

A Distributed Approach to Efficient Model Predictive Control of Building HVAC Systems

Vamsi PUTTA^{1*}, Guangwei ZHU¹, Donghun KIM², Jianghai HU¹, James BRAUN²

¹School of Electrical and Computer Engineering, Purdue University,
West Lafayette, IN, United States
{vputta, guangwei, jianghai}@purdue.edu

²School of Mechanical Engineering, Purdue University,
West Lafayette, IN, United States
{kim1077, jbraun}@purdue.edu

* Corresponding Author

ABSTRACT

Model based predictive control (MPC) is increasingly being seen as an attractive approach in controlling building HVAC systems. One advantage of the MPC approach is the ability to integrate weather forecast, occupancy information and utility price variations in determining the optimal HVAC operation. However, application to large-scale building HVAC systems is limited by the large number of controllable variables to be optimized at every time instance. This paper explores techniques to reduce the computational complexity arising in applying MPC to the control of large-scale buildings. We formulate the task of optimal control as a distributed optimization problem within the MPC framework. A distributed optimization approach alleviates computational costs by simultaneously solving reduced dimensional optimization problems at the subsystem level and integrating the resulting solutions to obtain a global control law. Additional computational efficiency can be achieved by utilizing the occupancy and utility price profiles to restrict the control laws to a piecewise constant function. Alternatively, under certain assumptions, the optimal control laws can be found analytically using a dynamic programming based approach without resorting to numerical optimization routines leading to massive computational savings. Initial results of simulations on case studies are presented to compare the proposed algorithms.

1. INTRODUCTION

The subject of optimal control of building heating, ventilation and air-conditioning (HVAC) systems has been receiving increased attention in the wake of climate change and soaring energy prices. With commercial floor space predicted to grow, the emphasis on reduction of HVAC energy consumption is warranted. However, operating building HVAC systems in an “optimal” way can be infeasible in real time, primarily due to the large number of decision variables to be controlled and the non-linear models involved. Additionally, utilizing weather and occupancy information to achieve higher savings in HVAC operation is not straight forward.

Model predictive control (MPC) has long been viewed as a practical solution for complex control problem involving non-linear dynamics and general cost functions. Efforts have been made to formulate and solve the optimal HVAC operation problem in an MPC framework. MPC based approaches also have the benefit of being capable of incorporating weather forecasts, utility pricing and occupancy profiles into the optimization. However, the large number of decision variables involved can make such approaches prohibitively slow for implementation in large buildings.

In this paper, we approach the problem of optimal HVAC control from the perspective of a distributed optimization problem. Such an approach enables us to decompose the original problem with a large number of decision variables into smaller optimization problems that can be solved simultaneously. The resulting solutions can be aggregated to obtain the solution of the original problem. Previous works in this direction include (Oldewurtel et al., 2010; Prívvara et al.; Zavala et al., 2011; Ma, Y. et al., 2010) We also propose an MPC algorithm that uses move-blocking to effectively reduce the dimension of the optimization problem to be solved at each time step. A dynamic

programming based MPC is also studied that under certain assumptions yields an easily computable analytic solution to the optimization problem. Initial results from simulations of a simplified case study are presented to compare the performance of the proposed methods.

The paper is organized as follows. In Section 2, we discuss the building and HVAC system models considered. The optimal control problem is formulated as a MPC problem in Section 3. This formulation is subsequently extended to a multi-zone scenario in a distributed optimization framework. Section 4 considers the use of dynamic programming to solve HVAC optimal control. Conclusions are drawn and future directions are given in Section 5.

2. MODELING

2.1 Thermal Zone Model

For the purpose of illustration, we consider a simplified state-space model of thermal zone and the accompanying air-handling unit (AHU). To reduce complexity, we only consider the thermal dynamics of the zone envelope and neglect spatial variations in the indoor air. This assumption enables us to use a single temperature, referred to as the zone temperature T , to represent the indoor air state. Using a detailed energy balance at all discrete wall nodes and the internal air node, we can obtain a state-space model that incorporates all the transients within walls and solar radiation through windows. Through linearization, discretization and appropriate model order reduction techniques, the model may be expressed in the form

$$\begin{aligned} x_{t+1} &= Ax_t + Bu_t + Fw_t, \\ T_t &= Cx_t \end{aligned} \quad (1)$$

where A, B, F and C represent the system matrices of reduced dimension obtained via model order reduction and t denotes the discrete time instant. The state vector x_t represents a transformed vector containing information about the temperatures of the wall and air nodes. Physical significance of each component of the state vector is not explicit due to the transformation. The vector u_t represents the input vector comprising of controllable inputs that act directly on the internal temperatures (rate of energy added by AHU, internal gains) and the matrix B encapsulates the effect of these inputs on the system. Vector w_t denotes the exogenous (uncontrollable) inputs acting on the envelope (solar radiation, ground radiation). The relation between the zone temperature T_t and the state vector x_t is modeled by the output matrix C . This particular formulation allows for flexibility in the treatment of the exogenous inputs—as disturbances with known parameters or completely known trajectories. A model of the form in (1), has been developed for the Purdue Living Lab Radiant Room (facility currently under construction). For the current study, the Living Lab model has been adopted as the test case.

2.2 Multi-zone Model

The above single-zone model can be extended to a more realistic scenario where the dynamics of neighboring zones are coupled. In the current study, we model output coupling where the heat flux into a zone from its neighbor is proportional to the temperature differential between the zones. This choice of coupling is motivated by a situation when the neighboring zones have a shared opening. We illustrate the idea using a two-zone building. Using the same principles as used in modeling the single thermal zone, we can write the coupled dynamics of both zones as

$$\begin{aligned} x_{t+1}^{(1)} &= A^{(1)}x_t^{(1)} + B^{(1)}[u_t^{(1)} + \alpha(T_t^{(2)} - T_t^{(1)})] + F^{(1)}w_t^{(1)}, \\ x_{t+1}^{(2)} &= A^{(2)}x_t^{(2)} + B^{(2)}[u_t^{(2)} + \alpha(T_t^{(1)} - T_t^{(2)})] + F^{(2)}w_t^{(2)}, \\ T_t^{(1)} &= C^{(1)}x_t^{(1)} \\ T_t^{(2)} &= C^{(2)}x_t^{(2)} \end{aligned} \quad (2)$$

where the numerical superscripts distinguish between the zones. The strength of the coupling is determined by the factor α . Similar models can be developed for more than two zones where the coupling strength α can be adjusted to model the relations between the dynamics as needed.

As a test scenario, we consider two identical zones based on the model of the Purdue Living Lab Radiant Room, with individual, identical AHU's. We also assume that the effect of weather and solar radiation to be identical on both the zones. Such a scenario presumes that the zones are small enough and close by to ignore variations in solar incidence and wind conditions. Individual AHU's imply the total power consumption is the sum of the individual power consumptions.

2.3 AHU model

The AHU units were assumed to comprise of heating and cooling coils along with the requisite pumps and fans. The governing equations for the power consumption of each component are written (simplifying assumptions such as constant efficiency and pressure ratios are assumed). The resulting equations are then summed up to generate the total power consumption P and AHU output u_t as a function of the mass flow rate (\dot{m}_{vent}) into the zone, the temperature of the air supplied by the AHU (T_{vent}), zone temperature T , ambient wet-bulb temperature T_{oa} :

$$\begin{aligned} P &= f(\dot{m}_{vent}, T_{vent}, T, T_{oa}), \\ u &= g(\dot{m}_{vent}, T_{vent}, T, T_{oa}). \end{aligned} \quad (3)$$

It must be noted that the actual controllable variables \dot{m} and T_{vent} affect the dynamics only through the rate at which heat is added (u). Hence we can reduce the problem of AHU operation to that of determining optimal value of u . Once the optimal trajectory is obtained, the requisite mass flow rate \dot{m}_{vent} and the air temperature T_{vent} that lead to the least power consumption can be back-calculated and applied. This observation allows us to succinctly capture the power consumption of the AHU in terms of the heat supplied to the zone using the relation

$$P^*(u; T, T_{OA}) = \min_{\dot{m}_{vent}, T_{vent}} P(T, \dot{m}_{vent}, T_{vent}, T_{OA}). \quad (4)$$

Here $P^*(u; T, T_{OA})$ represents the minimum power consumption of an AHU, when supplying heat energy at rate u to the zone at temperature T and ambient temperature T_{OA} . This representation of the power consumption will be used throughout the simulations.

The next section describes the formulation of the problem in the MPC framework. We define the objective function and explore the need for efficient MPC.

3. CONTROLLER DESIGN

3.1 Model Predictive Control Formulation

Model predictive control utilizes a predictive model to anticipate the behavior of the system over a prediction horizon L , and uses this information to decide upon the optimal course of action. The optimality of the decision is highly sensitive to the accuracy of the model and the forecast. Receding horizon control, where the forecast is updated every time instant, has proved highly effective in reducing this sensitivity.

In the current application, model predictive control allows us to accurately incorporate the uncontrollable factors such as variations in the occupancy, utility rates and weather conditions. Throughout the study, we assume availability of forecasts for all the exogenous inputs over the prediction horizon L . We use the inherent robustness of the receding horizon controller to handle inaccuracies in the forecasts. For the two-zone model presented in Section 2.2 the predicted trajectory can be represented as

$$\begin{aligned} x_{k+t+1|t}^{(i)} &= A^{(i)} x_{k+t|t}^{(i)} + B^{(i)} \left[u_{k+t}^{(i)} + \alpha \Delta T_{k+t|t}^{(i)} \right] + F^{(i)} w_{k+t|t}^{(i)} \\ T_{k+t|t}^{(i)} &= C x_{k+t|t}^{(i)}; \quad x_{t|t}^{(i)} = x_t^{(i)} \end{aligned} \quad (5)$$

Here $x_{k+t|t}^{(i)}$ represents the predicted state trajectory at time $t+k$ starting from time t under the action of $u^{(i)}$ and exogenous input $w^{(i)}$. The term $\Delta T_{k+t|t}^{(i)}$ denotes the predicted temperature differential.

In the two zone coupled model in equation (3), optimal HVAC operation would entail minimizing the total power consumption of both the AHUs while maintaining occupant comfort, or equivalently

$$\begin{aligned} J_t &= J_t^{(1)} + J_t^{(2)} \\ J_t^{(i)} &= \sum_{t=0}^L \left[r_t P^* \left(u_{k+t}^{(i)}, T_{k+t|t}, T_{OA_{t+k|t}} \right) + \gamma D(t+k, x_{t+k|t}^{(i)}) \right], \quad i = 1, 2 \end{aligned} \quad (6)$$

The term P_i^* represents the power consumption of the AHU of zone i as described earlier while r_t denotes the utility price at time t . This factor models the realistic scenario where electricity costs more during peak hours and encourages savings through load shifting. $D(t, x_t^{(i)})$ is a measure of the occupant discomfort in zone i at time t . Adjusting the factor γ prioritizes one of the two competing objectives. The integral nature of energy costs is

reflected in the summation over a look ahead horizon of L . The cost function is to be minimized subject to the dynamics given in equation (2) over the space of all admissible inputs $u^{(i)}$ that do not violate any physically imposed constraints (capacity of AHU). At time t the optimal trajectory of both the AHUs ($u_{t+k}^{(i)}, k = 0, 1, \dots, L$) are determined. The first input of the sequence $u_t^{(i)}$ is applied to the corresponding system. At time $k + 1$ the cost function and forecasts are updated to reflect the information available and the process repeated. The prediction horizon L is chosen to be large enough to sufficiently capture the behavior (such as periodicity) of the exogenous factors. We also presuppose knowledge of the state vectors $x_t^{(i)}$. The cost function can be extended to more than two zones.

The choice of the occupant discomfort metric $D(\cdot, \cdot)$ is an important factor in ensuring realistic results. Direct comfort models such as the Predicted Mean Vote (PMV) and Predicted Percentage Dissatisfied (PPD), though accurate, do not lend themselves to easy optimization. A more convenient method to quantify discomfort is to measure the deviation of the zone temperature from a predetermined comfort interval. As occupant discomfort needs to be prioritized during working hours, a factor representing the occupation profiles is also incorporated. The resulting discomfort measure can be written as

$$D(t, x_t) = \begin{cases} 0, & \text{if } T_t \in [T_L, T_U] \\ n_t(T_L - T_t), & \text{if } T_t < T_L \\ n_t(T_t - T_U), & \text{if } T_t > T_U \end{cases} \quad (7)$$

where $[T_L, T_U]$ denotes the comfortable temperature range. The occupancy profile n_t takes the values 1 or 0 depending upon whether the building is occupied or not respectively. In the special case, when $T_L = T_U$ the problem becomes one of temperature tracking.

Despite the apparent independence of the AHU units (and their cost functions) the coupling between the zone dynamics causes the optimal AHU controls to be coupled as well. If the coupling magnitude is small enough, each zone is effectively independent of the other and the optimization is performed individually for each zone. Alternatively, we can treat the coupling as an exogenous input that can be forecast at every instant and utilize the inherent robustness of receding horizon control to independently optimize the AHU controls simultaneously. We describe such a distributed optimization based algorithm in the following section.

3.2 Distributed MPC Formulation

Distributed optimization approaches have proved to be successful in large scale optimization problems. Recently, researchers have tried to apply distributed approaches to optimizing building system operation (Ma, Y. et al., 2009, 2011; Morosan et al., 2010; Zaheer-Uddin et al., 1993). Distributed approaches reduce computational times by simultaneously solving reduced dimensional problems at a subsystem level independently. The independent solutions so obtained can be improved via information exchange between the subsystems.

In this paper, we propose a distributed MPC algorithm that handles coupled dynamics between zones as in equation (2). Each controller uses an estimate of the temperature differential to predict future trajectories and obtain an optimal control. Using the receding horizon principle, the temperature differential estimate is updated at every time step. This allows the controllers to attain reasonable performance with a significant gain in optimization time.

An important consideration while implementing a distributed optimization problem is the information to be exchanged between the controllers. Ideally to save networking costs between the controllers, the information to be exchanged must be kept to a minimum while maintaining reasonable control performance. In the current two-zone test case, we found that passing the average zone temperature of one zone to another was sufficient to obtain reasonable performance. In particular, each zone computes a moving average of its zone temperature over a period of length N_{pred} and passes it to the other zone at every instant. The other zone then uses this information to compute an average temperature differential that is assumed to remain constant during the look-ahead interval for the predicted trajectory computation. Using the notation from equation (6), the predicted state trajectories have dynamics given by

$$\begin{aligned} x_{k+t+1|t}^{(i)} &= A^{(i)}x_{k+t|t}^{(i)} + B^{(i)}[u_{k+t}^{(i)} + \alpha\Delta T_t^{(i)}] + F^{(i)}w_{k+t|t}^{(i)} \\ T_{k+t|t}^{(i)} &= Cx_{k+t|t}^{(i)}; & x_{t|t}^{(i)} &= x_t^{(i)}, \quad i = 1, 2 \end{aligned} \quad (8)$$

where $\Delta T_t^{(i)}$ represent the past average differential over a period of length L time steps.

$$\Delta T_t^1 = \frac{1}{L} \sum_{k=1}^L (T_{t-k}^{(2)} - T_{t-k}^{(1)}), \quad \Delta T_t^{(2)} = -\Delta T_t^{(1)} \quad (9)$$

As each controller performs its optimization simultaneously, the total time taken would remain the same irrespective of the number of zones (allowing for time taken to exchange required information). This makes it an attractive approach for large buildings with several zones.

The optimization problem can be extended to more than two zones but incurs a linear increase in the number of decision variables. An m zone building would require optimization with respect to mL control inputs over every prediction horizon. The values of the prediction horizon L are typically large enough to prohibit real-time optimization in commercial buildings.

To overcome this problem, we restrict the control input to remain constant for a few time steps. This procedure also called move-blocking reduces the number of decision variables thus speeding up computation (Ma, Y. et al., 2009). Using the occupancy and the price profile enables us to intelligently guess the time instants at which an extra degree of freedom is required. This procedure is described next.

3.3 Reduced Degree of Freedom Optimization

At any instant t , a single zone AHU controller has L degrees of freedom (corresponding to the u_{t+k} , $k = 0, 1, \dots, L$). However, changing the AHU control too frequently can be costly and wear out the machine. Traditionally, this problem is circumvented by trying to maintain the AHU control u at a constant level for the longest possible time. A useful heuristic is to allow a level change before occupancy or price changes to allow for load shifting via precool/heat and comfort maintenance. Using such time instants as an initial guess, the optimal level changes are computed. This problem has a significantly reduced dimension compared to the original problem as the number of occupancy and price changes in a look-ahead horizon are limited. Once the optimal control levels are obtained, the time instants for level changes are improved using a line search. This step determines the amount of time for optimal precooling/heating. The resulting piecewise constant control trajectory can be further refined by repeating the procedure till the trajectory converges. From the resulting trajectory, the optimal control required at the current instance is extracted and applied with the MPC updating its predictions as well. The whole process is repeated in real time to drive the system.

3.4 Simulation Results

As an initial case study, the distributed MPC proposed above is applied to the two zone model in Section 2.2. The zone dynamics are obtained from Purdue Living Lab model. After order reduction, each zone has a state space dimension of 10 and 16 exogenous inputs. The discretization time step is chosen to be 30 min. Internal gains were ignored; the only energy directly added to the air node was the output of the AHU and the coupling term due to the temperature differential. The coupling factor α is chosen to be 10.

Existing weather data (Indiana TMY2) from July 2010 was used to calculate the solar inputs shared by both zones for a 31 day period. Electricity prices variation was represented by a threefold increase in the cost per *kwh* during peak periods (10am-3pm everyday). Zone 1 was assumed to be occupied daily from 7am-6pm daily while Zone 2 was assumed to be occupied from 9:30am to 3:30pm daily. The occupancy and price variations are plotted in Figure 1.

Initial results of applying the distributed MPC discussed in Section 4 to the two-zone model are given in Figure 2. Zone 1 results are given in red while Zone 2 results are plotted in blue. The move-blocking principle is also applied during the computation of the optimal control law. A prediction horizon of $L = 24$ hours was used. The resulting control law maintains both the zone temperatures within the comfort region ($[T_L, T_U] = [22, 25]$). Due to the difference in occupancy profiles, there is an offset between the zone temperatures. The spiky artifacts in the control law of Zone 1 are presumably due to existence of local minima. It is observed that the late occupancy of Zone 2 enables its AHU to bring its temperature to the comfortable level much later. After the initial heating the AHU is turned off letting the solar inputs and the temperature differential to maintain the temperature. This is an intuitive choice for the most economical control whenever feasible.

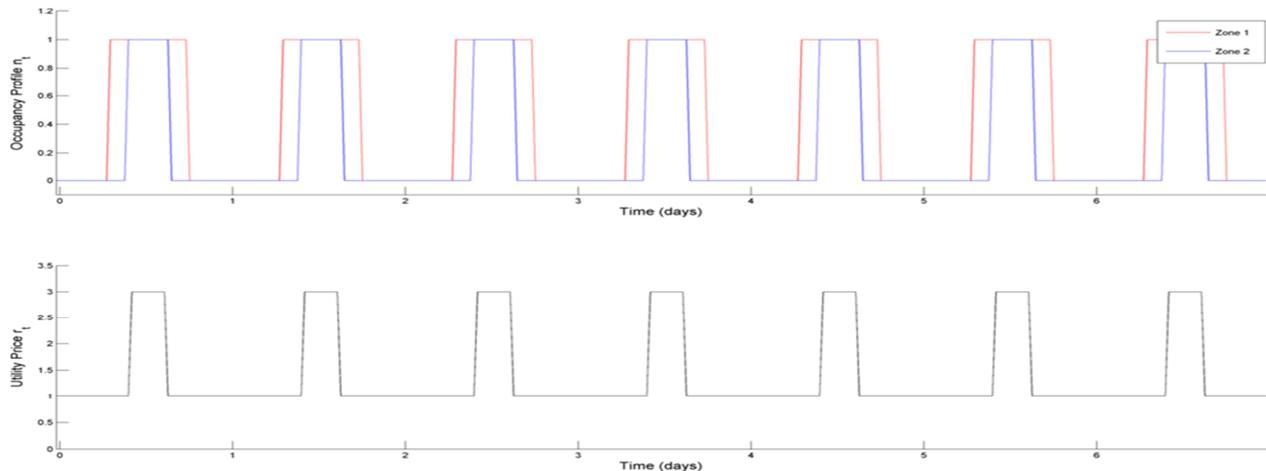


Figure 1: Occupancy and price profiles for the simulation.

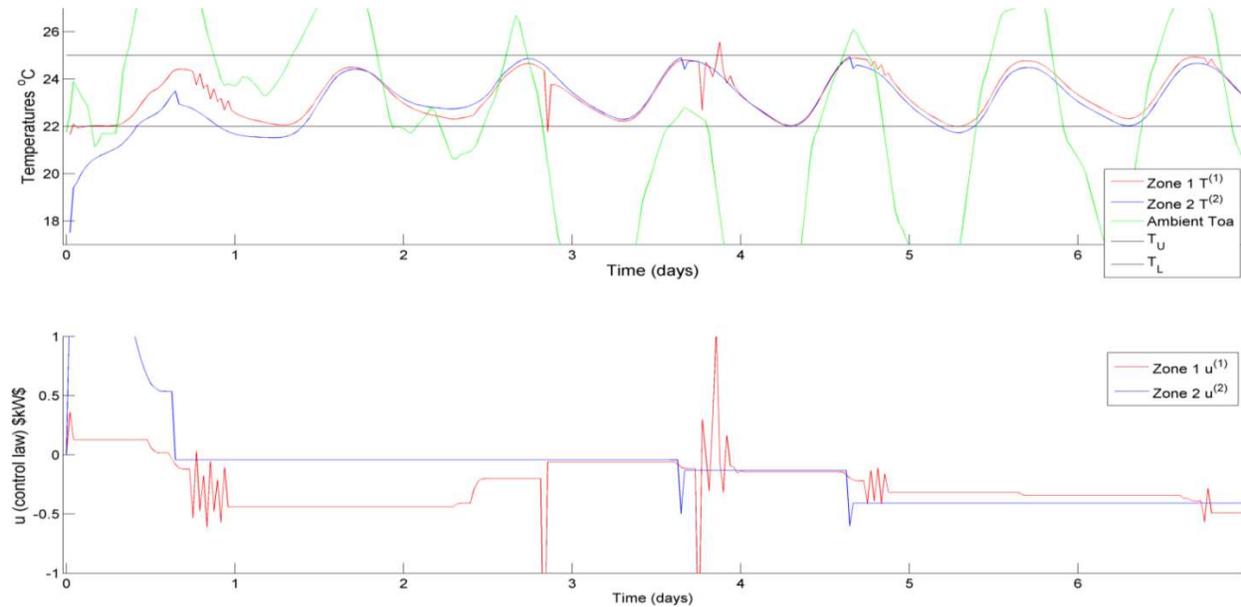


Figure 2: Simulation results of the distributed MPC

3.5 Effect of Information Exchange:

To see the effect of information exchange among the controllers, we compare the results of the previous simulation to those where the controllers ignored the coupling temperature differential while predicting future trajectories. This enables the controllers to work completely independently at the cost of reduced control performance. Figure 3 depicts the results of this simulation. Temperature fluctuation is higher than the previous case as expected. The lack of any information exchange frequently causes the predicted trajectories to deviate from real-life trajectories causing the controller to overcompensate.

The simulations of the distributed MPC with move-blocking requires 20 minutes on average to compute the complete optimal trajectory over a 7 day period on an Intel Core 2 Duo desktop. This shows that distributed optimization combined with move-blocking is a plausible candidate for real-time control design.

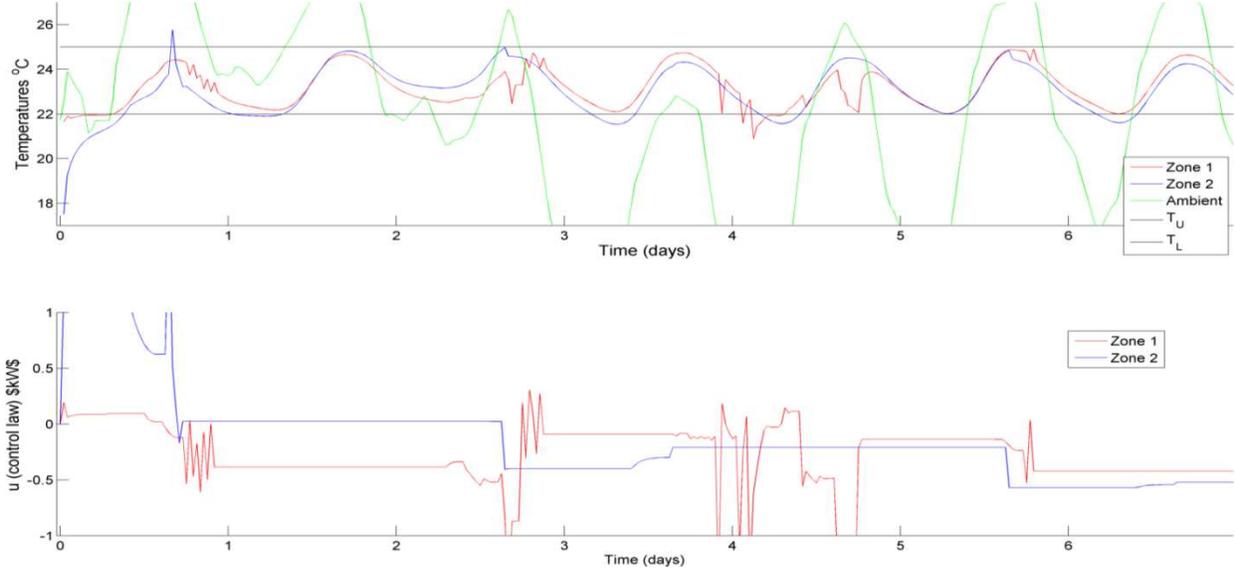


Figure 3: Distributed control with no information exchange between controllers

4. DYNAMIC PROGRAMMING APPROACH

In general, search-based optimization algorithms for HVAC MPC are computationally very costly due to the model complexity as well as the non-regularity of the cost functions. It is therefore interesting to look into another family of algorithms which are based on dynamic programming (DP) (Bertsekas, 1976). Dynamic programming provides an analytic solution to systems with linear dynamics and quadratic cost functions; its computational complexity is polynomial in the state space dimension. Due to its high efficiency, it forms an attractive approach in solving large scale problems. We now formulate the building HVAC control problem in a way more amenable for utilizing dynamic programming.

4.1 Affine Quadratic Regulator (AQR) Problem

Recall that the state-space model of a single zone is given by equation (1). We assume that at any time instant t , the cost function can be represented in the nonhomogeneous quadratic form shown as follows,

$$J_t(x_t, u_t) = (x_t - x_{t,ref})^T Q_t (x_t - x_{t,ref}) + (u_t - u_{t,ref})^T R_t (u_t - u_{t,ref}) \quad (10)$$

where Q_t, R_t are real symmetric positive semi-definite matrices with proper dimensions, and $x_{t,ref}, u_{t,ref}$ are offsets. We formulate the MPC as the following Affine Quadratic Regulator (AQR) problem:

$$\begin{aligned} \min_{u_{t+k|t}, k=0, \dots, L-1} J = & \left[\sum_{k=0}^{L-1} J_{t+k}(x_{t+k|t}, u_{t+k|t}) \right] + (x_{t+L|t} - x_{t+L,ref})^T Q_{t+L} (x_{t+L|t} - x_{t+L,ref}) \\ \text{subject to } & x_{t+k+1|t} = Ax_{t+k|t} + Bu_{t+k|t} + Fw_{t+k|t}, \forall k = 0, \dots, L-1 \\ & x_{t|t} = x_t \end{aligned} \quad (11)$$

where the notation $x_{s|t}$ denotes at current time t the forecast of the system state in a future time s , and $x_{s|t}, w_{s|t}$ are defined similarly. Positive integer L is the length of look-ahead horizon for the MPC. The forecast of the perturbations $w_{s|t}$ are always assumed known. At each time instant t , problem (11) is solved to obtain the best future input sequence $u_{t|t}, u_{t+1|t}, \dots, u_{t+L-1|t}$. Then the immediate input $u_{t|t}$, which is the first in the sequence, is imposed by the MPC to the system.

The AQR model provides a generic platform for MPC design based on dynamic programming. Our specific objective is to track the system output T_t around some given T_{ref} (reference temperature). Suppose the system output is given by $y = Cx$. The system input u_t is the heating/cooling power, which is a scalar. We consider the following cost-aware temperature tracking problem:

$$\min_{u_t} J = \left[\sum_{t=0}^{\infty} n_t (T_t - T_{ref})^2 + r_t u_t^2 \right] \quad (12)$$

Here the coefficients n_t and r_t encode the profiles of zone occupancy and utility price. To design the corresponding MPC based on the AQR framework, we can let

$$Q_t = n_t C C^T, \quad R_t = r_t \quad (13)$$

and pick any x_{ref} satisfying $Cx_{ref} = T_{ref}$. We use the dynamic programming approach to solve the above AQR problem, whose detailed procedures are omitted here due to space limit.

4.2 Simulation Result

We first tested the AQR-based MPC on a single-zone Purdue Living Lab model. The simulation result is depicted in Figure 4. The reference temperature is set to 23.5 degrees Celsius. The periodic solid and the dashed trajectories in top plot denote the *trend* of occupancy and utility price. We set $n_t = 500$ when occupancy is high, and $n_t = 0$ when low (meaning that there is no penalty for temperature deviation during that period). We set $r_t = 0.001$ when utility price is low, and $r_t = 0.003$ when high.

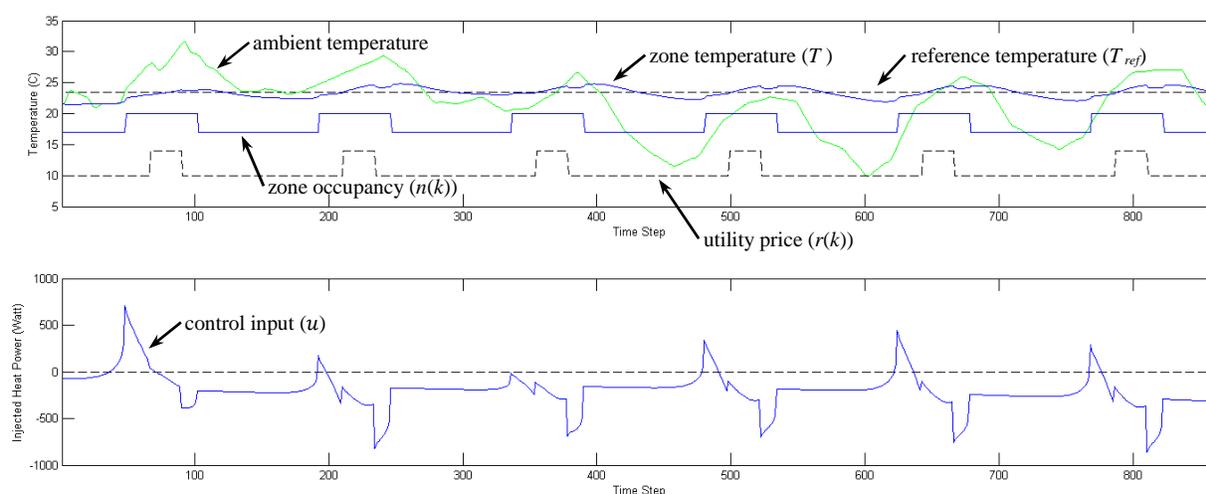


Figure 4: Simulation result of AQR-based MPC for single-zone Purdue Living Lab model

The simulation result indicates that the AQR-based MPC can effectively react to occupancy and utility price changes. To see this, we may note that there is a peak in heating power when the zone starts to be occupied. This is because the zone temperature is usually pre-cooled to below the reference temperature. When occupants move in, a shot of heat is injected to the zone to raise the zone temperature quickly. This approach enables the walls, which are cooled during night time, to cool down the zone when both ambient temperature and utility price rise later. We can also notice that when utility price is high, the MPC tends to save energy cost. Hence when the price drops, there is an increased amount of cooling to bring down the zone temperature. The MPC also tends to pre-cool the room using small amount of energy throughout the night time.

We kept the same parameters for utility prices and room occupancy, and applied the AQR-based MPC to the two-zone coupled model as described in equation (2), where the two zones differ in their initial states and occupancy profiles. In the two test cases adopted for simulation, we set $\alpha = 10$ for both cases. We also assume no information exchange between individual MPCs in one case, and perfect information exchange (i.e., equivalent to a centralized MPC for both zones) in the other. The simulation results are shown in Figure 5 and 6, where the blue curves and the red curves correspond to zone 1 and 2, respectively. It can be seen from the figures that with decent amount of coupling ($\alpha = 10$), the control outputs by two individual MPCs and by the centralized MPC are almost identical, whereas the computation load can be split in the first scenario, leading to further improvement of efficiency in the multi-zone scenario.

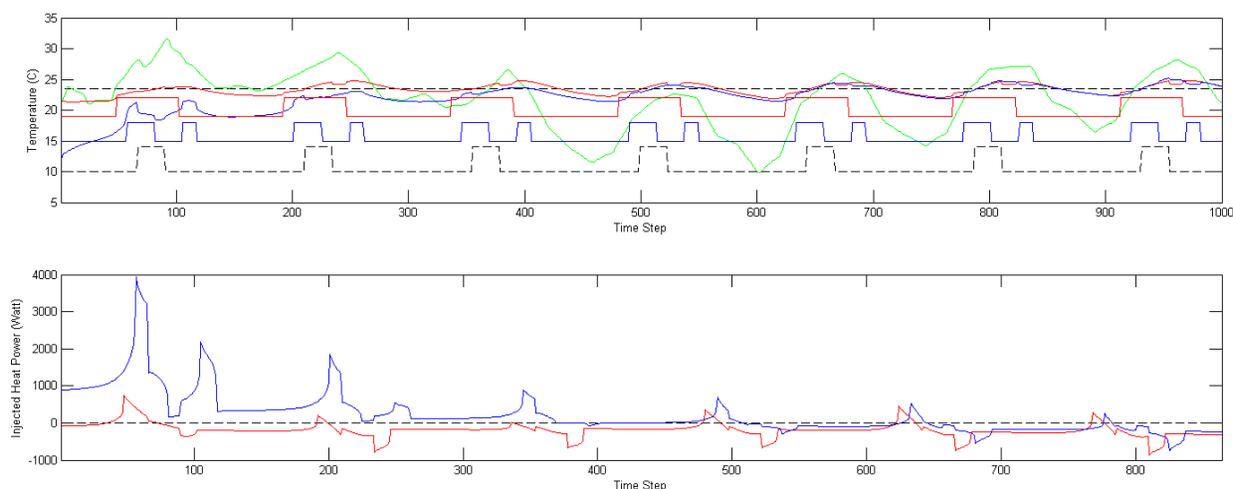


Figure 5: Simulation result of AQR-based MPC for two-zone model (no information exchange, $\alpha = 10$)

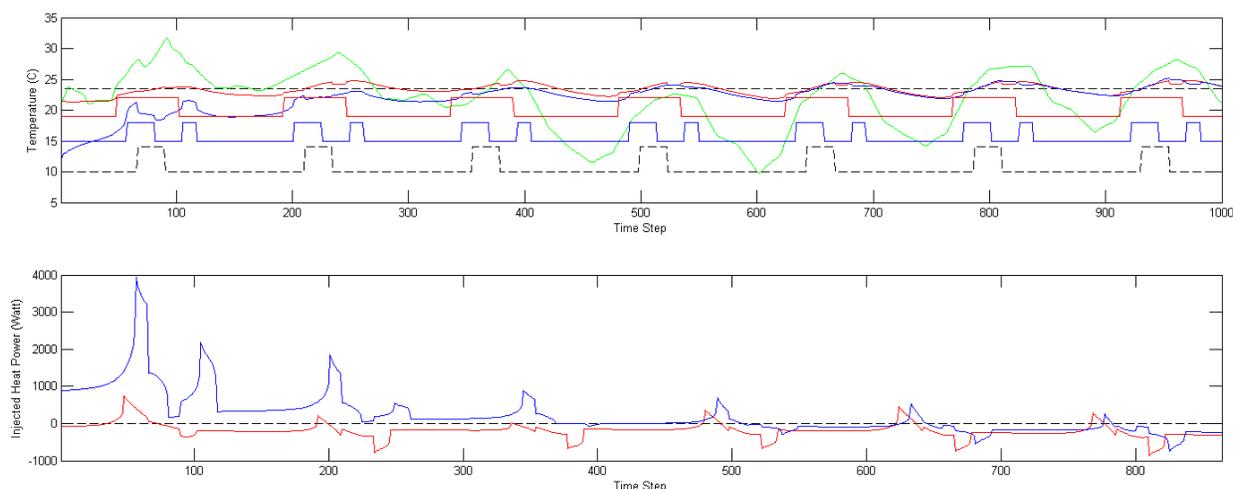


Figure 6: Simulation result of AQR-based MPC for two-zone model (perfect information exchange, $\alpha = 10$)

It is worth noting that the AQR-based MPC outperforms the majority of other MPC implementations in terms of computational efficiency. It takes less than three minutes to simulate 1000 time steps (1 time step = 10 min) with a look-ahead horizon of 432 time steps (3 days), while other implementations often take hours, or even days, to compute. This can be crucial if we extend the scenario to large-scale buildings, where controllers of several zones may need to perform local optimization and exchange information with each other. In such cases, computationally costly MPC implementations will not be able to afford multiple iterations for the information exchange process to converge in cooperative MPC schemes. We intend to investigate other information exchange schemes in more realistic scenarios in the future.

5. CONCLUSIONS

A distributed approach to optimal HVAC operation is presented. By exchanging information between two independent model predictive controllers, a computationally intractable problem can be solved simultaneously in real-time. An alternative approach based on dynamic programming is also presented. On proper formulation of the optimal control problem, the dynamic programming approach turns out to be highly efficient due to the existence of an analytical solution. Future directions include benchmarking the proposed approaches to test their optimality, simulations of more realistic scenarios and setting up distributed MPC in a multi-agent framework.

NOMENCLATURE

x	state vector		Superscripts
u	rate of heat addition	kw	(i) Zone i
w	solar and ambient inputs	kw	Subscripts
A,B,C,F	system matrices		t time
T	zone temperature	K	k+t t :predicted variable
r	utility rate	\$/kwh	
n	occupancy profile		
γ	discomfort weighting factor		
D	discomfort penalty		
L	prediction horizon		
ΔT	temperature differential	K	
J	cost function		
V	value function		

REFERENCES

- Bertsekas, D. P. (1976). *Dynamic programming and stochastic control* (Vol. 125). Academic Pr.
- Ma, Y., Anderson, G., & Borrelli, F. (2011). A distributed predictive control approach to building temperature regulation. *American Control Conference (ACC), 2011* (pp. 2089–2094). Presented at the American Control Conference (ACC), 2011, IEEE.
- Ma, Y., Borrelli, F., Hancey, B., Coffey, B., Bengoa, S., & Haves, P. (2010). Model predictive control for the operation of building cooling systems. *American Control Conference (ACC), 2010* (pp. 5106–5111). Presented at the American Control Conference (ACC), 2010, IEEE.
- Ma, Y., Borrelli, F., Hancey, B., Packard, A., & Bortoff, S. (2009). Model Predictive Control of thermal energy storage in building cooling systems. *Proceedings of the 48th IEEE Conference on Decision and Control, 2009 held jointly with the 2009 28th Chinese Control Conference. CDC/CCC 2009* (pp. 392–397). Presented at the Proceedings of the 48th IEEE Conference on Decision and Control, 2009 held jointly with the 2009 28th Chinese Control Conference. CDC/CCC 2009, IEEE.
- Morosan, P.-D., Bourdais, R., Dumur, D., & Buisson, J. (2010). A dynamic horizon distributed predictive control approach for temperature regulation in multi-zone buildings. *2010 18th Mediterranean Conference on Control & Automation (MED)* (pp. 622–627). Presented at the 2010 18th Mediterranean Conference on Control & Automation (MED), IEEE.
- Oldewurtel, F., Parisio, A., Jones, C. N., Morari, M., Gyalistras, D., Gwerder, M., Stauch, V., et al. (2010). Energy efficient building climate control using stochastic model predictive control and weather predictions. *American Control Conference (ACC), 2010* (pp. 5100–5105).
- Prívará, S., Široký, J., Ferkl, L., & Cigler, J. Model predictive control of a building heating system: The first experience. *Energy and Buildings*, 43(2-3), 564–572. Retrieved from <http://www.sciencedirect.com/science/article/pii/S0378778810003749>
- Zaheer-Uddin, M., Patel, R. V., & Al-Assadi, S. A. K. (1993). Design of decentralized robust controllers for multizone space heating systems. *Control Systems Technology, IEEE Transactions on*, 1(4), 246–261.
- Zavala, V. M., Skow, D., Celinski, T., & Dickinson, P. (2011). *Techno-Economic Evaluation of a Next-Generation Building Energy Management System*. Technical Report ANL/MCS-TM-313, Argonne National Laboratory.

ACKNOWLEDGEMENT

The authors wish to acknowledge Dr. George Chiu and Jie Cai for their valuable suggestions.