July 2018

Formulation and Application of an Economic Model Predictive Control Scheme for Thermostats

Matthew J Ellis
Johnson Controls, United States of America, matthew.j.ellis@jci.com

Anas Alanqar
Johnson Controls, United States of America, anas.alanqar@jci.com

Formulation and Application of an Economic Model Predictive Control Scheme for Connected Thermostats

Matthew J. ELLIS[1*], Anas ALANQAR[1]

[1] Johnson Controls Inc.; Technology and Advanced Development, Building Efficiency, Milwaukee, WI, United States of America
matthew.j.ellis@jci.com*, anas.alanqar@jci.com

* Corresponding Author

ABSTRACT

An economic model predictive control (MPC) scheme, which is an MPC scheme equipped with an economically-oriented objective function, is developed for a connected thermostat application. The economic MPC selects a setpoint from an occupant defined comfort range to minimize the HVAC power cost. Specifically, under a time-varying electric rate structure (e.g., time-of-use or real-time pricing), the economic MPC leverages the building mass as thermal energy storage to shift HVAC power consumption from high to low cost periods. The resulting economic MPC system includes a parameterized building thermal zone model, a parameter estimation procedure to identify the model parameters for a specific zone application, a state/disturbance estimator, a heat load disturbance forecaster, and an underlying optimal control problem formulation. Each of these features are tailored for broad application as a supervisory controller manipulating the zone temperature setpoint for a zone controlled by a thermostat. Given a lack of measurements available to estimate/measure HVAC power or load in a typical thermostat, the HVAC load is approximated via a filtered version of the thermostat equipment stage commands, which provides a normalized time-averaged approximation of the HVAC load. Simulation results are presented to demonstrate the effectiveness of the strategy.

1. INTRODUCTION

Within the last ten years, growing pressure to reduce energy consumption within buildings has led to an increased focus on the development and deployment of advanced control strategy capable of reducing energy consumption. One such control strategy, model predictive control (MPC), has gained significant attention as it computes the control inputs for a given system by repeating solving an optimal control problem on-line (e.g., (Mendoza-Serrano & Chmielewski, 2012; Ma, Qin, Salsbury, & Xu, 2012; Afram & Janabi-Sharifi, 2014; Patel, Risbeck, Rawlings, Wenzel, & Turney, 2016)). The problem formulation accounts for the system operating objective and constraints. To reduce energy consumption and/or energy costs, typically the MPC problem is formulated with an economically oriented cost function. Within the MPC community, this type of MPC is commonly referred to as economic MPC (e.g., (Rawlings, Angeli, & Bates, 2012; M. Ellis, Durand, & Christofides, 2014)).

Several application studies of MPC applied to buildings have been reported in the literature (e.g., (Mendoza-Serrano & Chmielewski, 2012; Ma *et al.*, 2012; Afram & Janabi-Sharifi, 2014; Patel *et al.*, 2016)). Additionally, industrial application of economic MPC schemes to HVAC systems is now a reality (Wenzel, Turney, & Drees, 2016). The applications reported have demonstrated clear realizable benefits of the application of MPC schemes to HVAC systems and buildings. However, one theme that appears in some of the literature is the lamentation on the difficulty to apply MPC broadly to buildings (e.g., (Henze, 2013)). Indeed, the design of a complete MPC system is a challenging problem because the MPC system design needs to include a broadly applicable system identification methodology. Additionally, in many building applications, the desired measurements for key variables are not available (e.g., a heat disturbance load and power consumption measurements are usually not available for residential zones controlled by a standard thermostat).

In this work, an economic MPC system that includes a system identification procedure, a state estimator, a heat disturbance load predictor, and a control problem formulation is developed for application as a supervisory control of a connected thermostat. Each of these features are tailored for broad application as a supervisory controller manipulating the zone temperature setpoint for a zone controlled by a thermostat. The economically oriented MPC seeks to minimize the utility bill by manipulating the setpoint to leverage the building mass as thermal energy storage while maintaining the zone temperature setpoint within a comfortable range. Given that the lack of a power or HVAC load measurement in a typical thermostat, the HVAC load is approximated via a filtered version of the thermostat stage commands, which provides a normalized time-average version of the HVAC load. The overall advantage of the approach is that it may be applied to any connected thermostat and potentially, achieve cost-reduction benefit. Thus, although one would expect improved closed-loop performance through additional sensors, the approach developed here does not require additional sensors. Simulation results are presented to demonstrate the effectiveness of the strategy.

2. ECONOMIC MODEL PREDICTIVE CONTROL FORMULATION

An overview of the economic MPC formulation applied to connected thermostats to minimize utility costs is provided in this section.

2.1 Notation

The set of integers is denoted by $\mathbb{I}$, the set of positive integers is denoted by $\mathbb{I}_{\geq 0}$, and the set of integers contained in the interval $[a, b]$ is denoted by $\mathbb{I}_{a:b}$. For a time-dependent vector $x(k) \in \mathbb{R}^n$ ($k \in \mathbb{I}_{\geq 0}$), $\hat{x}(i|k) \in \mathbb{R}^n$ denotes the estimated value of x at time step $i \in \mathbb{I}_{\geq 0}$ given the measurement at time step $k \in \mathbb{I}_{\geq 0}$ where $i \geq k$ and $\tilde{x}(i|k) \in \mathbb{R}^n$ denotes the (open-loop) predicted value of x at time step $i \in \mathbb{I}_{\geq 0}$ with the prediction initialized at time $k \in \mathbb{I}_{\geq 0}$ ($i \geq k$). For notational simplicity, the notation of the predicted value of x at time i starting from time k is abbreviated to $\tilde{x}(i)$. Boldface letters are used to represent a sequence with cardinality $N \in \mathbb{I}_{\geq 0}$ (i.e., $\mathbf{x} := \{x(0), \ldots, x(N-1)\}$).

2.2 Zone Thermal Model and Equipment Model

A two-resistance, two-capacitance (2R2C) control-oriented zone thermal model is considered to describe the thermal dynamics of a zone or space. The 2R2C model is given by:

$$\begin{aligned} C_{ia}\dot{T}_{ia} &= \frac{1}{R_{oi}}(T_{oa} - T_{ia}) + \frac{1}{R_{mi}}(T_m - T_{ia}) - K_{clg}\dot{\bar{Q}}_{clg} + K_{htg}\dot{\bar{Q}}_{htg} + \dot{Q}_{other} \\ C_m\dot{T}_m &= \frac{1}{R_{mi}}(T_{ia} - T_m) \end{aligned} \tag{1}$$

where T_{ia} is the indoor air temperature, T_m is the mass temperature, C_{ia} is the thermal capacitance of the indoor air, C_m is the thermal capacitance of the zone mass (e.g., the walls and furniture contained within the zone), R_{oi} is the thermal resistance between the indoor air and the outdoor air, R_{mi} is the thermal resistance between the indoor air and the zone mass, $\dot{\bar{Q}}_{clg} \geq 0$ and $\dot{\bar{Q}}_{htg} \geq 0$ is the time-averaged normalized sensible HVAC cooling and heating load, $K_{clg} > 0$ and $K_{htg} > 0$ are the normalization factors of the HVAC sensible load, $\dot{Q}_{other}$ is the sensible disturbance heat gained from, for example, operating electrical equipment within the zone, occupancy load, and solar radiation. The normalized sensible load and normalization factors are described further below.

Many standard thermostats that control up to two heating and cooling stages measure the indoor air temperature and humidity to decide the command (on or off) to send to the HVAC equipment such that the indoor air temperature tracks within the deadband of the temperature setpoint. Figure 1 gives an example of indoor air temperature of a zone and the first cooling stage on and off trajectories. Typically, measurements of air flows, power consumption, and other variables that could be used to estimate sensible loads are not available in this setting.

To effectively optimize energy consumption or cost, some knowledge (model or measurement) is required. The on/off trajectory of the HVAC equipment stages may be used to approximate the sensible HVAC load. Thermostats that control staged equipment dictate bang-bang type control actions over a time-scale of minutes. The time-scale of relevance to energy cost minimization is on the order of minutes-hours. To bridge these time-scales, a time-averaged version of the on-off trajectory may be used to provide an approximation of the HVAC load needed to force the temperature to and maintain it at the setpoint. Figure 2 shows an illustration of the time-averaged version of both the indoor temperature and the HVAC load trajectories. The time-averaging of the indoor temperature also removes the high frequency dither on the indoor temperature trajectory caused by the equipment being turned on and off, which may not be relevant over the time-scale of optimization.

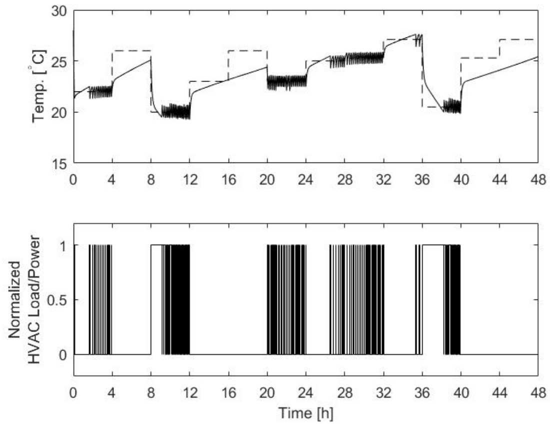

Figure 1: The indoor temperature and on/off of HVAC equipment trajectories over two days.

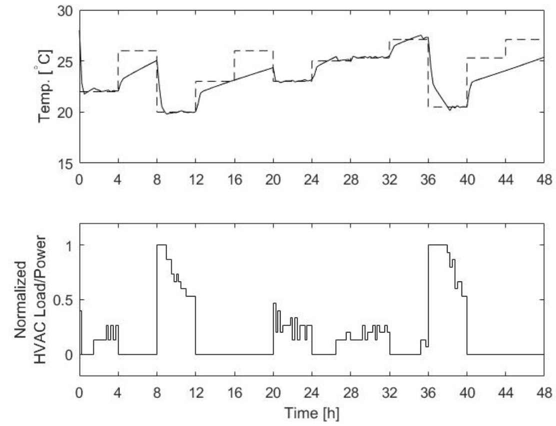

Figure 2: The time-averaged (over fifteen minute intervals) temperature and on/off trajectories.

Within the context of this work, a model predictive control scheme is developed to select a temperature setpoint within a user-specified comfort range to minimize the HVAC operating cost. Thus, a relationship between the temperature setpoint and the HVAC sensible load is needed. A proportional-integral (PI) controller is used to model this relationship. The PI controller model is given by:

$$\dot{\bar{Q}}_j = -K_{c,j}T_{ia} + \frac{K_{c,j}}{\tau_{I,j}}I_j + K_{c,j}T_{sp} \tag{2}$$

for $j \in \{htg, clg\}$ where $K_{c,j}$ and $\tau_{I,j}$ are the two parameters of the PI controller and T_{sp} is the temperature setpoint (note that $K_{c,clg} < 0$ and $K_{c,htg} > 0$ owing to the sign convention used). Two different PI control models are used to model the relationship between a temperature setpoint and the HVAC load because in general, the dynamics may be different between operating in heating and operating in cooling (e.g., different equipment may be used and/or the discharge air flow rates may be different in heating and cooling mode). All parameters (C_{ia}, C_m, R_{mi}, R_{oi}, $K_{c,j}$, $\tau_{I,j}$, and K_j) of the model of (1) are identified using a parameter identification scheme. The resulting parameter identification problem has six degrees of freedom for each operating mode. Due to space limitations, more details on the parameter identification scheme are not provided, but the interested reader is referred to (M. J. Ellis, Wenzel, & Turney, 2016; Alanqar, Ellis, Tapiero, & Wenzel, 2018) for more details.

For completeness, the resulting model for each mode are provided. In cooling mode, the model is given by:

$$C_{ia}\dot{T}_{ia} = \frac{1}{R_{oi}}(T_{oa} - T_{ia}) + \frac{1}{R_{mi}}(T_m - T_{ia}) - K_{clg}K_{c,clg}(T_{sp} - T_{ia}) - \frac{K_{clg}K_{c,clg}}{\tau_{I,clg}}I_{clg} + \dot{Q}_{other} \tag{3a}$$

$$C_m\dot{T}_m = \frac{1}{R_{mi}}(T_{ia} - T_m) \tag{3b}$$

$$\dot{I}_{clg} = T_{sp} - T_{ia} \tag{3c}$$

$$\dot{I}_{htg} = 0 \tag{3d}$$

In heating mode, the model is given by:

$$C_{ia}\dot{T}_{ia} = \frac{1}{R_{oi}}(T_{oa} - T_{ia}) + \frac{1}{R_{mi}}(T_m - T_{ia}) + K_{htg}K_{c,htg}(T_{sp} - T_{ia}) + \frac{K_{htg}K_{c,htg}}{\tau_{I,htg}}I_{htg} + \dot{Q}_{other} \tag{4a}$$

$$C_m\dot{T}_m = \frac{1}{R_{mi}}(T_{ia} - T_m) \tag{4b}$$

$$\dot{I}_{clg} = 0 \tag{4c}$$

$$\dot{I}_{htg} = T_{sp} - T_{ia} \tag{4d}$$

In off/standby mode, the model is given by:

$$C_{ia}\dot{T}_{ia} = \frac{1}{R_{oi}}(T_{oa} - T_{ia}) + \frac{1}{R_{mi}}(T_m - T_{ia}) + \dot{Q}_{other} \tag{5a}$$

$$C_m\dot{T}_m = \frac{1}{R_{mi}}(T_{ia} - T_m) \tag{5b}$$

$$\dot{I}_{clg} = 0 \tag{5c}$$

$$\dot{I}_{htg} = 0 \tag{5d}$$

Converting the continuous-time sets of ordinary differential equations of (3)-(5) to discrete-time, gives the following discrete-time switched system:

$$\underbrace{\begin{bmatrix} T_{ia}(k+1) \\ T_m(k+1) \\ I_{clg}(k+1) \\ I_{htg}(k+1) \end{bmatrix}}_{=:x(k+1)} = A_{\sigma(k)} \underbrace{\begin{bmatrix} T_{ia}(k) \\ T_m(k) \\ I_{clg}(k) \\ I_{htg}(k) \end{bmatrix}}_{=:x(k)} + B_{\sigma(k)} \underbrace{\begin{bmatrix} T_{sp}(k) \\ T_{oa}(k) \\ \dot{Q}_{other}(k) \end{bmatrix}}_{=:u(k)} \tag{6}$$

with measured outputs:

$$\underbrace{\begin{bmatrix} T_{ia}(k) \\ \dot{\hat{Q}}_{clg}(k) \\ \dot{\hat{Q}}_{htg}(k) \end{bmatrix}}_{=:y(k)} = C_{\sigma(k)}x(k) + D_{\sigma(k)}u(k) \tag{7}$$

where $k \in \mathbb{I}_{\geq 0}$ is the discrete time step and $\sigma(k) \in \{htg, clg, off\}$ denotes the mode at time step k. The discrete-time matrices $A_{\sigma(k)}$, $B_{\sigma(k)}$, $C_{\sigma(k)}$, and $D_{\sigma(k)}$ depends on the identified parameters for the corresponding mode at time k.

The current and forecasted outdoor air temperature may be available through a weather service. In this case, the outdoor air temperature is considered to be a measured/predicted disturbance. The heat load disturbance is an unmeasured disturbance. When a weather service is unavailable, the heat transfer between indoor and outdoor may be lumped into the heat load disturbance, but the predictions from the resulting model is expected to be worse than explicitly modeling the forcing due to the outdoor air temperature.

The power consumed by the HVAC equipment depends on the sensible load and may depend on outdoor air conditions. Owing to a lack of power measurement available, a constant efficiency model is assumed to model the relationship between power consumption and sensible HVAC load in the sense that:

$$P_i(k) = c_{eff,i}\dot{Q}_j(k) = c_{eff,i}K_j\dot{\hat{Q}}_j(k) \tag{8}$$

where $P_i(k)$ denotes the power consumption of utility i. The proportionality constant(s) $c_{eff,i}$ are not identified owing to a lack of measurements available to estimate or identify it. The efficiencies are not needed to optimize the cost (refer to Section 2.5). Future work could focus on developing representative equipment models that account for the effect of outdoor air conditions on the HVAC equipment efficiency. Given that the economic MPC works to optimize the energy cost by leveraging the building mass as thermal energy storage to shift part of the HVAC load from high to low cost periods, a time-varying electric rate structure (e.g., time-of-use or real-time pricing) is expected. For furnaces, cost savings may be achieved through shifting HVAC loads to optimize the fan power consumption.

2.3 State/Disturbance Estimation

Owing to the fact that the disturbance heat load $\dot{Q}_{other}(k)$ is not measured and to account for other sources of unmeasured disturbances, the model of (6)-(7) is augmented with an integrating disturbance model (see (Pannocchia & Rawlings, 2003), for example, for more details on the use of an integrating disturbance model within MPC). The update equations of the state/disturbance estimator are given by:

$$\begin{aligned}\begin{bmatrix}\hat{x}(k+1|k)\\ \hat{d}(k+1|k)\end{bmatrix} &= \begin{bmatrix}A_{\sigma(k)} & B_d\\ 0 & 1\end{bmatrix}\begin{bmatrix}\hat{x}(k|k-1)\\ \hat{d}(k|k-1)\end{bmatrix} + \begin{bmatrix}B_{\sigma(k)}\\ 0\end{bmatrix}u(k) + K_{\sigma(k)}(y(k)-\hat{y}(k|k-1))\\ \hat{y}(k|k-1) &= \begin{bmatrix}C_{\sigma(k)} & 0\end{bmatrix}\begin{bmatrix}\hat{x}(k|k-1)\\ \hat{d}(k|k-1)\end{bmatrix} + D_{\sigma(k)}u(k)\end{aligned} \tag{9}$$

where $\hat{d}(k|k-1)$ is the integrating disturbance estimated for time step k provided measurements up to $k-1$ and $K_{\sigma(k)}$ is the estimator gain (i.e., the Kalman gain estimated through the identification algorithm of (M. J. Ellis *et al.*, 2016; Tapiero Bernal & Ellis, 2017)). The measured indoor air temperature and the equipment commands are filtered using a Savitzky-Golay filter to perform the filtering described in Section 2.2.

2.4 Disturbance Load Predictor

The unmeasured disturbance heat load $\dot{Q}_{other}$ accounts for the heat load transferred to the space through lighting, other electrical equipment, occupants, and solar radiation. The heat load disturbance is time-varying and will have a periodic trend, which is deterministic in nature. To this end, the integrating disturbance state of (9) will account for $\dot{Q}_{other}$. When the model (6)-(7) describes the space dynamics well such that the most significant source of plant-model mismatch is the unmeasured disturbance, the low frequency component of the disturbance state may be used to train a predictor to predict the deterministic component of the heat load disturbance.

Owing to space restrictions, the method is only summarized here. Specifically, a sequence of historical values of $\hat{d}$ are collected over time, which is denoted by $\hat{\mathbf{d}}_{\mathbf{hist}} := \{\hat{d}(k|k-1)\}_{k=0}^{N_{hist}}$ where $N_{hist} \in \mathbb{I}_{>0}$ is the number of sampled collected and low-pass filtered to remove high frequency noise. The amount of data to be collected is a couple of weeks or more such that a day-type pattern recognition algorithm may adequately capture the different day type patterns. Given that the head load disturbance is correlated to the time-of-day and day-type (e.g., weekdays, weekends, and holidays), a function that describes the variation of $\hat{d}$ with the time may be fit using the data (e.g., using the algorithm of (ElBsat & Wenzel, 2016)). The resulting function is given by:

$$\tilde{d}(k) = f_{pred}(k\Delta + t_0, \ldots; \hat{\mathbf{d}}_{\mathbf{hist}}) \tag{10}$$

where $\Delta > 0$ is the sample period of the MPC, t_0 is some initial time (i.e., $k\Delta + t_0$ is the current time), f_{pred} is the fitted mapping using the historical data with time, outdoor air conditions, and potentially, cloud cover as inputs, and $\tilde{d}(k)$ is the predicted heat disturbance for the kth time step. An autoregressive (AR) model may be added to the predictor of (10) to account for discrepancies between the current estimated disturbance and the predicted values computed from (10) for the current time.

Using the predictor of (10) and adding the AR model, a forecast of the unmeasured disturbance may be utilized in the MPC. This approach is different than that typically used in MPC. Specifically, the typical approach for disturbance estimation and prediction is to estimate the current disturbance and for the disturbance forecast over the horizon, use the current disturbance estimate such that the resulting disturbance forecast is constant over the horizon (Pannocchia & Rawlings, 2003).

2.5 Economic MPC Problem Formulation

The purpose of the economic MPC applied to the thermostat is to manipulate the temperature setpoint within a user-defined comfort range to minimize the HVAC electric power consumption costs. Even though the MPC does not directly control the HVAC equipment, the MPC still should account for the predicted indoor air temperature evolution with respect to the comfort range. Thus, soft constraints are imposed within the MPC problem on the indoor air temperature in an attempt to help ensure the indoor air temperate is maintained within the comfort range. This means that a penalty is included in the cost function of the MPC problem when the predicted zone temperature deviates from the comfort bounds. Additionally, a change penalty on the temperature setpoint may be added to the cost function to prevent the MPC from making large and/or frequent changes in the temperature setpoint.

The input data to the economic MPC problem includes:

- The current state estimate from (9).

- The forecasted weather over the horizon.
- The heat load disturbance forecast over the horizon from (10).
- The electric rate over the MPC prediction horizon.
- The comfort range over the horizon. If the comfort range is defined on the basis of the occupancy, an occupancy schedule or forecasted occupancy is needed.
- The indoor air comfort violation penalty.
- The setpoint change penalty.
- The current and forecasted mode over the horizon. If the forecasted mode is unavailable, the current mode is assumed over the horizon.
- Previous temperature setpoint.

The economic MPC problem is given by:

$$\min_{\mathbf{T_{sp}},\epsilon_\mathbf{T},\delta\mathbf{T_{sp}}} \sum_{k=0}^{N-1} r_{elec}(k)\left(\dot{\tilde{Q}}_{\sigma(k)} + p_{\epsilon_T}\epsilon_T(k) + p_{\delta T}\delta T_{sp}(k)\right) \tag{11a}$$

$$\text{s.t.} \quad \tilde{x}(k+1) = A_{\sigma(k)}\tilde{x}(k) + B_{\sigma(k)}\begin{bmatrix}T_{sp}(k)\\ \tilde{w}(k)\end{bmatrix} \tag{11b}$$

$$\tilde{y}(k) = C_{\sigma(k)}\tilde{x}(k) + D_{\sigma(k)}\begin{bmatrix}T_{sp}(k)\\ \tilde{w}(k)\end{bmatrix} \tag{11c}$$

$$T_{\sigma(k),min}(k) - \epsilon_T(k) \le \tilde{T}_{ia}(k) \le T_{\sigma(k),max}(k) + \epsilon_T(k) \tag{11d}$$

$$T_{sp,\sigma(k),min}(k) \le T_{sp}(k) \le T_{sp,\sigma(k),max}(k) \tag{11e}$$

$$\tilde{T}_{sp}(k+1) - \tilde{T}_{sp}(k) \le \delta T_{sp}(k) \tag{11f}$$

$$\tilde{T}_{sp}(k) - \tilde{T}_{sp}(k+1) \le \delta T_{sp}(k) \tag{11g}$$

$$\epsilon_T(k) \ge 0,\ \delta T_{sp}(k) \ge 0 \tag{11h}$$

$$\forall\, k \in \mathbb{I}_{0:N-1}$$

$$\tilde{x}(0) = \hat{x} \tag{11i}$$

where N is the number of fifteen minute sampling periods in the prediction horizon, which was selected to be two days such that the MPC can effectively shift part of the HVAC load from high to low cost periods. The predicted trajectories $\tilde{T}_{ia}(k)$, $\tilde{x}(k)$, $\tilde{y}(k)$, $\dot{\tilde{Q}}_{\sigma(k)}(k)$ are the predicted zone (indoor) temperature, state, output, and normalized HVAC load (heating or cooling depending on the current mode), respectively, at time step k. The notation $\hat{x}$ denotes the current state estimate from the estimator of (9). The setpoint at time step k, which is the decision variable of the optimization problem is denoted by $T_{sp}(k)$, and $T_{\sigma(k),min}(k)$, $T_{sp,\sigma(k),min}(k)$, $T_{\sigma(k),max}(k)$, $T_{sp,\sigma(k),max}(k)$ are the minimum and the maximum indoor temperature and setpoint, respectively, defined by the comfort range at time step k. It is important to note that under cooling mode (and vice versa for heating mode), no actuation is available to increase the indoor air temperature, and therefore, one could consider not including a penalty on deviations from the lower range of the comfort range. Moreover, when the temperature setpoint is at its minimum and maximum, it is expected that there will be some violations of the indoor air temperature beyond the setpoint comfort range given that the thermostat controls the indoor air temperature within a deadband of the temperature setpoint. For both situations, the comfort ranges of the setpoint and indoor air temperature are different, and hence, why different comfort ranges are defined for both the indoor air and setpoint temperatures. The variable $\epsilon_T(k)$ is the predicted violation of the indoor air temperature from the comfort range at time step k, $\delta T_{sp}(k)$ is the change in temperature setpoint from time step $k-1$ to time step k, $\tilde{w}^T(k) := [\tilde{T}_{oa}(k), \tilde{d}(k)]$ are the forecasted disturbances including the outdoor air temperature and disturbance heat load at time step k, and $r_{elec}(k)$ is the electric rate at time step k. If the electric rate does not vary with time, the resulting problem becomes an energy minimization problem. The penalty weight $p_{\epsilon_T} \ge 0$ is the soft constraint penalty imposed on predicted violations of the zone temperature above the maximum or minimum comfort bound, $p_{\delta T} \ge 0$ is the change penalty imposed on the change in temperature setpoint

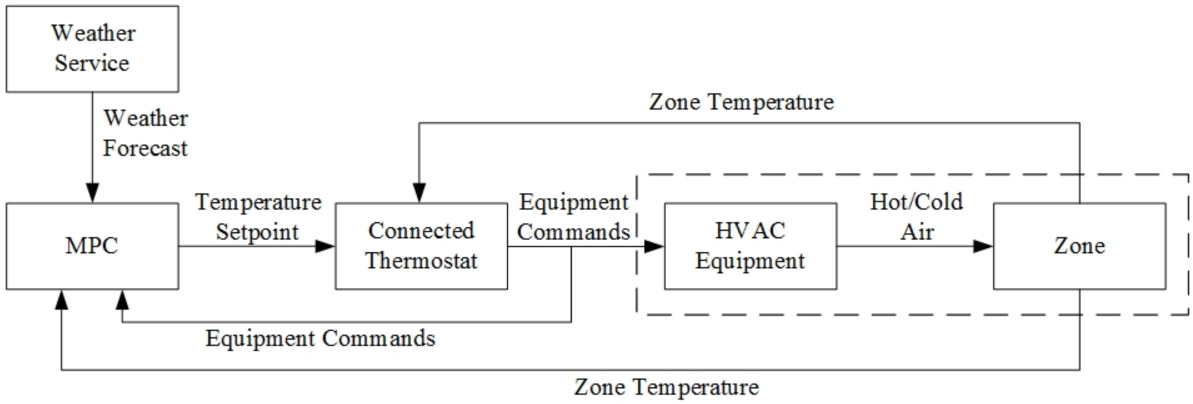

Figure 3: A block diagram depicting the closed-loop system with the economic MPC system commanding the temperature setpoint and a connected thermostat dictating the HVAC equipment commands to force the zone temperature to track the desired temperature setpoint.

The components of the problem of (11) are: a cost function of (11a) that accounts for the HVAC energy consumption costs plus the penalties on deviations of the indoor air temperature from the comfort bound and the change penalty on the temperature setpoint. The efficiency considered in (8) is omitted in the cost function because it is assumed that it is unknown. Nonetheless, omitting the efficiency does not change the optimal solution of the problem as it is only a scaling factor (assuming that the penalties are also scaled by the efficiency factor). The zone thermal dynamic model of (11b)-(11c) is used to predict the indoor air temperature and HVAC load over time. The constraints of (11d) are used for the soft constraints on the indoor temperature violations from the comfort range. The constraints of (11e) limits the temperature setpoint be from within the occupant defined comfort range. A set of constraints ((11f)-(11g)) ensures that $\delta T_{sp}(k) = |T_{sp}(k+1) - T_{sp}(k)|$. Finally, the constraint of (11i) is the initialization of the state trajectory at the current state estimate.

3. ECONOMIC MODEL PREDICTIVE CONTROL FOR CONNECTED THERMOSTATS

The operation of the economic MPC system for the connected thermostat application is described in this section. Figure 3 provides a block diagram of the closed-loop system with the economic MPC acting as a supervisory controller. The economic MPC system includes two modes of operation: (1) system excitation for parameter identification, which is referred to as parameter identification mode and (2) operational mode. Under the parameter identification mode, the zone temperature setpoint is manipulated to excite the zone temperature. Under operational mode, the MPC problem is solved to determine the optimal temperature setpoint trajectory over the prediction horizon. The computed setpoint is subsequently sent to the connected thermostat device to be implemented over the sampling period of the MPC.

3.1 Parameter Identification Mode

To identify the parameters of the zone thermal model, the zone is excited by changing the temperature setpoint of the zone, which induces a dynamic response in the zone temperature. This procedure of inducing a dynamic response is referred to as the parameter identification experiment. Over this period, a pseudo-random binary signal (PRBS) is generated that is subsequently used to produce a signal that takes values in the time-dependent set $\{T_{sp,1}(k), T_{sp,2}(k)\}$ where $T_{sp,1}(k)$ and $T_{sp,2}(k)$ is the maximum and minimum allowable temperature setpoint at some time k ($T_{sp,1} < T_{sp,2}$). The choice of the minimum and maximum temperature setpoints are made by the zone occupants.

As described in Section 2.5, the minimum and maximum temperature setpoints may be time-varying to account for the different occupancy states. For example, the thermostat may have three occupancy states: home, away, and sleep. The zone occupant may have different temperature comfort ranges for each of these states. Over the course of the parameter identification experiment, zone conditions and device commands are collected. Depending on how the minimum and maximum temperature setpoints are determined, feedback from the thermostat should also be received such that the temperature setpoint excitation signal is adjusted to account for any change in the occupancy state, for example. The collected data is subsequently filtered using the filtering strategy described in Section 2.2 and the resulting filtered data is used to compute the zone thermal model parameters with the parameter identification algorithm (M. J. Ellis *et al*., 2016).

The parameter identification mode is expected to only be executed during the commissioning of the MPC for both heating and cooling modes. Future work is required to understand the need to re-execute the parameter identification at some point in the future as well as how this decision should be made. Ideally, this choice could be made in a completely automated fashion with little or no user intervention.

3.2 Operational Mode

After the heating or cooling model of (6)-(7) is identified for the zone, the MPC may operate in its main operational mode, which is responsible for computing the zone temperature setpoint within the comfort bounds to minimize the energy costs of operating the HVAC equipment. To determine the optimal setpoint, the thermal model is used to predict the zone temperature and HVAC fuel consumption over a prediction horizon given a forecast of the weather and heat disturbance load on the zone. The zone temperature is subject to the comfort constraints. The predicted trajectories of the zone temperature, setpoint, and the HVAC fuel consumption could be sent to the thermostat to be displayed to the occupants allowing the occupants to better understand the setpoint decision made by the MPC. The type of user interaction allows for the user to gain trust and better understand the MPC decision, which has been pointed out in the literature as an important component in MPC for HVAC systems (e.g., (Henze, 2013; Wenzel *et al*., 2016)).

The setpoint computed for the first sample period of the prediction horizon is sent to the thermostat control application to be implemented over the sampling period of the MPC. If the thermostat has a heating and a cooling setpoint the heating or cooling setpoint is set to its minimum or maximum, respectively, if the active operation mode is cooling or heating, respectively (i.e., the non-active mode setpoint is not a decision in the MPC problem). The thermostat uses measurements of current zone conditions to decide and send commands to turn on/off the HVAC equipment such that the zone temperature is forced to and then maintained at the temperature setpoint. At the next time step, the MPC receives updated measurements and on/off history information, resolves the optimal control problem of (11), and sends the updated setpoint trajectory back to the thermostat.

If the MPC is not implemented in the local device (e.g., is implemented in the cloud), communication disruptions between the thermostat and the cloud may be addressed by implementing the setpoint trajectory computed the last time communication between the device and the cloud was established (i.e., applying the setpoints computed for the subsequent sampling periods of the prediction horizon in an open-loop fashion until connectivity to the cloud is restored). Nevertheless, it is not desirable to maintain this type of open-loop operation for a prolonged period of time given that the incorporation of feedback measurements is instrumental for the robustness of the MPC to plant-model mismatch and unmeasured disturbances. Under periods of prolonged communication disruptions between the cloud and the thermostat, the regulatory control application within the thermostat may fall-back to using a standard scheduled setpoint scheduling that is equipped on the local thermostat application.

The operational mode algorithm is summarized by the following steps:

1. At time step k, the zone temperature and equipment stage commands are received over the last 15 minutes and filtered according to the filtering approach of Section 2.2.
2. The current and forecasted outdoor air temperature is received from the weather service.
3. The current state and disturbance estimate is computed using (9).
4. The forecasted disturbance heat load is computed from (10).
5. The optimal control problem of (11) is solved to compute the setpoint trajectory defined over the prediction horizon.
6. The setpoint for the first time step of the prediction horizon is sent to the thermostat to be implemented from k to $k+1$.
7. $k \leftarrow k+1$. Go to Step 1.

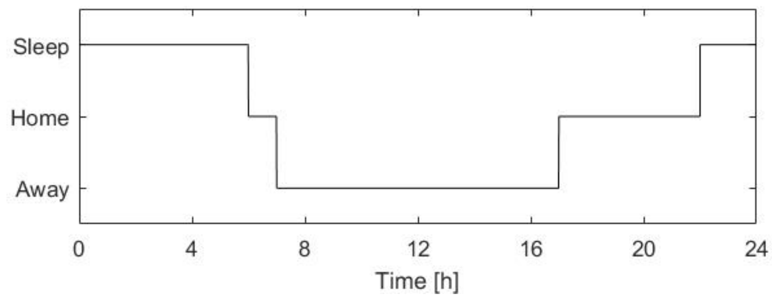

Figure 4: Occupancy schedule over the weekdays.

4. CLOSED-LOOP SIMULATION RESULTS

To demonstrate and test the MPC, several closed-loop simulations were conducted. To model the thermal dynamics of the zone and the thermostat temperature control, a detailed simulator was used along with a proprietary thermostat control application. The detailed simulator includes HVAC equipment and zone models, which consist of detailed first-principals based models. The zone model includes the zone temperature dynamics and is driven by realistic disturbances including heat load to the space due to the occupants, operating electrical equipment and lighting, and solar radiation.

The zone is first operated in parameter identification mode over a two day period. Under this operating mode, a PRBS signal is used to generate a temperature setpoint trajectory to excite the temperature dynamics of the zone. The indoor air temperature, equipment commands, and outdoor air temperature are collected during over the operating period. The data is filtered using the filtering methodology described in Section 2.2 and the resulting filtered data is provided to the parameter identification algorithm of (Tapiero Bernal & Ellis, 2017) to identify model parameters for the zone. After the model parameters are identified, the MPC switches to operational mode.

One example case study under operational mode is presented here, which considers the first week of August for a residential building in Phoenix, Arizona controlled by a thermostat with a residential air conditioning unit (1 cooling stage). The utility provides its residential customers a time-of-use electricity rate structure as an option. The rate structure includes an on-peak rate ($0.4652/kWh) during peak periods of weekdays (3:00PM-6:00PM), partial peak rate ($0.2448/kWh) during partial peak periods of weekdays (12:00PM-3:00PM and 6:00PM-7:00PM) and an off peak rate ($0.0552/kWh) during off peak periods of weekdays (7:00PM-12:00PM) and weekends. The sample period and prediction horizon of the MPC is set to fifteen minutes and two days, respectively.

To assess the performance of the MPC relative to a baseline case, three cases were considered. For the MPC case, the home, sleep, and away cooling comfort bound was defined as $74°\text{F} \pm 1°\text{F}$, $72°\text{F} \pm 1°\text{F}$, and $76°\text{F} \pm 3°\text{F}$. Additionally, two baseline cases were considered: (1) scheduled occupancy setpoint with a home, sleep, and away setpoint of 74°F, 72°F, and 76°F and (2) a home, sleep, and away setpoint of 75°F, 73.5°F, and 79°F. The two baseline cases represent the median and the maximum temperature of the comfort range, respectively. The weekday occupancy schedules used are shown in Figure 4. Over the week, the HVAC operating cost under the three cases was $46.84, $62.69 (25.3 percent savings with MPC), and $58.10 (19.4 percent savings with MPC), respectively. Figures 5-6 shows operation over one day and how the MPC is able to manipulate the temperature setpoint to achieve the cost by precooling.

5. CONCLUSIONS

A economic MPC system was developed for application as a supervisory controller to a connected thermostat application to minimize the HVAC power costs. The system includes a parameterized thermal model, parameter identification scheme, state estimator, disturbance forecaster, and an optimal control problem. To address the lack of measurements, the HVAC load is approximated from the equipment commands. The zone temperature and equipment command trajectories are filtered to remove the high frequency dither induced by the on-off operation and to yield a time-averaged HVAC load, respectively. The approach is demonstrated on a simulated residential zone application.

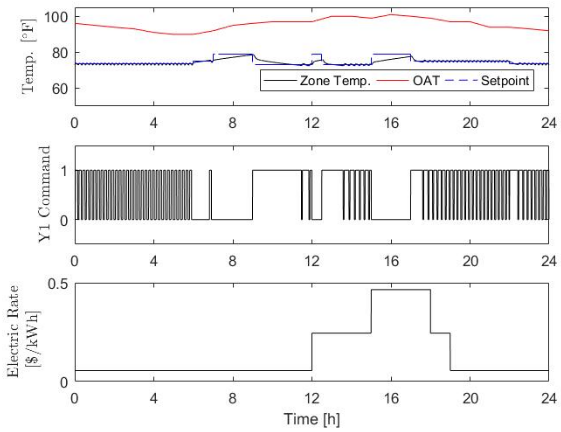

Figure 5: Simulated closed-loop temperatures, Y1 commands (cooling stage on/off command), and electric rate trajectories under MPC.

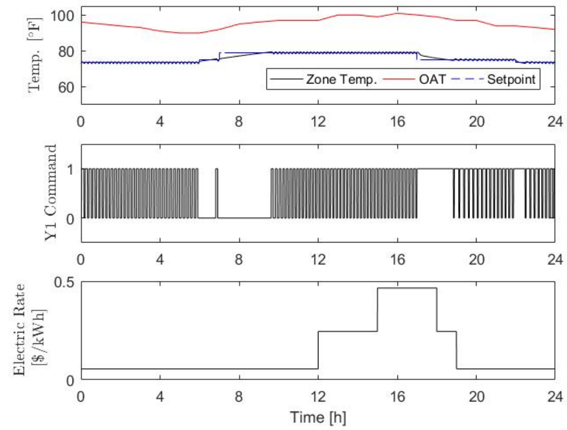

Figure 6: Simulated closed-loop temperatures, Y1 commands (cooling state on/off command), and electric rate trajectories with the scheduled occupancy setpoint.

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.

Alanqar, A., Ellis, M. J., Tapiero, J. E., & Wenzel, M. J. (2018). Practice-oriented system identification strategies for MPC of building thermal and HVAC dynamics. In *Proceedings of the 5th International High Performance Buildings Conference* (p. 3142). West Lafayette, IN.

ElBsat, M. N., & Wenzel, M. J. (2016). Load and electricity rates prediction for building wide optimization application. In *Proceedings of the 4th International High Performance Buildings Conference* (p. 3647). West Lafayette, IN.

Ellis, M., Durand, H., & Christofides, P. D. (2014). A tutorial review of economic model predictive control methods. *Journal of Process Control*, *24*, 1156–1178.

Ellis, M. J., Wenzel, M. J., & Turney, R. D. (2016). System identification for cascaded model predictive control of building region temperature applications. In *Proceedings of the 4th International High Performance Buildings Conference* (p. 3583). West Lafayette, IN.

Henze, G. P. (2013). Model predictive control for buildings: a quantum leap? *Journal of Building Performance Simulation*, *6*, 157–158.

Ma, J., Qin, S. 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 *Proceedings of the 51st IEEE Conference on Decision and Control* (pp. 6963–6968). Maui, HI.

Pannocchia, G., & Rawlings, J. B. (2003). Disturbance models for offset-free model-predictive control. *AIChE Journal*, *49*, 426–437.

Patel, N. R., Risbeck, M. J., Rawlings, J. B., Wenzel, M. J., & Turney, R. D. (2016). Distributed economic model predictive control for large-scale building temperature regulation. In *Proceedings of the 2016 american control conference* (pp. 895–900). Boston, MA. doi: 10.1109/ACC.2016.7525028

Rawlings, J. B., Angeli, D., & Bates, C. N. (2012). Fundamentals of economic model predictive control. In *Proceedings of the 51st IEEE Conference on Decision and Control* (pp. 3851–3861). Maui, HI.

Tapiero Bernal, J. E., & Ellis, M. J. (2017, May). *Building thermal mass energy storage control-oriented model and parameter identification for glas mpc* (Tech. Rep.). Johnson Controls White Paper.

Wenzel, M. J., Turney, R. D., & Drees, K. H. (2016). Autonomous optimization and control for central plants with energy storage. In *Proceedings of the 4th International High Performance Buildings Conference* (p. 3548). West Lafayette, IN.