand response resources were dispatched system-wide; we assume that the simulated building was called upon to reduce load from 2 to 6 PM. The prices of energy, demand, demand response and under- or over-cooling for the simulation day are shown in Figure 1.

The building studied is the "Large Office" prototype available in [17], a 3-story, 14,200 m^2 office building with occupancy, lighting and plug load schedules following

[2]In practice, ConEd computes each hour's baseline by averaging past consumption in the same hour of a number of similar days. The OLOC method is nonstandard, but captures the intent of approximating the energy that would have been used, absent a demand response event.

Table 1: Time-Invariant Simulation Parameters

Quantity	Units	Value
Control duration, T	hours	24
Time step, Δt	hours	0.5
MPC horizon, H	-	48
Tank capacity, $\bar{x}_1$	kWh_{th}	1,760
Maximum ice chiller power, $\bar{u}_1$	kW	94
Maximum ice chiller ramp, $\Delta\bar{u}_1$	kW	70.5
Maximum ice melt per stage, $\bar{u}_2$	kW_{th}	510
Maximum ice melt ramp, $\Delta\bar{u}_2$	kW_{th}	510
Maximum main chiller power, $\bar{u}_3$	kW	73
Maximum main chiller ramp, $\Delta\bar{u}_3$	kW	54.75
Tank depletion penalty, c_t	$\$/kWh_{th}$	0.18
Ice retention rate, β	-	0.98
Ice extraction efficiency, η	-	0.9
Initial peak demand (all hours), $\bar{p}_1(0)$	kW	200
Initial peak demand (8 AM - 10 PM), $\bar{p}_2(0)$	kW	200
Initial peak demand (8 AM - 6 PM), $\bar{p}_3(0)$	kW	200
Initial state, $(x_1(0), x_2(0))$	kWh_{th}	(176, 0)

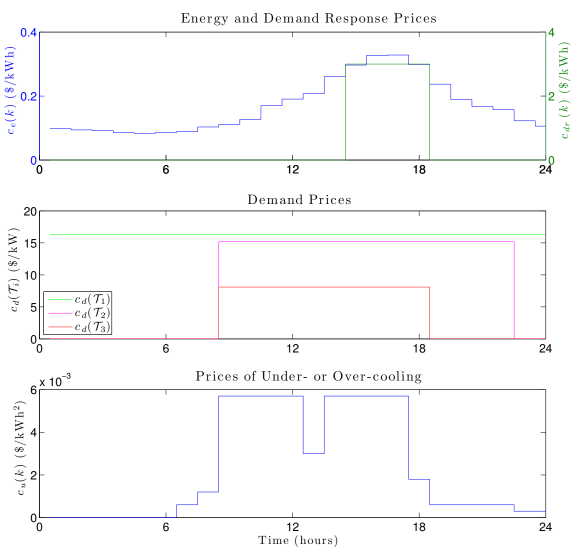

Fig. 1. Prices vs. time. The top plot shows that the demand response price c_{dr} = 3 \$/kWh exceeds the maximum energy price of 0.328 \$/kWh by nearly an order of magnitude. The middle plot shows the off-peak, shoulder and peak tiers of the demand charge ($\mathcal{T}_1, \mathcal{T}_2$ and $\mathcal{T}_3$, respectively) and prices. Note that $\mathcal{T}_3 \subset \mathcal{T}_2 \subset \mathcal{T}_1$, so if the peak demand occurs on $\mathcal{T}_3$ it costs the building operator $c_d(\mathcal{T}_1) + c_d(\mathcal{T}_2) + c_d(\mathcal{T}_3) = 39.54$ \$/kW. The bottom plot shows the price of under- or over-cooling, which varies with building occupancy (hence the decrease over the lunch hour).

ASHRAE Standard 90.1. In order to estimate cooling loads, other electrical loads and chiller coefficients of performance, the building was simulated in TRNSYS under a variety of weather conditions. A regression model was developed to relate the desired quantities to time of day, day of week and outdoor air temperature. The values of the static physical parameters are given in Table 1. Figure 2 shows the time-varying parameters.

It can be shown that all of the cost and constraint functions defined in §II are convex in the controls, so the optimization problems discussed in §III are also convex. The optimization runs are done in CVX [18], [19], calling the SDPT3 solver. [20], [21] An OLOC problem takes about six seconds to run on a 2.8 GHz Intel Xeon processor. A 24-hour MPC simulation with half-hour time steps takes about three minutes.

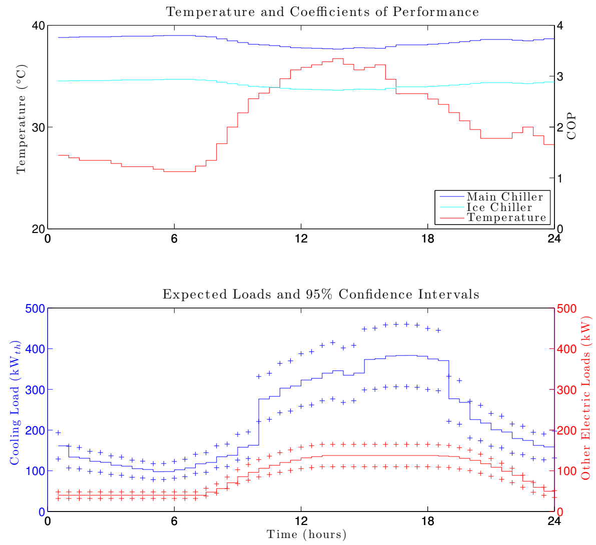

Fig. 2. Time-varying physical parameters. The top plot shows that the simulation day is particularly hot, with temperature peaking at 36.7 °C (98.1 °F). The coefficients of performance of both chillers decrease as the temperature increases, which encourages ice-making overnight. Note that the ice chiller is about 25% less efficient than the main chiller. The bottom plot uses two axes to distinguish thermal from electrical power.

Figure 3 shows the states and controls under the MPC policy for a particular Monte Carlo run. Figure 4 shows the total power and costs under the MPC policy, along with the baseline and the power the building would have consumed if it had no thermal storage. Figure 5 shows cost histograms of the oracle and MPC policies for 200 Monte Carlo runs. Figure 6 shows the power and cost profiles under the MPC policy, with and without including the demand charge in the objective function. Please see the figure captions for details.

V. Discussion

In this paper, we studied the problem of controlling cooling systems under dynamic energy prices, a multi-tiered demand charge, and a demand response program. To our knowledge, this is the first paper to apply MPC to a building under all three objectives. We demonstrated that MPC can successfully navigate the trade-offs inherent in this complex – and really existing – economic environment. This result is enabled by the framework of convex optimization, which admits objectives that, compared to the linear and quadratic costs often used in MPC, more accurately reflect the true economic incentives. As Figures 5 and 6 demonstrate, the value of accurate economic modeling can be substantial.

A second conclusion, relevant to utility regulators and system operators, is that the dynamic energy prices, the demand charge, and the demand response call fight each other for flexible load. Each of these incentives is intended, at least in part, to encourage building operators to shift load away from hours of peak system demand. As Figure 4 shows, however, the midday energy prices are high enough

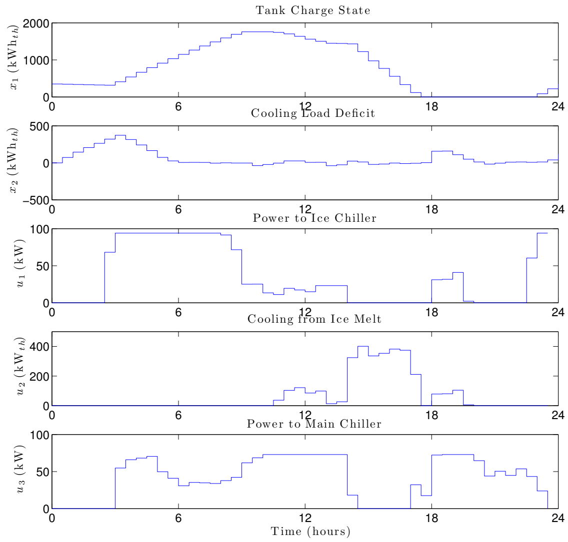

Fig. 3. States and controls under the MPC policy. Plots 2 and 5 show that the building is allowed to drift overnight when occupancy is low. Around 3 AM, both chillers turn on (plots 3 and 5) to fill the ice tank (plot 1) and cool the building back down (plot 2). During the demand response event (hours 14 to 18 in plots 3-5), both chillers are turned off and all cooling is provided by melting ice.

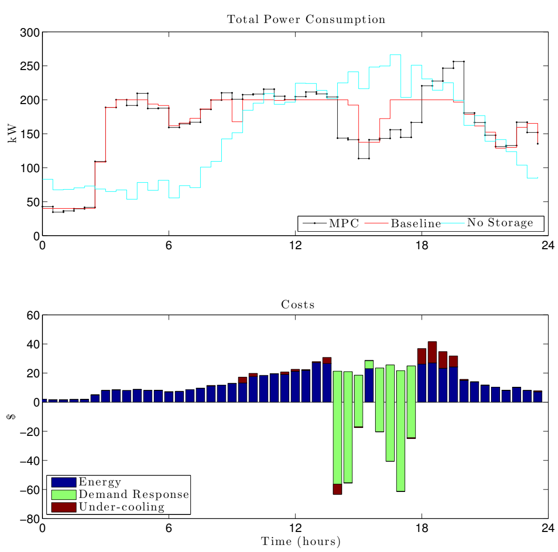

Fig. 4. Power and costs under the MPC policy. The baseline, shown in red in the top plot, is computed through OLOC with demand response prices set to zero. In the baseline calculation, the energy prices from 3-5 PM are high enough to shut down the main chiller. This low midday baseline significantly reduces revenues from the (overlapping) demand response event.

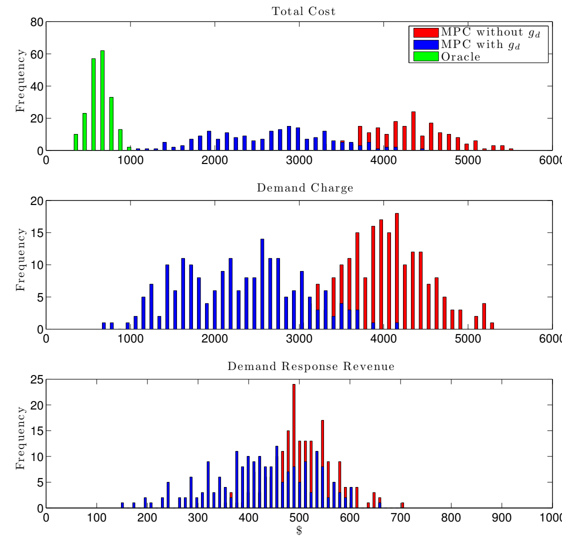

Fig. 5. Histograms of the total cost, demand charge and demand response revenues for 200 Monte Carlo runs. The average total cost of the oracle policy ($630) is 23% that of the MPC policy ($2,660). The MPC policy that neglects the demand charge incurs an average cost increase of 61% ($1,610). In the absence of a demand charge, however, the average demand response revenue (and the corresponding energy reduction) increases by about 20% ($80).

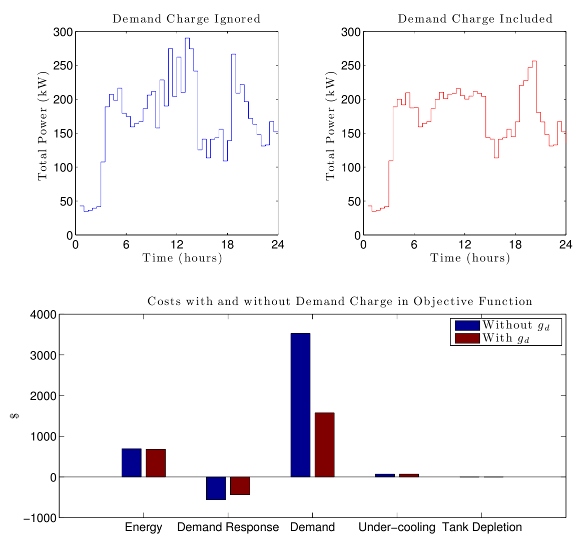

Fig. 6. Power profiles and costs of daily operation under the MPC policy with and without the demand charge in the objective function. Leaving g_d out of the objective function more than doubles the demand cost. For this Monte Carlo run, the value of directly modeling the demand charge is about $1,950, or 53% of the daily operating cost. Note the tension between the demand charge and demand response incentives, however. With no demand charge, the ice tank can be replenished and the building can be pre-cooled between 10 AM and 2 PM. This preparation for the demand response event enables more curtailment, generating more revenue.

to prompt full curtailment of the main chiller, even without a demand response call. This significantly decreases the amount of cooling load that can be curtailed during the demand response event. In Figure 6, a similar tension can be seen between the demand charge and the demand response call. With no demand charge, MPC discovers two strategies between 10 AM and 2 PM: it pre-cools the building, and it refills the ice tank. These strategies enable more curtailment during the 2 PM to 6 PM demand response event, but require that both chillers work at full capacity, causing a spike in demand. With the demand charge in place, the cost of this spike is prohibitive, so the building is not pre-cooled and the tank is not refilled. A rate structure with less internal conflict would make the building easier to control and more responsive to system needs.

An obvious extension of this work is sensitivity analysis. As the tank size $\bar{x}_1$ increases, for example, the trade-offs between the competing economic incentives become less pronounced. This raises the interesting question of optimal tank sizing. For another example, as the covariance matrices $\Sigma(k)$ decrease in norm, the cost of the MPC policy also decreases until, in the limiting case of degenerate disturbances, the oracle and MPC policies agree exactly. In other words, better load predictions give better controllers. Load prediction, which involves both thermal modeling of the building and stochastic modeling of weather and occupant behavior, is not a trivial problem. Quantifying the sensitivity to $\Sigma(k)$ could help building operators decide how much to invest in load prediction.

Another direction for improvement is the physical model developed in §II-A, which is a quasi-steady-state approximation of the underlying chiller and storage dynamics. This approximation, while valid for the long time steps in this paper, worsens as the time step decreases and the nonlinear chiller dynamics become significant. Additional complications arise from the chiller coefficients of performance, which vary nontrivially with chiller load. In practice, chillers are typically allowed either to be off or within some range $[\underline{u}, \bar{u}]$, where $\underline{u} > 0$. These nonconvex constraints would require an appropriate convex relaxation.

There are other extension opportunities. Increasing the simulation horizon to a month or a cooling season would enable more rigorous analysis of the effects of realistic economic modeling. Initialization schemes for the target demand limits $\bar{p}_i$ could be compared. Other mechanisms for providing grid services, such as critical peak pricing, ancillary service markets, or contracts with aggregators, could be explored. Finally, while the MPC algorithm presented in this paper appears to give a good solution to Problem 1, it is suboptimal for the reasons discussed in §III-B. This raises the question of whether a different algorithm might more closely approximate the optimal causal policy.

Acknowledgments

The authors are grateful to Dr. Brandon Hencey of the Air Force Research Laboratory and Brendan Kirby, P.E., for their formative input on this project. We also wish to thank Dr. Santiago Naranjo Palacio of Cornell University, who kindly shared both his building data and his practical perspective on operating thermal storage. This study was funded by the Consortium for Electric Reliability Technology Solutions.

References

[1] "Assessment of Demand Response and Advanced Metering." Federal Energy Regulatory Commission Staff Report, October 2013.

[2] A. Arteconi, N.J. Hewitt and F. Polonara. "State of the Art of Thermal Storage for Demand-Side Management." *Applied Energy* **93** (2012) 371-389.

[3] R. Hensley, J. Newman, and M. Rogers. "Battery Technology Charges Ahead." *McKinsey Quarterly* **3** (2012): 5-50.

[4] J.E. Braun. "A Comparison of Chiller-Priority, Storage-Priority and Optimal Control of an Ice Storage System." *ASHRAE Transactions* **98/1** (1992): 893-902.

[5] A.Afram and F. Janabi-Sharifi. "Theory and applications of HVAC control systems – a review of model predictive control (MPC)." *Building and Environment* **72** (2014): 343-355.

[6] G.P. Henze, R. H. Dodier and M. Krarti. "Development of a Predictive Optimal Controller for Thermal Energy Storage Systems." *HVAC&R Research* **3.3** (1997): 233-264.

[7] J.E. Braun. "Load Control Using Building Thermal Mass." *Journal of Solar Energy Engineering* **125.3** (2003): 292-301.

[8] G.P. Henze, C. Felsmann and G. Knabe. "Evaluation of Optimal Control for Active and Passive Building Thermal Storage." *International Journal of Thermal Sciences* **43** (2004): 173-183.

[9] Y. Ma, F. Borrelli, B. Hencey, A. Packard and S. Bortoff. "Model Predictive Control of Thermal Energy Storage in Building Cooling Systems." *Control Systems Technology, IEEE Transactions on* **20.3** (2012): 796-803.

[10] V.R. Dehkordi and J.A. Candanedo. "A Model-Based Predictive Control Approach for a Building Cooling System with Ice Storage." In *Proceedings of the ASME 2013 Dynamic Systems and Control Conference*.

[11] G.P. Henze, M.J. Brandemuehl, C. Felsmann, H. Cheng, A.R. Florita and C.E. Waters. "Optimization of Building Thermal Mass Control in the Presence of Energy and Demand Charges." *ASHRAE Transactions*, **114** (2), 2008.

[12] J. Ma, J. Qin, T. Salsbury and P. Xu. "Demand Reduction in Building Energy Systems Based on Economic Model Predictive Control." *Chemical Engineering Science* **67** (2012): 92-100.

[13] "Rider M - Day-Ahead Hourly Pricing." *General Rule 24: Service Classification Riders.* Consolidated Edison Company of New York, Inc, 2014.

[14] "Commercial System Relief Program." *Demand Response Program Details,* Consolidated Edison Company of New York, Inc, 2014.

[15] J.J. Kim and S. Kiliccote. "Price Responsive Demand in New York Wholesale Electricity Market using OpenADR." Lawrence Berkeley National Laboratory report 5557E, June 2012.

[16] Day-Ahead Market LBMP, Zone J. *Market & Operational Data*, New York Independent System Operator, 2013.

[17] Commercial Building Prototype Models: "Large Office." *Building Energy Codes Program*, U.S. Department of Energy, 2011.

[18] M. Grant and S. Boyd. "CVX: Matlab Software for Disciplined Convex Programming, version 2.0 beta." http://cvxr.com/cvx, September 2013.

[19] M. Grant and S. Boyd. "Graph Implementations for Nonsmooth Convex Programs, Recent Advances in Learning and Control (a tribute to M. Vidyasagar)," V. Blondel, S. Boyd, and H. Kimura, editors, pages 95-110, Lecture Notes in Control and Information Sciences, Springer, 2008. `http://stanford.edu/~boyd/graph_dcp.html`

[20] K.C. Toh, M.J. Todd, and R.H. Tutuncu. "SDPT3 – a Matlab Software Package for Semidefinite Programming." *Optimization Methods and Software* **11** (1999): 545–581.

[21] R.H Tutuncu, K.C. Toh, and M.J. Todd. "Solving Semidefinite-Quadratic-Linear Programs Using SDPT3." *Mathematical Programming Ser. B* **95** (2003): 189–217.