

2016

# Design and Application of Distributed Economic Model Predictive Control for Large-Scale Building Temperature Regulation

Nishith R. Patel

*University of Wisconsin-Madison, Madison, WI, nishith.patel@wisc.edu*

James B. Rawlings

*University of Wisconsin-Madison, Madison, WI, james.rawlings@wisc.edu*

Michael J. Wenzel

*Johnson Controls, Milwaukee, WI, mike.wenzel@jci.com*

Robert D. Turney

*Johnson Controls, Milwaukee, WI, Robert.D.Turney@jci.com*

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

---

Patel, Nishith R.; Rawlings, James B.; Wenzel, Michael J.; and Turney, Robert D., "Design and Application of Distributed Economic Model Predictive Control for Large-Scale Building Temperature Regulation" (2016). *International High Performance Buildings Conference*. Paper 181.  
<http://docs.lib.purdue.edu/ihpbc/181>

This document has been made available through Purdue e-Pubs, a service of the Purdue University Libraries. Please contact [epubs@purdue.edu](mailto: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>

# Design and Application of Distributed Economic Model Predictive Control for Large-Scale Building Temperature Regulation

Nishith R. PATEL<sup>1\*</sup>, James B. RAWLINGS<sup>1</sup>, Michael J. WENZEL<sup>2</sup>, Robert D. TURNEY<sup>2</sup>

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

<sup>2</sup> Johnson Controls Inc., Technology and Advanced Development, Building Efficiency,  
Milwaukee, WI, United States of America  
mike.wenzel@jci.com, robert.d.turney@jci.com

\* Corresponding Author

## ABSTRACT


Although recent research has suggested model predictive control as a promising solution for minimizing energy costs of commercial buildings, advanced control systems have not been widely deployed in practice. Large-scale implementations, including industrial complexes and university campuses, may contain thousands of air handler units each serving a multiplicity of zones. A single centralized control system for these applications is not desirable. In this paper, we propose a distributed control system to economically optimize temperature regulation for large-scale commercial building applications. The decomposition strategy considers the complexities of thermal energy storage, zone interactions, and chiller plant equipment while remaining computationally tractable. One of the primary benefits of the proposed formulation is that the low-level airside problem can be decoupled and solved in a distributed manner; hence, it can be easily extended to handle large applications. Peak demand charges, a major source of coupling, are included. The interactions of the airside system with the waterside system are also considered, including discrete decisions, such as turning chillers on and off. To deploy such a control scheme, a system model is required. Since using physical knowledge about building models can greatly reduce the number of parameters that must be identified, grey-box models are recommended to reduce the length of expensive identification testing. We demonstrate the effectiveness of this control system architecture and identification procedure via simulation studies.

## 1. INTRODUCTION

### 1.1 Motivation

Commercial buildings account for 20% of the total U.S. energy consumption, corresponding to roughly \$200 billion a year in primary energy expenditures (Department of Energy, 2012). 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. Average energy prices, on the other hand, have been, and are expected to remain, relatively stable; therefore, the amount spent on energy in commercial buildings will continue to increase significantly. Given the significance of these energy cost values and their projected growth, buildings have become a clear target for control strategies designed to reduce costs and improve efficiency, particularly in the area of temperature control.

Almost all heating and cooling systems in commercial buildings and educational facilities use on/off and proportional-integral-derivative (PID) controllers for control of their heating, ventilation, and air conditioning (HVAC) systems (Afram & Janabi-Sharifi, 2014). These systems rely on controllers whose goal is to converge to the desired temperature set point. However, a better goal is to minimize the total energy cost. In a utility market with time-varying prices, the potential exists for cost savings by temporally shifting heating or cooling loads using thermal energy storage. To achieve these savings, predictive optimization is required using a model of the system for forecasting. Load shifting decreases the burden on power plants during peak hours, allowing them to operate more efficiently. Furthermore, chillers can operate more efficiently at night than during the day.



**Figure 1: Schematic of large-scale commercial application with airside system shown on the left and waterside system shown on the right.**

## 1.2 Model Predictive Control

Model predictive control (MPC) is an advanced control method that has been highly successful over the past two decades with thousands of applications in the chemical and petroleum industries alone (Qin & Badgwell, 2003). MPC relies on a model of the system to predict the process variables based on the actions taken by the controller (Rawlings & Mayne, 2009). At each step, an online optimization problem is solved. In economic MPC, the objective for this optimization is to minimize total cost (Rawlings & Amrit, 2009). Numerous studies have shown that MPC outperforms existing control systems due to its ability to forecast into the future and anticipate events before they occur (Ma *et al.*, 2012; Mendoza-Serrano & Chmielewski, 2012). MPC enables shifting of the energy load from peak hours to off-peak hours by using the mass of the building for passive thermal energy storage (TES) (Oldewurtel *et al.*, 2010; Avci *et al.*, 2013). Active TES, e.g., chilled water tanks, can also be used to further facilitate load shifting. Through a combination of active and passive storage systems, energy costs can be decreased by concentrating equipment usage to times of low resource prices while maintaining comfort limits within the building.

## 1.3 Large-Scale Commercial Applications

Economically optimal control systems have not been deployed widely in the HVAC industry. One fundamental obstacle to the successful deployment of MPC is the large number of zones. To implement MPC in HVAC systems, an optimization problem must be solved in a reasonably short time (on the order of a few minutes). Campus wide implementations may contain hundreds of buildings and thousands of air handler units each serving tens of zones. A schematic of such a large-scale application is shown in Figure 1. The airside system consists of all zones and the air-handler units (AHUs) that are responsible for temperature regulation. The waterside system consists of the equipment, such as chillers, that are required to provide cooling and/or heating. The control objective is to determine temperature setpoints for all zones in the airside system and to operate the waterside equipment to meet the corresponding load. A single, combined control system for these applications is impractical and undesirable as the number of zones and equipment increases because the resulting single optimization problem is difficult to maintain and also too large to solve in real time.

Distributed MPC can be implemented so that the large combined problem can be divided into smaller problems, which can be solved in parallel to reduce the overall computation time (Rawlings & Mayne, 2009; Rawlings & Stewart, 2008). Iterative methods are typically used in distributed MPC (Scherer *et al.*, 2013; Lamoudi *et al.*, 2011). However, iterative

methods with many information exchanges between the distributed controllers are undesirable in practice because the controllers must wait on other controllers at each iteration, which can result in a greater computational time than that of solving the combined problem directly. Additionally, there are potential practical limitations to information exchanges based on existing HVAC communication protocols.

Power companies not only charge customers according to their total energy usage with time varying electricity costs, but they often also levy a peak demand charge based on the peak power usage during the month. Since peak demand charges comprise a significant portion of total cooling costs in buildings, excluding these charges leads to suboptimal performance with respect to minimizing energy costs. Existing distributed methods do not address these peak demand charges. Coordination is required; otherwise, zone controllers may cool at similar times leading to a high peak power usage. Furthermore, these existing methods often do not consider active TES or include detailed waterside equipment models, where significant cost savings can be achieved.

In this paper, we design a distributed economic MPC architecture for building temperature regulation to address these issues. The formulation is general enough to handle a large class of HVAC systems including interactions between the airside and waterside systems. The overall single optimization problem is decomposed into smaller subproblems for large-scale applications. Aggregate models are used in the high-level problem to reduce the computational burden. Iterations and communication between the low-level airside subsystems are not required. Demand response is addressed for optimal performance in the presence of peak demand charges. Regulatory controller dynamics are modeled to account for the slow closed-loop responses to temperature setpoint changes in commercial buildings.

## 2. DECOMPOSITION

### 2.1 Problem Statement

The goal is to develop a solution to the overall problem: given a forecast of electricity prices, ambient conditions, and disturbance loads, decide how to operate the HVAC system to minimize the total operating cost. The decision variables include temperature setpoints sent to the zone controllers, chiller operation schedule and storage tank operation schedule. In making these decisions, comfort bounds on the zone temperatures, capacity constraints on the equipment (air-handlers, chillers, and storage tank), and switching constraints for chillers must be respected.

Given this problem statement, we can formulate a single overall optimization problem for the combined airside and waterside system. As discussed in Section 1, solving this single optimization problem becomes difficult for large-scale applications. The airside system can consist of thousands of zones, and the waterside system can contain multiple pieces of equipment. Inclusion of the discrete decisions (turning equipment on or off) results in a mixed-integer optimization problem, which further increases the problem complexity. Since the single problem cannot be solved in real-time implementations, a decomposition is required to formulate smaller, manageable optimization problems.

### 2.2 Control Architecture Design

Our decomposition of the overall problem for HVAC control is shown in Figure 2. The MPC layer sits on top of the regulatory layer, which consists of the building automation system (BAS) and PID controllers for the airside and waterside systems. The design is expected to work with any existing BAS, allowing one to retrofit an existing HVAC control architecture with optimization. Only setpoint and measurement communication between the layers is required. In addition to a high- and low-level separation for the MPC problem, we split the low-level problem into *airside* and *waterside* following industrial convention.

For large-scale applications with hundreds or thousands of zones, the low-level airside problem can be further decomposed into separate distributed MPC controllers with each controlling a subsystem of the entire airside system as shown in Figure 2. The information exchanges between the layers are identical for all of the distributed airside controllers. The high-level problem contains an aggregate model of each airside subsystem and allocates cooling to each subsystem. The subsystem MPC controllers contain more detailed zone-level models. The choice of dividing the airside system into separate subsystems is made for computational reasons so that the subsystem MPC problems can be solved within a few minutes. Each subsystem can be anything from a single zone to a collection of zones to an entire building or even multiple buildings. If there is significant coupling between the subsystems (e.g., heat exchange between zones in separate subsystems), then performance may deteriorate as the controllers do not coordinate their



### 3. PROBLEM FORMULATIONS

#### 3.1 System Model

The dynamics of cooling a single zone or building are described by the energy balance

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

This model lumps all zone mass and air properties into a single zone temperature. Other models can also be used to describe the airside system dynamics. Differential equations can be written to model the zone air and mass temperatures separately, and the mass temperature can be further separated into a shallow mass temperature and deep mass temperature. However, for simplicity, we considered the lumped model in the simulation studies.

The system model given in Equation (1) can be extended to handle buildings with multiple zones. To this end, the coupling between the air spaces needs to be included. As energy is transferred due to a temperature gradient, the coupled model is 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  $\beta_{ij}$  characterizes the degree of coupling between zones  $i$  and  $j$ . If zones  $i$  and  $j$  are not adjacent, then  $\beta_{ij} = 0$ . Since buildings are separate and do not directly exchange energy,  $\beta = 0$  in the high-level problem. Active TES is modeled using first order linear dynamics

$$\frac{ds}{dt} = -\sigma s + \dot{Q}_{\text{storage}} \quad (3)$$

For convenience, we convert all models to state-space form for use in the MPC optimization problem. Equation (2) is combined with Equation (3) and discretized according to the sample time of the controller to yield a discrete-time state-space model for the high-level problem. Similarly, using Equation (2) and the cooling duty model, which is discussed later, a discrete-time state-space model is obtained for the low-level airside problem.

#### 3.2 High-Level Optimization Problem

As shown in Figure 2, the data for the high-level problem includes electricity pricing, weather forecast, building use model, and disturbance estimates from low-level problem. Since most pricing structures include a peak demand charge based on the peak power usage during the month in addition to the time varying electricity costs, both components are included in the economic cost objective. The decision variables are building temperatures, storage utilization, and thermal loads for waterside system. Dynamics models for the building temperature and active TES are required. The aggregate waterside system model is represented by the coefficient of performance (COP) in the high-level objective. Constraints include comfort bounds on the air temperature and bounds on the cooling duty. The comfort zone constraints are implemented as soft constraints. Mathematically, the high-level MPC problem is formulated as

$$\begin{aligned} \min_{\mathbf{x}, \mathbf{u}, \dot{Q}_{\text{peak}}} \quad & \sum_{k=0}^{N-1} c_k \eta \dot{Q}_{\text{HVAC},k} \Delta + c_{\text{peak}} \eta \dot{Q}_{\text{peak}} \\ \text{s.t.} \quad & C_b \frac{dT_b}{dt} = -H_b(T_b - T_a) - \dot{Q}_{c,b} + \dot{Q}_{\text{other},b} \\ & \frac{ds}{dt} = -\sigma s + \dot{Q}_{\text{storage}} \\ & \dot{Q}_{\text{HVAC},k} = \sum_b \dot{Q}_{c,b,k} + \dot{Q}_{\text{storage},k} \\ & 0 \leq \dot{Q}_{\text{HVAC},k} \leq \dot{Q}_{\text{peak}} \\ & \dot{Q}_{\text{peak,past}} \leq \dot{Q}_{\text{peak}} \leq \dot{Q}_{\text{HVAC,max}} \\ & 0 \leq s_k \leq s_{\text{max}} \\ & T_{\text{min}} \leq T_b \leq T_{\text{max}} \end{aligned} \quad (4)$$

The states in the high-level MPC problem are the building temperatures and storage tank level. The inputs are the cooling duties to each building and the amount of storage charged or discharged. A rate of change penalty can be added to the inputs in the MPC problem. The disturbances are the ambient temperature and external loads placed on each building. A forecast of this disturbance vector is obtained and provided to the optimization problem at each MPC step. In the simulation study, the states are assumed to be measured. In real applications, state estimation is required to reconstruct the states from the measurements. To obtain building models from the zone models that are typically available or identified, model aggregation is employed. The average building temperature is defined by summing over all zones in that building using

$$T_b = \frac{\sum_i C_i T_i}{\sum_i C_i} \quad (5)$$

With this definition, the building parameters for the high-level airside model are obtained by summing parameters from the low-level airside models using

$$C_b = \sum_i C_i \quad H_b = \sum_i H_i \quad \dot{Q}_{\text{other},b} = \sum_i \dot{Q}_{\text{other},i} \quad (6)$$

Model aggregation also helps to reduce the amount of information that needs to be exchanged, thus, addressing the potential practical limitations with existing HVAC communication protocols.

### 3.3 Low-Level Airside Optimization Problem

As shown in Figure 2, the data for the low-level airside problem includes the cooling duty allocated from the high-level problem. Its objective is to minimize energy usage over the horizon by varying zone temperature setpoints. The zone air temperature and temperature setpoint to cooling duty model is required ( $\dot{Q}_{c,i} = f(T_i, T_{sp,i})$ ), which is discussed in the following section. The constraints include comfort bounds on the zone air temperatures and bounds on the cooling duty obtained from the high-level problem; they are all implemented as soft constraints with comfort zone violations penalized more than high-level cooling duty violations.

In each building, the low-level airside MPC problem is formulated as


$$\begin{aligned} \min_{\mathbf{x}, \mathbf{u}} \quad & Q_{\text{total},N} \\ \text{s.t.} \quad & C_i \frac{dT_i}{dt} = -H_i(T_i - T_a) - \sum_{j \neq i} \beta_{ij}(T_i - T_j) - \dot{Q}_{c,i} + \dot{Q}_{\text{other},i} \\ & \frac{dQ_{\text{total}}}{dt} = \sum_i \dot{Q}_{c,i} \\ & T_{\min} \leq T_i \leq T_{\max} \\ & \frac{Q_{\text{total},k+1} - Q_{\text{total},k}}{\Delta} \leq \dot{Q}_{\text{HighLevel},k} \\ & Q_{\text{total},k+1} - Q_{\text{total},k} \geq 0 \\ & \dot{Q}_{c,i} = f(T_i, T_{sp,i}) \end{aligned} \quad (7)$$

To obtain the cooling duty as a function of the temperature setpoint and temperature, two separate models are required:


- 1). A model of the regulatory zone temperature controller to determine its control action in the air handler unit (AHU) as a function of the zone temperature and temperature setpoint:  $v_{\text{air},i} = f_1(T_i, T_{sp,i})$
- 2). An energy balance relating the control action to the actual cooling duty:  $\dot{Q}_{c,i} = f_2(v_{\text{air},i})$

Assuming an ideal proportional-integral (PI) controller and a linear relationship between  $\dot{Q}_{c,i}$  and  $v_{\text{air},i}$ , the simplified linear cooling duty controller model is

$$\begin{aligned} \dot{Q}_{c,i} &= \dot{Q}_{ss,i} + K_{c,i} \left[ \varepsilon_i + \frac{1}{\tau_{1,i}} \int_0^t \varepsilon_i(t') dt' \right] \\ \varepsilon_i &= T_{sp,i} - T_i \end{aligned} \quad (8)$$



**Figure 3: Ambient Temperature and Power Cost.**



**Figure 4: High-Level Optimization Result.**

Saturation is represented as constraints on  $\dot{Q}_{c,i}$ . If a linear model is not sufficiently accurate to model the PI controller and heat transfer in air handler units, a nonlinear cooling duty model can be used instead. In this case, (7) becomes a nonlinear optimization problem. While it may seem atypical to model the regulatory controllers in the MPC optimization problem, the dynamics of zone PI controllers cannot be ignored as their responses are sluggish. In building temperature regulation, it may take zones up to an hour to reach a new temperature setpoint, which does not permit a time-scale separation in the decomposition. Since electricity prices are time-varying, ignoring these dynamics leads to suboptimal operation.

For the low-level airside problem, the states are the zone temperatures, the integrals of the zone tracking errors, and the total amount of cooling delivered. The inputs are the zone temperature setpoints. A rate of change penalty can be added to these inputs as well in the MPC problem to discourage rapid setpoint changes. The disturbances are the ambient temperature, external loads, and steady-state cooling rates.

### 3.4 Alternative Formulation: Temperature Tracking

An alternative strategy for the low-level airside problem is to track the average building temperature from the high-level problem rather than to minimize energy consumption. To track the temperature computed from the high-level problem, the high-level cooling duty constraint from (7) is removed, and the objective is replaced with

$$\sum_{k=1}^N \frac{1}{2} \|T_{b,k} - T_{\text{HighLevel},k}\|^2 + \mu Q_{\text{total},N} \quad (9)$$

in which the average building temperature is defined by Equation (5) and  $\mu$  is a small penalty placed on energy usage. This penalty on energy usage ensures uniqueness of the solution.

## 4. SIMULATION RESULTS

A simulation study was performed for a three-zone building with active TES. The weather and electricity pricing data used in the simulation are shown in Figure 3. The weather data were obtained from Atlanta, Georgia for a three-day period in the summer (July 25-27, 2012). The time-varying electricity price and peak demand charge data are representative of data from Johnson Controls. Figure 4 shows the results of the high-level optimization problem, and Figure 5 shows the results of the low-level optimization airside problem. The top graph in each figure shows the temperatures with the comfort zone denoted by dotted lines. The bottom graph in each figure shows the inputs computed from the corresponding optimizations. With passive TES available, the economically optimal strategy is to precool the building before the peak period when electricity prices are highest. The flatness of the cooling duty profile from the high-level problem is a result of the peak demand charge. The low-level airside problem determines the order of precooling based on the dynamics of the individual zones in the building to minimize energy usage. Since Zone 2 is not as well insulated as Zone 3, the low-level airside controller chooses to precool Zone 3 before the Zone 2.




Figure 5: Low-Level Airside Optimization Result.




Figure 6: Low-Level Waterside Optimization Result.

Figure 6 shows the results of the low-level waterside optimization problem. The top graph shows how the demand profile from the high-level problem is met using a combination of production from chillers and charging/discharging of the storage tank. The bottom graph shows the corresponding equipment operation schedule for the two chillers. The shaded region depicts when the equipment is turned on with the approximate loading fraction between 0 and 1 depicted by the solid line. With active TES available, chillers produce more chilled water during the night to charge the storage tank so that they can be turned off during the peak period when electricity is more expensive.

The concept can be extended to larger systems with more buildings and more zones. This decomposition strategy has also been applied to a campus with 25 buildings, each with 20 zones (for a total of 500 zones). The resulting optimization problems at each MPC step can be solved within a few minutes when distributed controllers are implemented for the low-level airside problem; however, the results are not shown due to space limitations.

## 5. AIRSIDE SYSTEM IDENTIFICATION

To implement this control system, an airside model is required. A few issues must be considered in the building problem that are not present in typical chemical process control applications. During testing, exciting inputs using large pseudo-random, binary signals (PRBS) can result in frequent saturation of the PID controllers in air handler units, hence small PRBS or integrated PRBS are used instead. Moreover, one of the primary disturbances affecting buildings, ambient temperature, can be easily measured and predicted. In order to use weather forecasts during online optimization, the effect of ambient temperature on zone temperature dynamics is identified by adding ambient temperature to the input vector before passing the data to the identification routine. The model between other disturbances, such as solar radiation, and the zone temperature can also be identified if measurements and predictions of these disturbances are available. Furthermore, temperatures must be kept in the comfort zone if the building has occupants during testing period. Sample identification results are presented in the next section that account for these issues.

### 5.1 Black-Box and Grey-Box Identification

In a linear black-box model, no knowledge about the internal structure of the system is assumed. While these techniques can be applied to any system using some knowledge of the system to add structure to the model can greatly increase the accuracy of the identified model, while decreasing data requirements. Models with this added structure are called grey-box models. To determine whether black-box and grey-box identification perform equally well for the building problem, both types of identification were performed on the same dataset using MATLAB's System Identification Toolbox. The free parameters in the grey-box approach include the thermal capacitances of the zones and heat transfer coefficients from Equation (2) as well as the PI controller parameters from Equation (8). The results are shown in Figure 7. Based on the figure, both methods fit the validation data well. However, the two methods do not scale equivalently for large-scale systems.




Figure 7: Identification Results.




Figure 8: Scaling Analysis.

## 5.2 Scaling Analysis

Figure 8 shows how the number of free parameters to be identified from data increases as the number of zones increases. The figure shows the worst case scenario for grey-box identification when all heat transfer coefficients between zones need to be determined. Typically, many of these coupling coefficients can be set to zero if zones are not adjacent, which leads to a more linear scaling for the grey-box case. As large-scale commercial building applications may have thousands of zones, identifying black-box models for these systems involves determining many parameters. More data is required when there are more free parameters to identify. Hence, grey-box identification requires less data than black-box identification. In practice, less data means a shorter testing period and faster implementation of the MPC controller, which makes grey-box identification more suitable for large-scale commercial applications.

## 6. CONCLUSIONS

An economically optimal and practically implementable method of temperature regulation and equipment operation for large-scale commercial building applications was presented. The control system architecture considers both airside and waterside optimization as well as storage, and the architecture can be easily extended to handle more complex models. Since the time scale of regulatory controllers is significant, their dynamics are modeled. Grey-box system identification is preferred over black-box identification for large systems to decrease the testing period. We show the effectiveness of the proposed method through the simulation of a three-zone building.

## 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
$\dot{Q}_{total}$	total amount of cooling delivered
$\dot{Q}_{HVAC}$	cooling rate from HVAC system
$\dot{Q}_{peak}$	peak HVAC system cooling rate
$s$	amount of cooling potential in storage tank
$\dot{Q}_{storage}$	cooling rate delivered to storage tank
$\dot{Q}_{HighLevel}$	cooling rate from high-level problem
$T_{HighLevel}$	building temperature from high-level
$\varepsilon_i$	tracking error
$v_{air,i}$	air flow rate in AHU

**Parameters**

$\Delta$	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
$\beta_{ij}$	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
$\eta$	inverse of the aggregate COP
$\dot{Q}_{HVAC,max}$	max cooling capacity of HVAC system
$\dot{Q}_{peak,past}$	peak cooling rate previously achieved
$\sigma$	decay constant for storage tank
$s_{max}$	max cooling capacity of storage tank
$T_{min}$	lower bound of comfort region
$T_{max}$	upper bound of comfort region
$\mu$	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.
- Department of Energy. (2012, March). *2011 Buildings Energy Data Book*. <http://buildingsdatabook.eren.doe.gov/ChapterIntro3.aspx>.
- 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.
- Oldewurtel, F., Parisio, A., Jones, C. N., Morari, M., Gyalistras, D., Gwerder, M., ... Wirth, K. (2010). Energy efficient building climate control using stochastic model predictive control and weather predictions. In *American Control Conference* (pp. 5100–5105). Baltimore, MD, USA.
- 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.
- Rawlings, J. B., & Stewart, B. T. (2008). Coordinating multiple optimization-based controllers: New opportunities and challenges. *Journal of Process Control*, 18, 839–845.
- 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*.