

July 2018

An Economic Model Predictive Control Framework for Distributed Embedded Battery Applications

Nishith R. Patel

University of Wisconsin-Madison, United States of America, nishith.patel@wisc.edu

James B. Rawlings

University of Wisconsin-Madison, United States of America, james.rawlings@wisc.edu

Matthew J. Ellis

Johnson Controls, Inc., United States of America, matthew.j.ellis@jci.com

Michael J. Wenzel

Johnson Controls, Inc., United States of America, mike.wenzel@jci.com

Robert D. Turney

Johnson Controls, Inc., United States of America, Robert.D.Turney@jci.com

Follow this and additional works at: <https://docs.lib.purdue.edu/ihpbc>

Patel, Nishith R.; Rawlings, James B.; Ellis, Matthew J.; Wenzel, Michael J.; and Turney, Robert D., "An Economic Model Predictive Control Framework for Distributed Embedded Battery Applications" (2018). *International High Performance Buildings Conference*. Paper 254.

<https://docs.lib.purdue.edu/ihpbc/254>

This document has been made available through Purdue e-Pubs, a service of the Purdue University Libraries. Please contact epubs@purdue.edu for additional information.

Complete proceedings may be acquired in print and on CD-ROM directly from the Ray W. Herrick Laboratories at <https://engineering.purdue.edu/Herrick/Events/orderlit.html>

An Economic Model Predictive Control Framework for Distributed Embedded Battery Applications

Nishith R. PATEL^{1*}, James B. RAWLINGS¹, Matthew J. ELLIS², Michael J. WENZEL², Robert D. TURNEY²

¹ University of Wisconsin–Madison, Department of Chemical and Biological Engineering,
Madison, WI, United States of America
nishith.patel@wisc.edu*, james.rawlings@wisc.edu

² Johnson Controls Inc., Technology and Advanced Development, Building Efficiency,
Milwaukee, WI, United States of America
matthew.j.ellis@jci.com, mike.wenzel@jci.com, robert.d.turney@jci.com

* Corresponding Author

ABSTRACT

Since building heating, ventilation, and air conditioning (HVAC) systems are significant consumers of primary energy, considerable efforts are being made to improve energy efficiency and decrease energy costs in these applications. Notably, substantial opportunities in the area of HVAC control exist for decreasing energy costs by shifting loads from peak periods to off-peak periods in the presence of time-varying utility prices. Economic model predictive control (MPC) has been shown to significantly decrease the energy costs of commercial HVAC systems via load shifting. Typically, thermal energy storage (TES) is used for this purpose; however, with batteries becoming less expensive to manufacture, electrical energy storage in batteries is becoming a viable option for load shifting. In this work, large-scale embedded battery applications are considered in which the batteries are directly packaged with airside equipment such as air handler units (AHUs), roof-top units (RTUs), and variable refrigerant flow systems (VRFs).

In this paper, we propose a hierarchical control system framework for the economic optimization of distributed embedded battery units. The architecture considers both building mass storage as well as the electrical energy storage of the battery units. A high-level problem performs an economic optimization over the entire system using aggregate models. The low-level layer is broken into subsystems, each optimizing its local decisions with higher fidelity models. Advantages of this framework include: (i) no iterative communication required between subsystems, (ii) decreased computational complexity in the high-level problem allowing for real-time online implementation, and (iii) management of total demand across the entire system to reduce peak demand charges. We conclude with a simulation study demonstrating the benefits of the proposed control architecture.

1. INTRODUCTION

In 2009, commercial buildings accounted for approximately \$200 billion a year in primary energy expenditures with heating, ventilation, and air conditioning (HVAC) systems comprising the bulk of these expenditures. The Energy Information Administration projects that commercial floor space and primary energy consumption will continue to grow between 2009 and 2035, at 28% and 22%, respectively, hence the expenditures will increase as additional buildings are constructed (Department of Energy, 2012). The magnitude of these values along with their projected growth have made HVAC systems an attractive target for economic optimization in the research community.

1.1 MPC for HVAC

Model predictive control (MPC) is a paradigm in which a system model is used to make predictions over a horizon and an optimization problem is solved to meet control objectives while respecting constraints (Rawlings & Mayne, 2009). The first input of the optimized sequence is implemented, and the problem is resolved at the next time step. MPC has had a tremendous impact in the chemical and petroleum industries (Qin & Badgwell, 2003). In economic MPC, the

objective being optimized is *economic*, such as minimizing cost or maximizing profit (Rawlings & Amrit, 2009).

For HVAC systems, economic MPC has been shown to outperform other control strategies, including preexisting ones, in terms of reducing energy costs by many researchers in the HVAC community (Afram & Janabi-Sharifi, 2014). Power companies often employ a time-of-use pricing structure that charges customers more for electricity purchased during hours of greater demand, such as the afternoon (peak hours), and less during hours of smaller demand, such as overnight (off-peak hours). Power companies use such a model to encourage a more constant load profile so they can operate more efficiently. MPC can decrease energy costs by shifting the HVAC load from peak hours to off peak hours using thermal energy storage (TES) (Mendoza-Serrano & Chmielewski, 2012). Since power companies also levy a peak demand charge to customers based on the maximum rate of power consumption over a period of time, MPC can also coordinate when the various pieces of equipment are drawing power to reduce this peak demand charge incurred (Ma *et al.*, 2012).

Depending on the resources available, several methods of storage can be utilized to shift the load from peak hours to off-peak hours. Passive thermal energy storage (TES), in the form of building mass, can be used to precool the zones in the building during periods of low cost and discharge the cooling energy during period of high cost. (Avci *et al.*, 2013). Active TES, such as water and ice tanks, are more efficient methods of storage thermal energy but requires capital investment (Risbeck *et al.*, 2015). Alternatively, the electricity can be stored directly in batteries to shift the power load as well (Kumar *et al.*, 2018). Most past efforts has been focused on using thermal energy storage.

1.2 Previous Work and Industrial Impact

In previous work, the aim was to develop an implementable optimal control system for a large-scale commercial HVAC system with active and passive TES (Patel *et al.*, 2016). These concepts have already been implemented in industry.

Stanford University recently replaced their central plant which services the heating and cooling needs of their 155-building campus. The Stanford Energy System Innovations (SESI) project replaced an aging cogeneration plant with a new more efficient heat-recovery system (Blair, 2016). They installed hot and chilled water thermal energy storage tanks. The HVAC control system for the campus project was designed by Johnson Controls using economic MPC.

The optimization-based control system has been in operation since December 2015 (Wenzel *et al.*, 2016). The central plant was run in autonomous mode about 90% of the time, including time off-line for plant maintenance. MPC achieved 10% to 15% additional savings in operating costs compared to control by the best team of trained human operators (Stagner, 2016). This large-scale implementation demonstrates the significant economic benefits to applying model-based optimization to large HVAC systems with available storage for load shifting.

1.3 Extension to Batteries

In this work, we consider systems with electrical storage instead of active TES. The cost of manufacturing batteries is becoming cheaper. Batteries can be used for both load shifting to decrease costs and revenue generation if the incentives on the electricity market are appropriate (Dowling *et al.*, 2017). Batteries can generate revenues from programs such as frequency regulation due to their ability to act on a faster time scale to meet demands. However, revenue generation is not the focus of this paper. Instead, the aim is to analyze the use of batteries for load shifting in HVAC systems. In particular, HVAC equipment with embedded batteries are considered.

The rest of the paper is organized as follows. In Section 2, the problem statement for embedded battery applications is presented. In Section 3, the control system to address this problem is proposed. In Section 4, mathematical formulations for each of the MPC optimizations is shown. In Section 5, a sample simulation is performed to demonstrate the control system. Conclusions and future directions are stated in Section 6.

2. EMBEDDED BATTERIES

In embedded battery applications, batteries are directly packaged with individual pieces of HVAC equipment. They can be embedded in any equipment that requires power, ranging from air handler units (AHUs), roof-top units (RTUs),

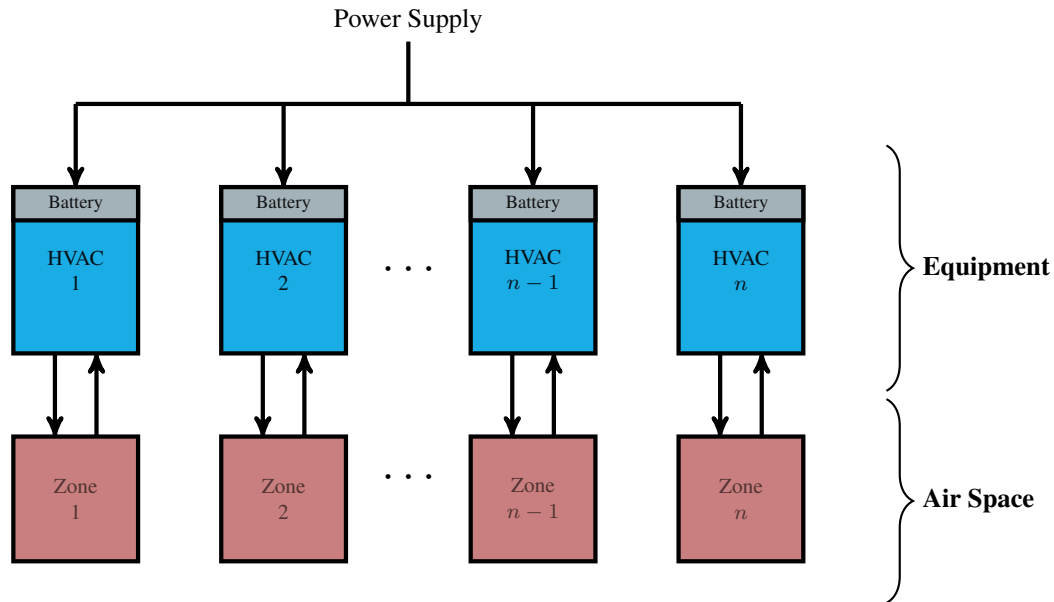


Figure 1: Embedded Battery Applications.

and variable refrigerant flow systems (VRFs) to auxiliary equipment such as fans and pumps.

In this paper, we consider the system depicted in Figure 1: a collection of n zones whose temperatures can be independently manipulated. The formulation is general, including even large campuses with many buildings each with hundreds of zones. Each zone has an associated piece of HVAC equipment for temperature regulation. Each HVAC equipment provides the cooling necessary for the zone by drawing power from the grid. While only cooling is considered in this work, a similar analysis can also be performed for heating. Each piece of HVAC equipment also has an embedded battery that it can charge or discharge as necessary. In this arrangement, the batteries are accessible only to the local unit and not to other units in different zones.

The control problem statement is to design a supervisory system to determine the temperature setpoints for all of the zones as well as the charging and discharging schedules for the embedded batteries and equipment operation that minimize the total operating cost using knowledge of the system dynamics, variable electricity prices, and disturbances. Constraints on temperature for comfort and capacity limits for equipment must be satisfied. Economic MPC can be used to formulate a single optimization problem with this objective and set of constraints along with the system model to make predictions. However, as discussed in the next section, a single optimization for large systems is neither desirable nor practical.

3. CONTROL SYSTEM

Large-scale systems, including university campuses, may contain hundreds or thousands of zones. Solving the MPC problem as a single optimization for such applications is often infeasible in real-time since solution times exceed the sample time of the supervisory controller. Furthermore, these centralized control systems are difficult to maintain.

Such systems have motivated the need to create decompositions of the centralized problem for real-time online implementations. In this section, we propose a decomposition to address these needs.

3.1 Decomposition

In the proposed decomposition, the MPC layer that sits above the existing regulatory control is separated into high and low levels as depicted in Figure 2. First, the individual zones in the systems are grouped into a small set of subsystems. The high-level problem performs a system-wide economic optimization using disturbance forecasts, electricity pricing

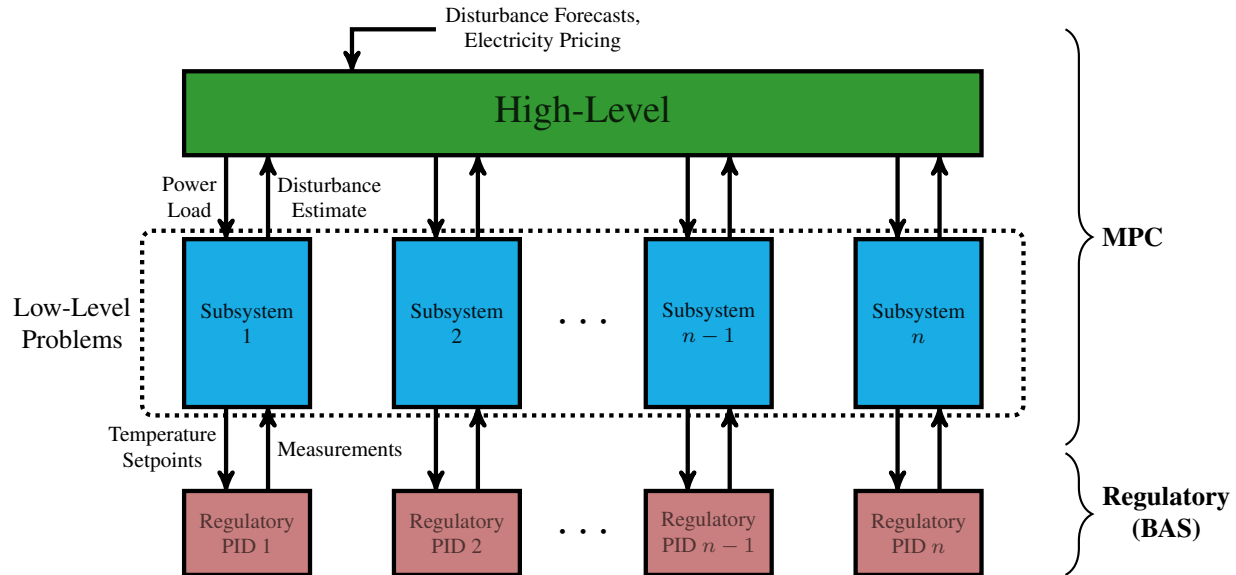


Figure 2: Proposed Control System.

data, and aggregate models of these subsystems. It computes the power allocation for each subsystem and sends that down to the low-level layer. Each subsystem has a low-level controller that uses this load along with more detailed zone models to determine all of the zone setpoints for that subsystem which minimize energy usage subject to satisfying the constraints. These setpoints are sent down to the regulatory level. Measurements collected are passed up to the MPC layer and used to update the disturbance estimate for the next time step. The details of each of these problem formulations are discussed in the next section.

3.2 Advantages over Alternative Strategies

The proposed system has several advantages. The main sources of coupling are handled in the high-level problem, such as the total load and peak demand charge based on the maximum rate of power consumption for the entire system. The high-level problem uses aggregate zone and battery models to reduce the computational burden so the problem is tractable online.

The subsystems in the low-level problem can be chosen to be decoupled (e.g., by building) if zones in one subsystem do not exchange heat with zones in other subsystems. Decoupling these low-level problems can drastically reduce the communication required since iterations between them are not required as they would be in other formulations of distributed MPC (Scherer *et al.*, 2013; Lamoudi *et al.*, 2011). Communication can also be a limiting factor with the existing network infrastructure of HVAC systems, further increasing the solve times.

Additionally, the problem formulation is general enough to handle various types of HVAC applications from air handler units (AHUs) to roof top units (RTUs) to variable refrigerant flow (VRF) systems. The corresponding model components change from application to application but the same problem decomposition remains.

4. PROBLEM FORMULATIONS

Several choices can be made for choosing the model equations to represent system dynamics, ranging from complex and more accurate to simple and less accurate. Since there is feedback to account for model errors and disturbances, simple linear models have been shown often to be effective even when controlling complex, nonlinear chemical processes due to feedback and maintaining the system near the same operating point (Qin & Badgwell, 2003). In this section, we present a computationally tractable formulation; however, the specific models can be modified to increase accuracy should it be necessary for the application.

4.1 System Model

The dynamics of cooling a single zone or building are represented using the energy balance given by

$$C \frac{dT}{dt} = -H(T - T_a) - \dot{Q}_c + \dot{Q}_{\text{other}} \quad (1)$$

The notation used throughout this paper is defined in the nomenclature section at the end. This particular single-state model lumps all zone mass and air properties into a single temperature variable. However, other models can also be used. Higher-order linear models have also been used to describe the temperature dynamics. The zone air and mass temperatures could be modeled separately. Additionally, the mass temperature can be further separated into shallow and deep masses. These choices increase the number of states for the MPC problem and also increase the complexity of the system identification problem. In this paper, lumped temperature model is considered both for zones and also aggregated for buildings.

The temperature model given in Equation (1) is extended to handle systems with multiple zones by considering the coupling between zones as given by

$$C_i \frac{dT_i}{dt} = -H_i(T_i - T_a) - \sum_{j \neq i} \beta_{ij}(T_i - T_j) - \dot{Q}_{c,i} + \dot{Q}_{\text{other},i} \quad (2)$$

in which β_{ij} characterizes the degree of coupling between zones i and j . If zones i and j are not adjacent, then $\beta_{ij} = 0$. Similarly, $\beta = 0$ in the high-level problem since buildings do not exchange heat with one another.

Electrical storage in the batteries is modeled using

$$\frac{dB}{dt} = -\sigma B + \sqrt{\eta_b} \dot{W}_{\text{in}} - \frac{1}{\sqrt{\eta_b}} \dot{W}_{\text{out}} \quad (3)$$

in which σ represents the decay term and η_b represents the round-trip efficiency.

For use in the MPC optimization problems, all models are converted to state-space form and discretized according to the sample time of the controller to yield discrete-time state-space models.

4.2 High-Level Optimization Problem

In the high-level optimization, the objective is to minimize the total energy cost. Both the time-varying price of electricity, c_k , as well as the peak demand charge c_{peak} are considered as shown in Equation (4).

$$\begin{aligned} \min_{\mathbf{x}, \mathbf{u}, \dot{W}_{\text{peak}}} \quad & \sum_{k=0}^{N-1} c_k \dot{W}_{\text{grid},k} \Delta + c_{\text{peak}} \dot{W}_{\text{peak}} \\ \text{s.t.} \quad & C_b \frac{dT_b}{dt} = -H_b(T_b - T_a) - \dot{Q}_{c,b} + \dot{Q}_{\text{other},b} \\ & \frac{dB}{dt} = -\sigma B + \sqrt{\eta_b} \dot{W}_{\text{in}} - \frac{1}{\sqrt{\eta_b}} \dot{W}_{\text{out}} \\ & \dot{W}_{\text{HVAC},k} = \eta_{\text{HVAC}} \sum_b \dot{Q}_{c,b,k} \\ & \dot{W}_{\text{grid},k} = \dot{W}_{\text{HVAC},k} + \dot{W}_{\text{in},k} - \dot{W}_{\text{out},k} \leq \dot{W}_{\text{peak}} \\ & \dot{W}_{\text{peak,past}} \leq \dot{W}_{\text{peak}}, 0 \leq \sum_b \dot{Q}_{c,b,k} \leq \dot{Q}_{\text{HVAC,max}} \\ & 0 \leq B_{k,b} \leq B_{\text{max},b} \\ & T_{\text{min}} \leq T_b \leq T_{\text{max}} \end{aligned} \quad (4)$$

In this example, the subsystems are chosen to be the individual buildings. Aggregate building temperature dynamics are described using the lumped temperature model. An aggregate battery model per building is also used to model the

available electrical storage. The power purchased from the electricity market, \dot{W}_{grid} , must satisfy the load generated by running the HVAC equipment as well as charging and discharging the batteries. Constraints are also added to respect physical limits such as equipment capacity and battery capacity as well as comfort bounds to ensure the building temperatures are within the prespecified limits to keep occupants comfortable. The power loads computed from this optimization are sent to the low-level problems.

4.3 Low-Level Optimization Problem

For each subsystem, the low-level optimization problem solved is given by

$$\begin{aligned}
 & \min_{\mathbf{x}, \mathbf{u}} \quad W_{\text{total}, N} \\
 \text{s.t.} \quad & C_i \frac{dT_i}{dt} = -H_i(T_i - T_a) - \sum_{j \neq i} \beta_{ij}(T_i - T_j) - \dot{Q}_{c,i} + \dot{Q}_{\text{other},i} \\
 & \frac{dB_i}{dt} = -\sigma_i B_i + \sqrt{\eta_{b,i}} \dot{W}_{\text{in},i} - \frac{1}{\sqrt{\eta_{b,i}}} \dot{W}_{\text{out},i} \\
 & \frac{dW_{\text{total}}}{dt} = \sum_i \eta_{\text{HVAC},i} \dot{Q}_{c,i} + \dot{W}_{\text{in},i} - \dot{W}_{\text{out},i} \\
 & T_{\min} \leq T_i \leq T_{\max} \\
 & 0 \leq B_{k,i} \leq B_{\max,i} \\
 & \frac{W_{\text{total},k+1} - W_{\text{total},k}}{\Delta} \leq \dot{W}_{\text{HighLevel},k} \\
 & W_{\text{total},k+1} - W_{\text{total},k} \geq 0 \\
 & \dot{Q}_{c,i} = f(T_i, T_{\text{sp},i})
 \end{aligned} \tag{5}$$

The objective of each low-level controller is to minimize energy usage. It must meet all constraints while not exceeding the power allocation from the high-level problem. Constraints are softened to prevent infeasibility, however violations are used as feedback to adjust the disturbance estimate for the subsequent time step to update the forecasts in the high-level problem to improve its accuracy. As shown in Equation (5), detailed zone models are used in the low-level problem to model temperature dynamics and batteries as opposed to the aggregate models using in the high-level problem.

One way to aggregate low-level model parameters into a subsystem model is given as follows. The average building temperature is defined by summing over all zones in that building using

$$T_b = \frac{\sum_i C_i T_i}{\sum_i C_i}$$

Then, the aggregate building parameters for the high-level model are obtained by using

$$C_b = \sum_i C_i \quad H_b = \sum_i H_i$$

Another model required for the low-level problem is the cooling duty model $\dot{Q}_{c,i} = f(T_i, T_{\text{sp},i})$. Since dynamics of regulatory temperature controllers are sluggish, they cannot be ignore and assumed to be instantaneous in the MPC layer. To obtain the zone cooling duty as a function of the setpoint and temperature, two models are required:

1. A model of the regulatory zone temperature controller to determine the output of the regulatory controller as a function of zone temperature and temperature setpoint

$$v_{\text{air},i} = f_1(T_i, T_{\text{sp},i})$$

2. An energy balance relating the regulatory controller output to the actual cooling duty

$$\dot{Q}_{c,i} = f_2(v_{\text{air},i})$$




Figure 3: Ambient Temperature and Electricity Pricing Data.

As an example, if an ideal PI controller and a linear relationship between $\dot{Q}_{c,i}$ and $v_{\text{air},i}$ is assumed, a linear cooling duty controller model is obtained as

$$\dot{Q}_{c,i} = \dot{Q}_{\text{ss},i} + K_{c,i} \left[\varepsilon_i + \frac{1}{\tau_{1,i}} \int_0^t \varepsilon_i(t') dt' \right]$$

$$\varepsilon_i = T_{\text{sp},i} - T_i$$

If such a linear model is not accurate for a specific application, a nonlinear cooling duty model can be used at the expense of increasing complexity.

5. SIMULATION RESULTS

To demonstrate the costs savings of using the proposed control system on a system with embedded batteries, a simulation study is performed. The system studied consists of ten five-zone buildings; each of the 50 zone HVAC units has its own embedded battery.

The weather and electricity pricing data used in the simulation are shown in Figure 3. The weather data were obtained from a city in the southeastern US for a three-day period in the summer. The time-varying electricity price and peak demand charge data are representative of data from Johnson Controls.

Optimizations were solved using Gurobi 6.0 via MATLAB R2017b on a machine with 8GB RAM and 2.66GHz Intel Core 2 Quad Processor Q8400. The high-level problem took approximately 10 seconds to solve. Each of the low-level subproblems took about 4–6 seconds to solve, and they can be computed in parallel and easily extended to handle even larger applications. Hence, with a time step of 15 minutes, the system can be implemented online with these real-time computational requirements.

Figure 4 shows the results of the high-level optimization problem. The top graph shows how the HVAC equipment demand is met using a combination of production (purchase) from the electricity market and charging/discharging of the electrical storage (batteries). With electrical storage available, more power is purchased during the night to charge the batteries so that less power is purchased during the peak period when electricity is more expensive. The flatness

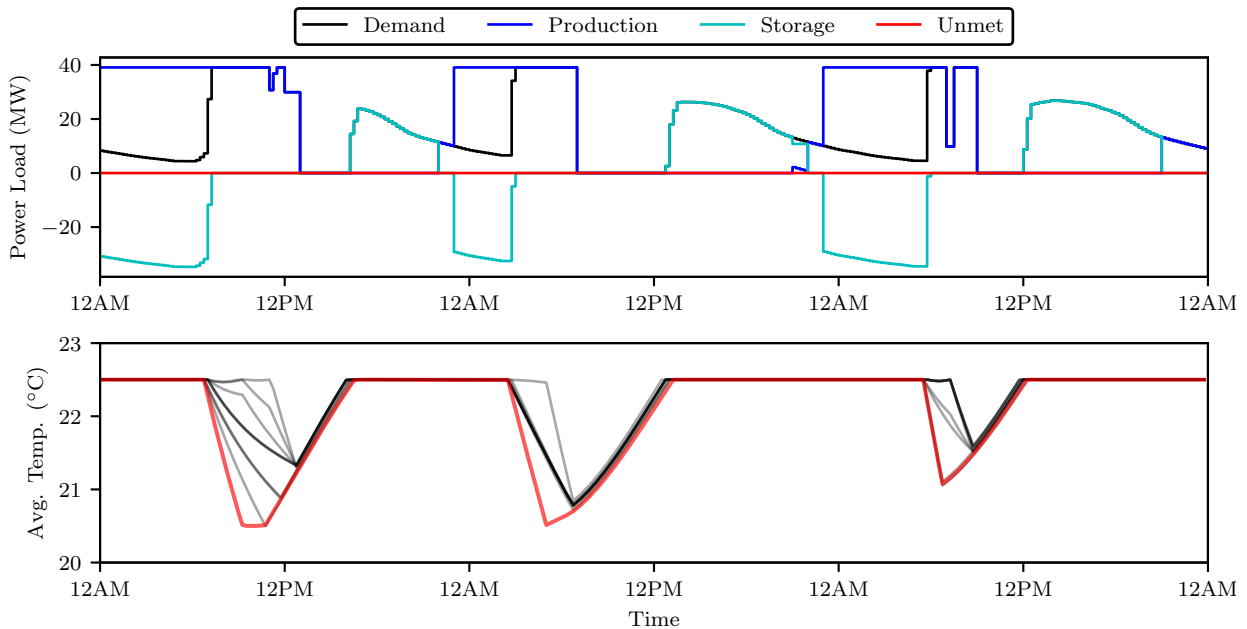


Figure 4: High-Level Optimization Result.

of the power purchase profile from the high-level problem is a result of the peak demand charge in the objective. Eliminating spikes in power purchased reduces the peak demand cost incurred. The bottom graph shows the average building temperatures during this period. With passive TES available, the economically optimal strategy is to precool the buildings before the peak period when electricity prices are highest. The red line highlights one particular building profile in the campus. Both the electrical storage and passive TES are utilized in this system.

Figure 5 shows the results of the low-level optimization problem. The top graph shows the zone temperatures with the comfort zone between 20.5 and 22.5 °C. The middle graph shows the temperature setpoints sent to the zones to achieve the profile. Notice that MPC is able to account for the sluggishness of building temperature responses by preemptively sending a lower setpoint when precooling since it has a model of the regulatory response. The bottom graph in each figure shows the charging and discharging profiles of the individual zone batteries. The batteries are charged during the night and early morning hours when power is cheaper and discharged during the peak hours when it is more expensive. The low-level problem determines the exact schedules to prevent a spike in the total peak demand. The red line in each plot highlights one particular zone in the system for clarity. MPC utilizes the electrical storage first, then the passive TES since it is more efficient and able to retain energy for longer time periods.

Energy cost savings range from 5–40% depending on the amount of energy storage available, the efficiency of the available storage, and the market incentives to shift load. In the presented example with modest battery capacities, the savings is 10% due to the utilization of the passive thermal energy and electrical storage. With larger battery capacities, the savings increase and the optimizer decreases the usage of the passive TES since it is less efficient compared to electrical storage.

6. CONCLUSIONS

An economically optimal method of temperature regulation for HVAC systems with embedded batteries was proposed. The control system framework considers the optimization of both thermal and electrical energy storage for a large-scale application. Economic MPC provides significant cost savings over steady-state setpoint-based operation.

Future research includes performing an economic analysis of the effect of battery size on savings to determine how large a battery must be in order to provide a desirable payback period for the capital investment.

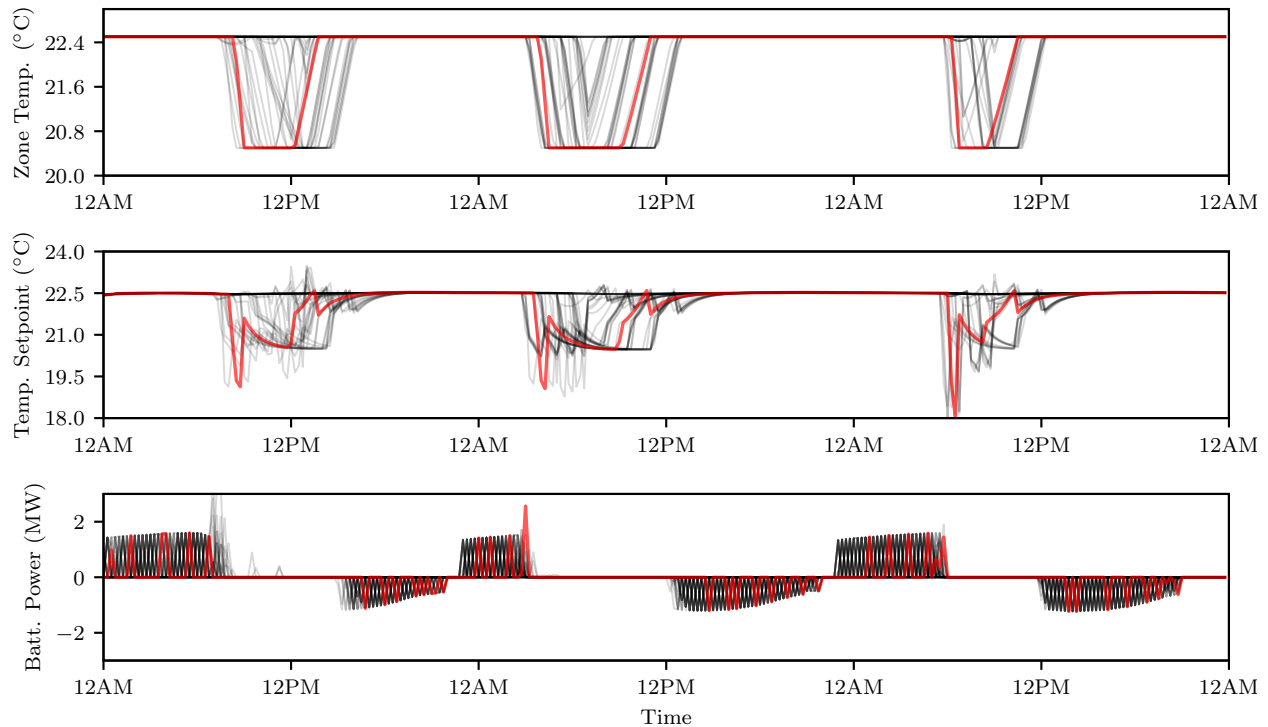


Figure 5: Low-Level Optimization Result.

NOMENCLATURE

Indices

b	building index (high-level)
i	zone index (low-level)
k	time index

Variables

T	temperature
T_a	ambient temperature
$T_{sp,i}$	zone temperature setpoint
\dot{Q}_{other}	external load, radiation, or disturbance
\dot{Q}_c	cooling rate delivered
\dot{W}_{total}	total amount of power delivered
\dot{W}_{peak}	peak power demand
B	power level in battery
\dot{W}_{in}	charging rate of battery
\dot{W}_{out}	discharging rate of battery
$\dot{W}_{HighLevel}$	power allocation from high-level problem
$T_{HighLevel}$	building temperature from high-level
ε_i	tracking error
$v_{air,i}$	air flow rate in AHU

Parameters

Δ	sample time of controller
N	horizon length
c_k	cost of electricity at time k
c_{peak}	peak demand charge
C	thermal capacitance
H	scaled ambient heat transfer coefficient
β_{ij}	scaled inter-zone heat transfer coefficient
$K_{c,i}$	scaled zone PI controller gain
$\tau_{I,i}$	integral time constant for zone PI controller
$\dot{Q}_{ss,i}$	steady-state rate of cooling
η_{HVAC}	inverse of the aggregate COP
$\dot{Q}_{HVAC,max}$	max cooling capacity of HVAC system
$\dot{W}_{peak,past}$	peak power demand previously achieved
σ	decay constant for battery
η_b	round-trip efficiency of battery
B_{max}	max capacity of battery
T_{min}	lower bound of comfort region
T_{max}	upper bound of comfort region
μ	penalty on energy usage

REFERENCES

- Afram, A., & Janabi-Sharifi, F. (2014). Theory and applications of HVAC control systems—a review of model predictive control (MPC). *Building and Environment*, 72, 343–355.
- Avci, M., Erkoc, M., Rahmani, A., & Asfour, S. (2013). Model predictive HVAC load control in buildings using real-time electricity pricing. *Energy and Buildings*, 60, 199–209.
- Blair, S. (2016). Editors' choice and best energy/industrial: Stanford Energy System Innovations. *Engineering News-Record*. Retrieved from <http://www.enr.com/articles/39005-editors-choice-best-energyindustrial-stanford-energy-system-innovations>
- Department of Energy. (2012, March). *2011 Buildings Energy Data Book*. <http://buildingsdatabook.eren.doe.gov/ChapterIntro3.aspx>.
- Dowling, A. W., Kumar, R., & Zavala, V. M. (2017). A multi-scale optimization framework for electricity market participation. *Applied Energy*, 190, 147–164.
- Kumar, R., Wenzel, M. J., Ellis, M. J., ElBsat, M. N., Drees, K. H., & Zavala, V. M. (2018). A stochastic model predictive control framework for stationary battery systems. *IEEE Transactions on Power Systems*.
- Lamoudi, M. Y., Alamir, M., & Béguery, P. (2011). Distributed constrained model predictive control based on bundle method for building energy management. In *50th IEEE Conference on Decision and Control and European Control Conference* (pp. 8118–8124). Orlando, FL, USA.
- Ma, J., Qin, J., Salsbury, T., & Xu, P. (2012). Demand reduction in building energy systems based on economic model predictive control. *Chemical Engineering Science*, 67, 92–100.
- Mendoza-Serrano, D. I., & Chmielewski, D. J. (2012). HVAC control using infinite-horizon economic MPC. In *51st IEEE Conference on Decision and Control* (pp. 6963–6968). Maui, Hawaii, USA.
- Patel, N. R., Rawlings, J. B., Wenzel, M. J., & Turney, R. D. (2016, July 11–14.). Design and application of distributed economic model predictive control for large-scale building temperature regulation. In *4th International High Performance Buildings Conference at Purdue*. West Lafayette, IN.
- Qin, S. J., & Badgwell, T. A. (2003). A survey of industrial model predictive control technology. *Control Engineering Practice*, 11(7), 733–764.
- Rawlings, J. B., & Amrit, R. (2009). Optimizing process economic performance using model predictive control. In *Nonlinear Model Predictive Control* (Vol. 384, pp. 119–138). Berlin: Springer.
- Rawlings, J. B., & Mayne, D. Q. (2009). *Model Predictive Control: Theory and Design*. Madison, WI: Nob Hill Publishing.
- Risbeck, M. J., Maravelias, C. T., Rawlings, J. B., & Turney, R. D. (2015). Cost optimization of combined building heating/cooling equipment via mixed-integer linear programming. In *American Control Conference* (p. 1689–1694). Chicago, IL.
- Scherer, H., Pasamontes, M., Guzmán, J., Álvarez, J., Camponogara, E., & Normey-Rico, J. (2013). Efficient building energy management using distributed model predictive control. *Journal of Process Control*.
- Stagner, J. (2016). *Enterprise optimization solution (EOS) cost savings vs. manual plant dispatching*. Report on Central Energy Facility, Stanford Energy System Innovations.
- Wenzel, M. J., Turney, R. D., & Drees, K. H. (2016). Autonomous optimization and control for central plants with energy storage. In *4th International High Performance Buildings Conference at Purdue*. West Lafayette, IN.

ACKNOWLEDGMENT

The authors gratefully acknowledge the financial support of the NSF through grant #CTS-1603768 and thank Johnson Controls, Inc. for providing sample data, models, and research funding.