The Effect of Input Uncertainty in Model-based Predictive Control

H. Burak Gunay[1], William O'Brien[1], Ian Beausoleil-Morrison[2], Brent Huchuk[1], Michael Palmer[3], Jim Fletcher[3], Alexandre Pavlovski[3]
[1] Carleton University, Dept. of Civil and Environmental Engineering, Ottawa, Ontario
[2] Carleton University, Dept. of Mechanical and Aerospace Engineering, Ottawa, Ontario
[3] Green Power Labs, Dartmouth, Nova Scotia

Abstract

The adoption of model-based predictive control (MPC) strategies for building systems, such as auxiliary heating/cooling systems, electric lighting systems, operable window and window shading systems, represents a significant potential to reduce the environment impact of buildings. However, the potential of MPC has been challenged with the limited nature of the information available to the controller to make optimal operational decisions. This paper proposes a self-adaptive MPC method: one in which the control system (1) learns the building physics via a sequential non-linear filter —processing the data coming from a network of sensors/meters; (2) predicts its response using a first-order discrete-time state-space model—processing data coming from forecasts; and, (3) adapts the operation of a heating/cooling system by employing the Newton's method in optimization. The potential of this MPC method is demonstrated through a simulation-based study focused on a hypothetical south-facing office. The performance of the MPC method is contrasted to a classical reactive controller. Its sensitivity to varying resolution and availability of sensors, meters and forecasts is studied. The simulation results indicate that such a method could lead to significant energy savings, even with the limitations/challenges related with the availability and resolution of these inputs. Such a system may also enhance the controllability of the indoor climate within the designated setpoints.

1 Introduction

1.1 Motivation

Buildings play a significant role in our environmental impact (e.g., 53% of electricity use, 35% of greenhouse gas emissions) and adoption of innovative operation strategies for building systems, such as auxiliary heating/cooling systems, electric lighting systems, operable window and window shading systems, represents a great potential to reduce this impact.

Oftentimes, these operation strategies involve the incorporation of heuristics (e.g. optimal start/stop, night purge) in classical controllers (e.g. on-off or proportional control). However, heuristics require having access to (and the incurred cost of) experienced personnel to implement and commission the strategies. In theory, a self-adaptive model predictive controller (MPC) should systematize the optimal operational decisions and reduce the dependency on the experienced personnel.

1.2 Background on Model-based Predictive Control

A self-adaptive MPC controller can typically be pictured as three agents handling three distinct tasks: (1) learning, (2) predicting, and (3) adapting. The learning agent is an inverse model. As schematically illustrated in Figure 1.a, it interprets the real-time measure-

ments and disturbances recursively to learn characteristics of the indoor environment (e.g. effective thermal resistance and capacitance). To this end, several researchers have studied various inverse model training methodologies. For example, Fux et al. (2014) studied extended Kalman Filters (EKF). They reported that EKF is capable of yielding robust convergence to the model parameters after three weeks. They also analyzed the influence of the model complexity and reported that a 1R1C model is adequate to yield short term forecasts of the building response. Radecki and Hencey (2012) compared the performance of EKF with an unscented Kalman Filter (UKF) using a 1R1C model and reported that UKF as a state and parameter estimation algorithm yields more stable estimates than EKF. In line with this, Maasoumy et al. (2012) compared the performance of EKF and UKF using a relatively more detailed building model (5R5C) and reported that the UKF algorithm yields better estimates than the EKF algorithm.

The predicting agent is a forward model which predicts the future states in the receding horizon; e.g. indoor temperature in the next six hours. As schematically illustrated in Figure 1.b., it inputs the parameters updated by the learning agent and forecasted disturbances. To this end, numerous researchers (Candanedo et al., 2013, Sourbron et al., 2012, Candanedo and Athienitis, 2010, Fraisse et al., 2002, Chen and Athienitis, 2003) studied the predictive accuracy of reduced order building models and their viability in MPC. They reported common findings that reduced order building models can accurately predict the short term temperature response of a control zone.

The adapting agent is an optimization algorithm which chooses the optimal control step by minimizing an objective function (e.g. energy use or utility cost) subject to comfort and equipment capacity-related constraints in the receding horizon. The principle of receding horizon control can be best illustrated through a descriptive example. For example, as shown in Figure 1.c, at each time-step the agent determines a sequence of discrete control states and it only actuates the first of this control sequence. In the following time-step, the agent repeats the identical process, while the discrete prediction horizon recedes to the neighbouring time-step. Unlike the reactive controllers, the adapting agent assigns weights to the predicted future states while determining the magnitude of the current control signal. This can be illustrated in Figure 1.c., at time=1, the heating was actuated despite the fact that the indoor temperature is above the heating setpoint. Here, the researchers investigated different approaches to achieve the optimal discrete control sequence in the receding finite horizon. For example, Oldewurtel et al. (2012) and Wang and Ma (2008) employed linear-programming methodologies; while Nghiem and Pappas (2011), Liu et al. (2013) Prívara et al. (2013), and Henze et al. (2005) employed non-linear programming approaches to solve for the optimal control sequence in a finite time horizon.

1.3 Problem Definition

The potential of MPC has been challenged with the sparsely and temporally limited nature of the information available for the controller to make the optimal operational decisions. Such information can be summarized as follow: (1) How much solar radiation is currently transmitting and how much will be transmitted into the control zone in the next a few hours? (2) How much internal gains is currently being emitted and how much will be emitted by the lighting, occupants, and equipments in the prediction horizon? (3) What is the current outdoor temperature and what will it be in the prediction horizon? (4) What is the effective thermal capacitance and resistance of the control zone?

While it is impossible to acquire absolute certainty in the validity of parameters, it is possible to gain increasing confidence with an increasing effort to acquire a finer grid of sensors/meters and using weather forecasts at a finer resolution. It is therefore crucial to quantify the sensitivity of a control strategy over these uncertainties and to choose the acceptable level

of detail in our controller models. The building performance simulation (BPS) lends itself (as a perfectly repeatable environment) to test the robustness of a MPC strategy over a wide-range of synthetic sources of uncertainty emerging from availability of the aforementioned inputs.

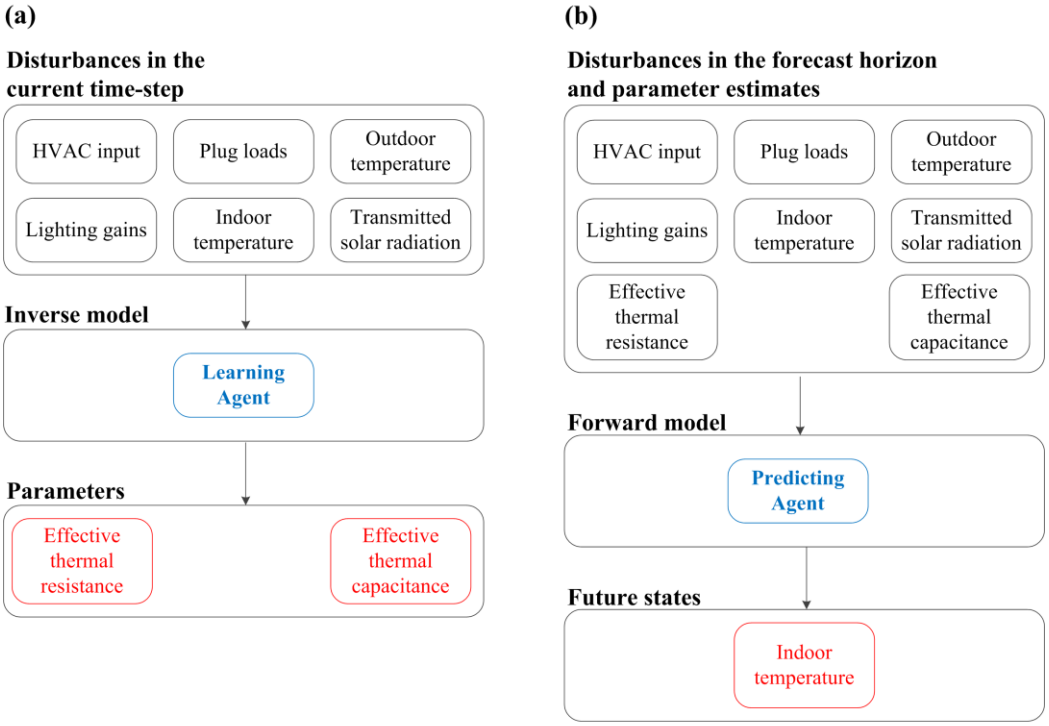

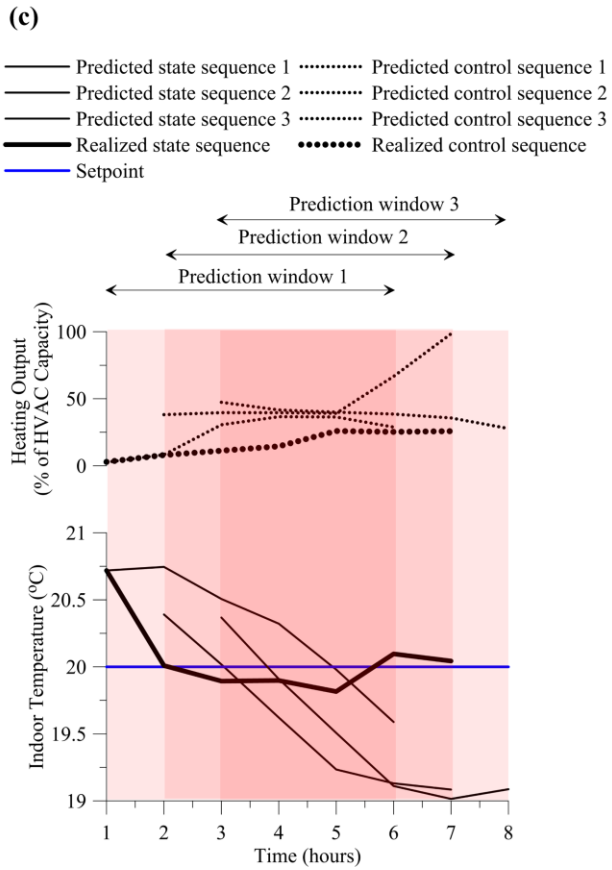

Figure 1: Schematics describing the (a) learning agent, (b) predicting agent, and (c) an example illustrating the receding horizon concept in the adapting agent of a self-adaptive MPC.

1.4 Document Structure

The objective of the current paper is to seek numerical support for the robustness of a self-adaptive MPC strategy by contrasting it to a classical reactive controller; while changing the resolution and availability of the sensors, meters, and forecasts accessible for it to undertake the optimal operational decisions. Section 2.1 describes the numerical approach whereby the run cases were presented. Section 2.2 describes the building and the occupant models whereby the EnergyPlus building model, idealized heating/cooling system, and the occupant behaviour models are presented. In Section 2.3, the control models are presented. The reactive control (in Section 2.3.1) model is a classical proportional (P) controller. The predictive controller (in Section 2.3.2) consists of a learning agent based on a particle filter, a predicting agent based on a first-order discrete-time state-space model, and an adapting agent based on the Newton's method in optimization. In the results section (Section 3), the performance indicators of the reactive model were summarized. Also, the learning accuracy of the predictive controller to varying resolution/availability of sensors and meters was analyzed. Furthermore, the predictive accuracy of the predicting agent to varying resolution/availability of forecasts was investigated. Subsequently, the performance indicators with the proposed controller for each run case were presented and discussed by contrasting them with the reactive controller. In Section 4, concluding remarks were stated and future research goals were defined.

2 Methodology

2.1 Numerical Approach

The run cases for the numerical experiments are summarized in Table 1 and 2. The resolution of the existing meters and sensors to measure the ambient temperature (T_{amb}), the indoor temperature (T_{in}), the transmitted solar radiation (Q_{sol}), the lighting (Q_{light}), HVAC (Q_{hvac}), and equipment (Q_{plug}) heat inputs to the control zone were varied as shown in Table 1. The Case 1L represents the ideal recursive learning case whereby inputs like transmitted solar radiation or heat inputs was (somehow) known. The run Case 2L and Case 3L represent situations whereby the application controller (i.e. lighting and blinds controller) was not integrated to the rest of the control network or the control network does not provide such high resolution inputs for the particular zone; and therefore plug loads and/or HVAC states were unknown for the learning agent. The run Case 4L represents a case whereby the transmitted solar radiation was only estimated using a sensor on the facade (without a feedback from the blind position).

Table 1: Run cases to study the influence of the existing meter and sensor resolution

Parameter Learning	Existing meters and sensors to measure					
	T_{amb}	T_{in}	Q_{sol}	Q_{light}	Q_{hvac}	Q_{plug}
Case 1L	Known	Known	Known	Known	Known	Known
Case 2L	Known	Known	Known	Known	Known	Unknown
Case 3L	Known	Known	Known	Known	Unknown	Unknown
Case 4L	Known	Known	Non-ideal	Unknown	Unknown	Unknown

The run cases Case 1P, 2P, and 3P (see Table 2) were selected to reveal the influence of the forecast availability for the disturbances; i.e. T_{amb}, Q_{sol}, Q_{light}, and Q_{plug} for the next six hours. The run Case 1P represents the ideal case whereby the observation data is somehow available to the controller. In the run Case 2P, it was assumed that the knowledge of Q_{light} and Q_{plug} in the next six hours is not predictable (thus unavailable to the controller). In Case 3P, instead of ideal forecasts, historical forecasts made by the North American Mesoscale Forecast System (NAM) were used. The NAM forecasts are publicly available four times a day at 13 km resolution. Individual historical NAM forecasts — for up to sixth hour — were merged

to create a continuous data record. Thus, in run Case 3P the short term forecast for the solar radiation intensity on the south-facade was taken instead of the exact amount of solar radiation that will transmit into the zone in the next six hours. Similarly, the ideal knowledge of dry bulb temperature in the next six hours was replaced with the short term forecasts available for the controller.

Table 2: Run cases to study the influence of the forecast availability

Predicting T_{in}	Forecasts available for			
	T_{amb}	Q_{sol}	Q_{light}	Q_{plug}
Case 1P	Ideal	Ideal	Ideal	Ideal
Case 2P	Ideal	Ideal	Unknown	Unknown
Case 3P	Non-ideal	Non-ideal	Unknown	Unknown

2.2 Building and Occupant Models

A simple thermal model of the hypothetical office in Halifax, Canada (see Figure 2) was created with EnergyPlus v8.0. The office had 9 m^2 of south-facing exterior envelope area, 3 m^2 of which were glazed. The area of the concrete floor slabs was 9 m^2 and it was 0.10 m thick; the ceiling slab was identical. A carpet was placed on top of the floor slab while the occupied volume of the office was separated from a plenum space by ceiling tiles. The three interior walls of the office were treated as adiabatic.

Observational data (dry bulb temperature, relative humidity, wind direction and speed, global horizontal short-wave solar radiation) representing the weather conditions were extracted from an on-site weather station in Dartmouth, Halifax for the time period between January 1st and October 31st, 2013; and converted to the EnergyPlus weather file format.

EnergyPlus' default conduction transfer function (CTF) algorithm was used to represent transient heat conduction through the opaque portion of the south-facing exterior envelope. Fenestration and heat transmission through the glazing was treated in a simplified fashion using a constant solar heat gain coefficient (SHGC) of 0.6, a constant U-factor of 2 W/m^2-K, and visible transmissivity of 0.75. Outdoor air was introduced into the office at a constant rate of 10 l/s during the occupied hours (1.3 air changes per hour).

The energy consumption of an idealized air-based HVAC system was approximated by assuming a constant coefficient of performance (COP) of 3 for the cooling system and a constant heating efficiency of 100%. Cooling capacity was taken as 750 W and the heating capacity was taken 500 W. These values were chosen based on a preliminary sizing run. The energy consumed by fans and pumps was neglected. This object was controlled via an actuator placed in the EMS application of EnergyPlus; and influence of control approach was studied by using the control strategies described in Section 2.3.

The office was assumed to be occupied by one person from 9:00am to 5:00pm during the weekdays. The occupant was assumed to generate a heat gain of 100 W (with a sensible fraction of 0.6). Lighting and equipment power densities were taken as 10 W.m^{-2} (with a radiant fraction of 0.5) and 5.4 W.m^{-2} (with a radiant fraction of 0.3), respectively (ASHRAE, 2010). Interior-mounted roller blinds were considered in the model. These were perfect diffusers with optical properties that were independent of angle of incidence (DOE, 2010) (transmittance of 0.05 and reflectance of 0.8). In the simulations the blinds were discretized as a five state manually operated object (i.e. fully retracted, quarter-deployed, half-deployed, quarter-retracted, and fully-deployed). Manual control of the interior roller blinds were modeled as described in Haldi and Robinson (2010) and as summarized in Haldi and Robinson (2011). These logistic regression models representing the decision making process of the occupants are also shown in Figure 2. These stochastic control actions were included to better represent the uncertainty propagating to the controller models due to occupant behaviours.

The random seed was kept constant to ensure a repeatable sequence of control actions between different runs. Lighting fixtures dimmable and photosensor/occupancy sensor automated to meet a setpoint of 500 lux in the centre of the workplane (0.8 m above floor).

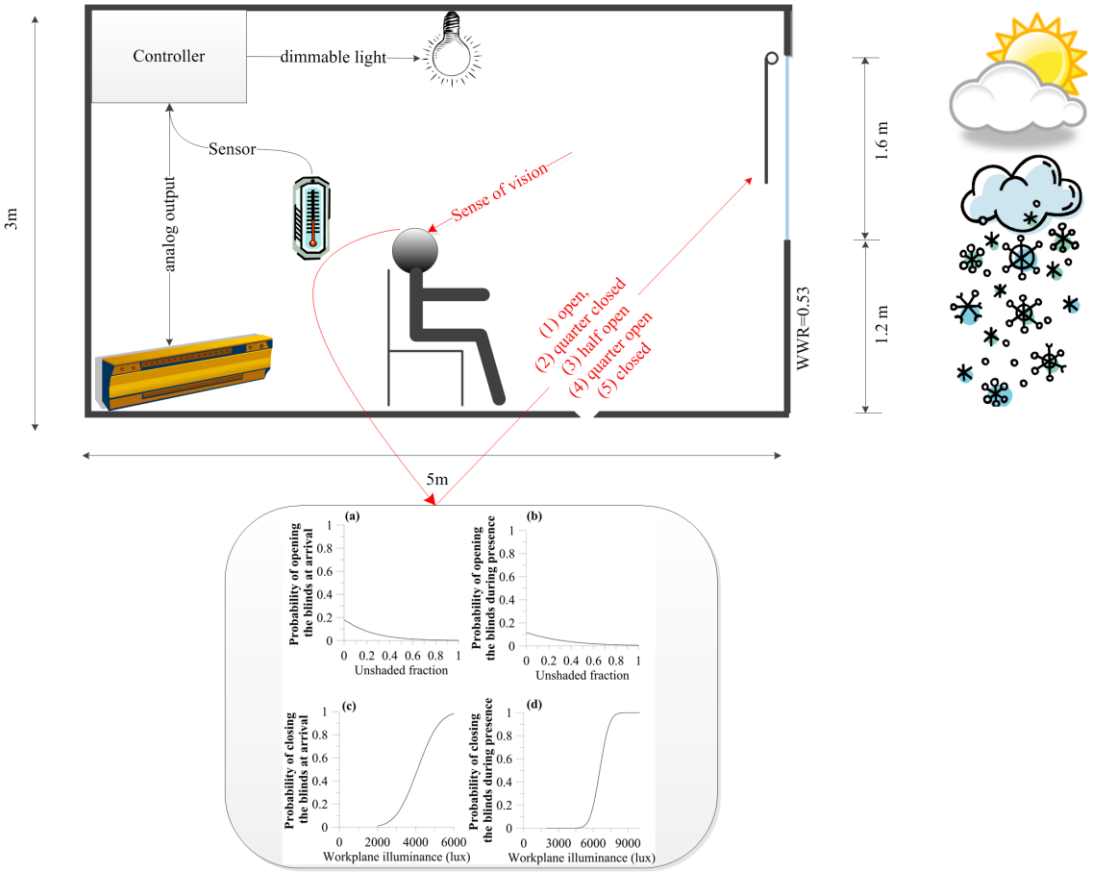

Figure 2: Building and occupant models.

2.3 Controller Models

2.3.1 Reactive controller model

The reactive controller model used in this study consists of a proportional controller with a heating setpoint of 20 °C, a cooling setpoint of 25 °C, and a throttling range of 1 °C. As shown in Figure 3, this indicates that below 20.5 °C, the ideal HVAC system initiates to provide heating. The heating input increases proportionally as the difference between the setpoint and control point increases. Below 19.5 °C, the heating system is actuated at its full capacity. Similarly, the cooling system turns on above 24.5 °C and proportionally reaches to its full capacity, if the temperature increases above 25.5 °C.

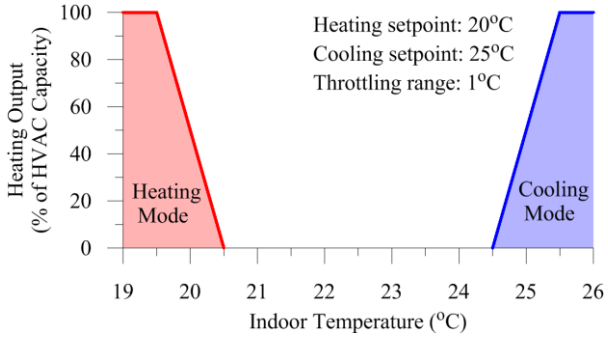

Figure 3: Reactive controller model.

2.3.2 Predictive controller model

The predictive controller model involves a learning agent developed using a particle filtering algorithm. Here, ten ensembles of parameters (time-constant and capacitance — denoted as τ and C in Figure 4) and state (temperature) were randomly sampled from the earlier (a priori) likelihood distribution. These ensembles were later inputted to a 1R1C explicit finite-difference model to forecast the state (T^{f}) in the following time-step (i.e. 5 minutes ahead). Concurrently, the state was measured — denoted as mean T^{m} and variance R in Figure 4. Based on a comparison between the measurement and the forecasts made based on each of the ensembles, the ensembles were assigned likelihood weights — denoted as W in Figure 4. The ensembles that yield forecasts closer to the measurements were assigned higher likelihood weights; while the ensembles that yield forecasts further away from the measurements were assigned lower likelihood weights. The normalized weights for the ensembles represent a posteriori distribution. At the end of each time step, ten new parameter ensembles were resampled from this new distribution forming a sharper and sharper estimate for the parameters (time-constant and capacitance). Application of the algorithm in the current study is summarized in Figure 4; for further information on particle filtering is available elsewhere (Doucet et al., 2000, Arulampalam et al., 2002).

This controller model also involves a predicting agent which is the state-space representation of a 1R1C model. Here, the temperature response of the control zone — for a forecast window of six hours — was predicted in absence of a control input (i.e., free-running response of the zone). This response array is denoted as $T_{k,i}$ in Figure 4, where subscript k represents the time at which the forecast was made and subscript i — zero to six hours with one hour increments — represents the time relative to k.

This controller model also contains an adapting agent which seeks for the optimal control sequence by employing the Newton's method in optimization. The square coefficient matrix for the control sequence array ΔU was denoted as Φ in Figure 4. Also, the tuning parameter (r in Figure 4) represents balance between the setpoint (SP in Figure 4) and the magnitude of the control signal. For example, for the case that r equals zero, the optimization goal is interpreted as the situation where we would not pay any attention to how large ΔU can be. Here, the tuning parameter was selected as 0.5. Then, Φ and r were used to compute the Hessian matrix — denoted as H in Figure 4. The optimal control sequence was then multiplied by a gain of 500 W/K for heating and -750 W/K for cooling loads. Further information about the optimization algorithm employed in the current study is available elsewhere (Wang, 2009).

3 Results

3.1 Reactive Controller

For the run period —from January 1st to October 31st 2013 —, simulated results indicate that the heating energy use per unit floor area was 28.8 kWh/m^2 and the cooling energy use was estimated as 19.3 kWh/m^2. This controller failed to keep the indoor climate within the designated setpoints for 457 ˚C.hours (excluding temperature deviations smaller than ±0.5). The cost of electricity for heating and cooling were estimated by assuming a hypothetical billing scheme, i.e. $10/kW for the monthly incremental peak due to heating/cooling energy use and $0.1/kWh due to the incremental energy use for heating/cooling. This hypothetical billing scheme, despite not being a true representation for a particular utility billing scheme, was taken to capture the relative variations between peak and usage-induced costs in different control strategies (Nghiem and Pappas, 2011). With this reactive controller, the incremental utility cost due to heating and cooling was estimated as $12.5/m^2. Of which, $4.8/m^2 was incurred due to electricity charges and $7.7/m^2 was incurred due to the demand charges.

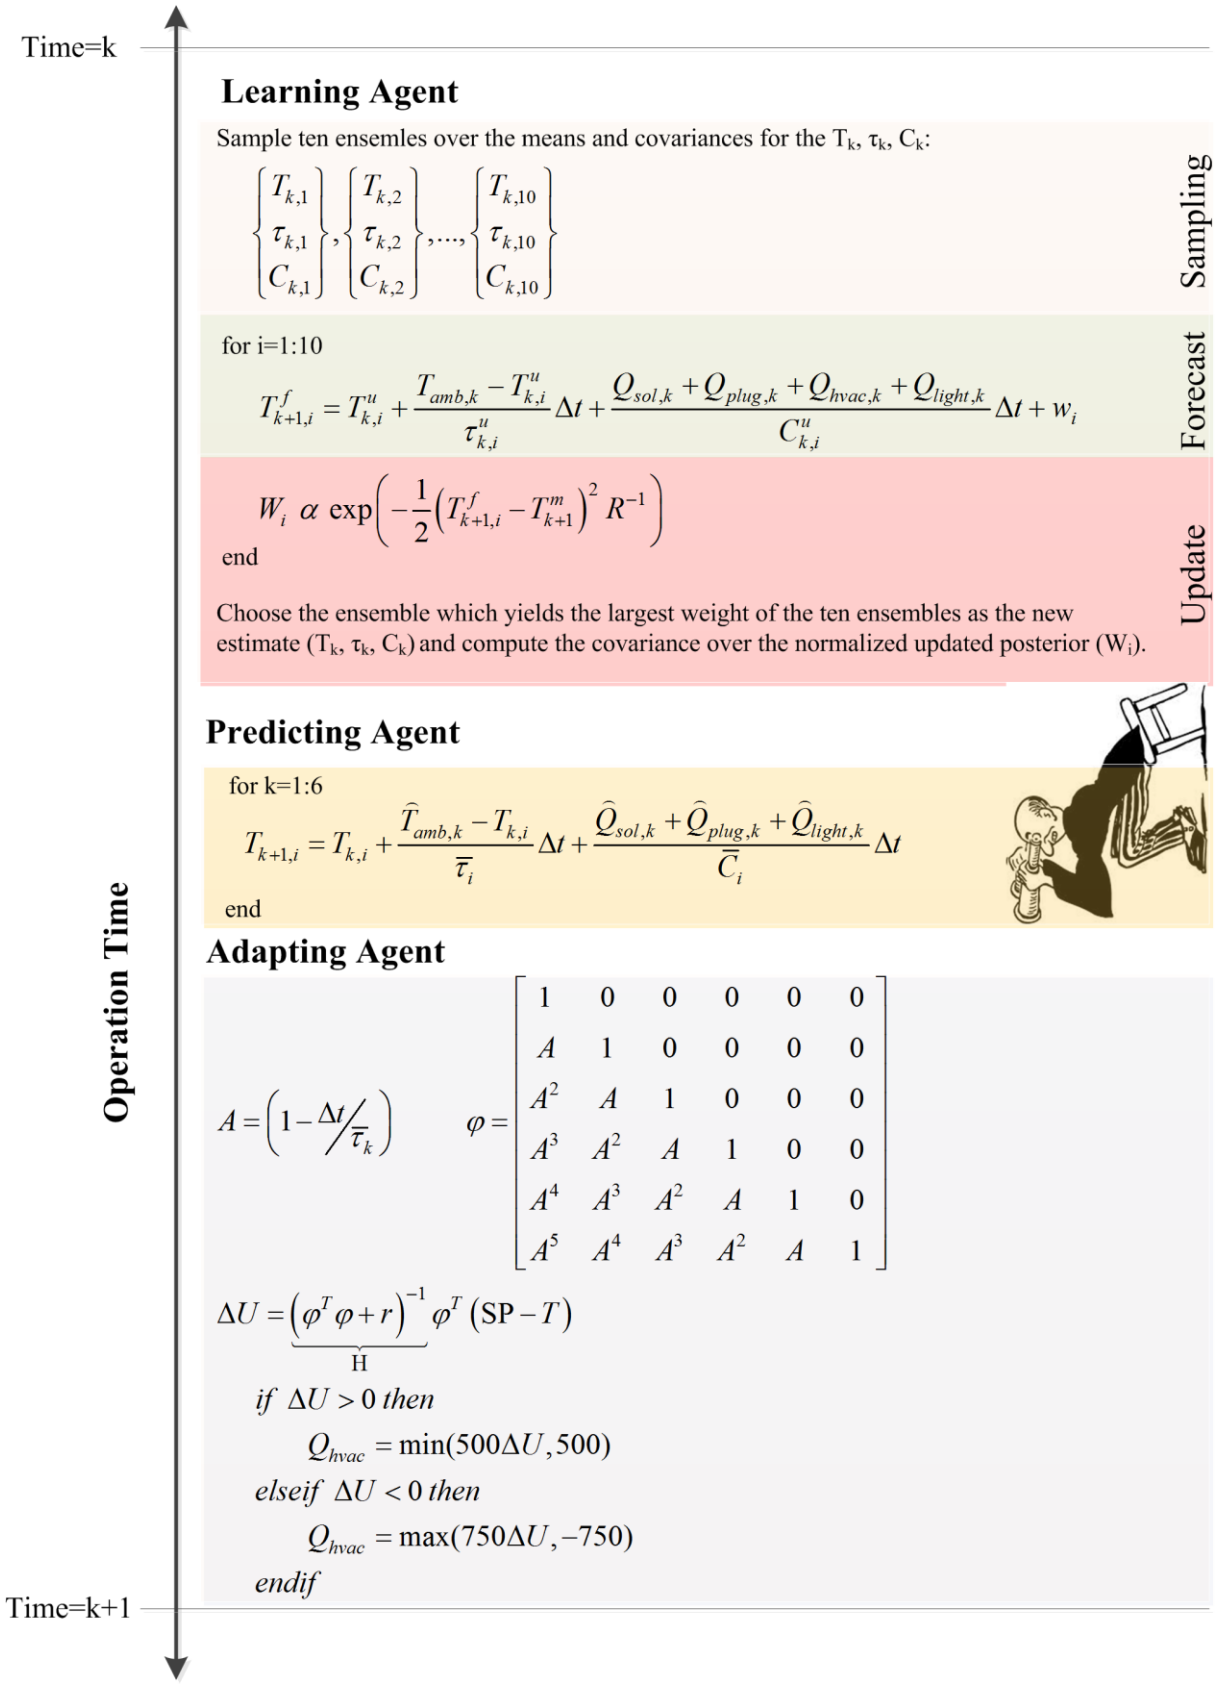

Figure 4: Predictive controller model

3.2 Predictive Controller

3.2.1 Learning Accuracy

Figure 5 shows the performance of the learning agent to back out the effective time-constant and effective thermal capacitance sequentially. The contour map represents the mean-squared errors resulted in while EnergyPlus building model results were fitted via the 1R1C controller model. A wide-range of parameter couples selected through a parameter sweep with fine incremental steps. This contour map was utilized to analyze the performance of the learning agent as whether or not it was capable to converge to the optimal region (enclosed by the smallest error contour line). The scatter plots represent the parameter couples selected by the learning agent at each time-step. The parameter couples were initialized at the origin of the plots (time-constant of 10 hours and capacitance of 0.1 MJ). In the end of the run periods, run Cases 1L, 2L, and 3L (see Figure 5 a to c) achieved convergence to the region enclosed by the smallest error contour line; while the run Case 4L failed to converge to a stable and accurate estimate for the parameters. This can be interpreted as the performance of the learning agent — in this hypothetical south facing perimeter office — is strongly tied with the availability of the sensors that measure/estimate the transmitted solar radiation. It is worth noting that the goal of estimating the solar gains in practice points out the necessity of having photosensors (or ideally pyranometers) and the importance of real-time feedback from blinds' positions. The rest of the paper accepts that a sensor and meter network exist to support the learning model in run case 1L and focuses on the predicting and adapting challenges and limitations.

3.2.2 Predictive Accuracy

Accuracy of the predicting agent to forecast the temperature response of the free-running office temperature an hour in advance in run cases 1P, 2P, and 3P is shown in Figure 6. Thus, the predictions for the natural response of the building were made in each time-step; however the adapting agent was not invoked to actuate the simulated-heating/cooling system. It is worth noting that all run cases resulted in identical 1-hour predictions. This can be explained as the simulated controller model — discrete-time state-space model — use the real-time measurements (not forecasts) to predict the temperature response an hour ahead. Therefore, results shown in Figure 6 indicate that regardless of the resolution/availability of forecasts (transmitted solar radiation, lighting or equipment energy use for the next six hours), the predicting agent can estimate one hour ahead response accurately (variance of the error spread is $1.1^{\circ}C^2$). This underlines an extremely crucial point in the application of MPC, because the adapting agent assigns reducing weights to the predictions further away in time. For example, for a control zone with a 30 hours of time-constant, the adaptive agent makes 74% of its decision based on the one hour prediction —and only depends on the previously received inputs. Second to sixth hour predictions receive diminishing weights in decision making; i.e. 19, 5.3, 1.3, 0.35, 0.05% respectively.

Figure 7 illustrates a comparison of the second hour and fourth hour predictions made in run cases 1P, 2P, and 3P. Results indicate that the variance of the error spread increases from 2.6 to $3.4\ ^{\circ}C^2$ (mean-squared error of 1.6 to $1.8\ ^{\circ}C$) for two hour and from 6.7 to $9.6\ ^{\circ}C^2$ (mean-squared error of 2.6 to $3.9\ ^{\circ}C$) for four hour predictions, when the idealized forecast files used in 1P were replaced with the NAM based forecast files used in 3P. The uncertainty introduced by several factors (non-ideal weather forecasts, as well as lack of knowledge on future interior shade positions or the state of the lighting equipment) introduce additional challenges for prediction. Considering the aforementioned diminishing roles of these predictions further away in time in controls decision making and the increase in error spread due to limitations related with the availability/resolution of forecasts, increasing prediction time ho-

rizon beyond a few hours may become both computationally tedious and may detrimentally influence control decisions.

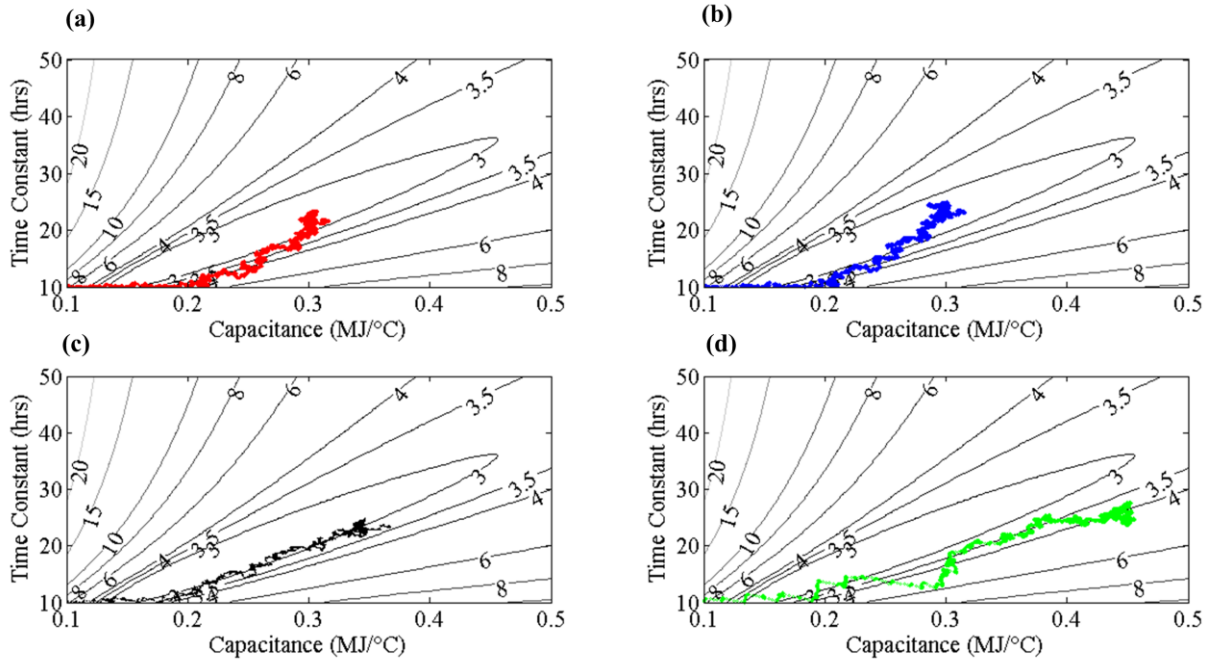

Figure 5: Performance of the learning agent in converging to the optimal region for run Cases (a) L1, (b) L2, (c) L3, (d) L4.

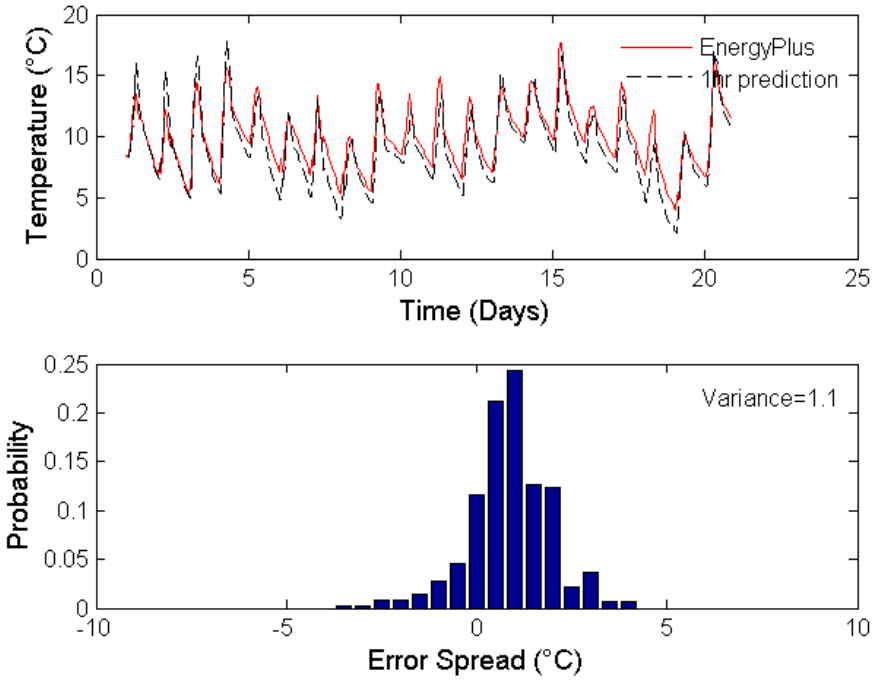

Figure 6: Accuracy of the predicting agent to predict the temperature response 1-hour in advance in run cases 1P, 2P, and 3P.

3.3 Performance of the predictive controller

Based on the predictions made by the predicting agent, the adaptive agent was invoked to actuate the simulated heating and cooling system. Figure 8 illustrates a comparison of the reactive system and the predictive system during a representative day (January 9th). Results indicate that the reactive system overshoots the heating setpoint causing excess in heating energy use (2.7 kWh). On the contrary, the predictive system (with ideal forecasts in run Case 1P) met the heating setpoint without overshooting and successfully reduced the heating energy use to 2.4 kWh for the representative day. To this end, predictive controller inputted variables such as outdoor temperature and transmitted solar radiation. In reality, the predictive controller cannot easily find forecasts for the future internal shade positions and, thus, it would rather use the weather forecasts as inputs. The influence of such uncertainties on the theoretical performance of the MPC was studied in run cases 2P and 3P.

Figure 9 demonstrates a comparison between the reactive run case and the predictive run cases P1, P2, and P3 in terms of the heating/cooling energy use and the incremental utility costs due to the heating/cooling energy use. Results suggest that the proposed predictive system in the hypothetical south-facing perimeter office can theoretically reduce the heating/cooling induced energy use by 13% and heating/cooling induced utility costs by 10%. Moreover, with this theoretical system ˚C.hours during which setpoints were not met reduced by about 35%. However, this theoretical potential was challenged with the limited nature of the sensors/meters to learn the physics of the control zone and the availability/resolution of forecasts to predict the response of the control zone. In Case P3, when the ideal forecasts were replaced with the NAM based weather forecasts, the simulated-energy savings reduced by 5% and the utility cost savings reduced by 2% from the theoretical maximum (Case P1). Furthermore, reduction in the ˚C.hours during which setpoints were not met was realized as 392 in Case P3 as opposed to 306 in Case P1. It is worth noting that despite these reductions in the improvements, all predictive run cases resulted in significant savings in energy and utility costs and showed promising enhancements in the controllability of the indoor climate within the designated setpoints. This suggests that the self-adaptive MPC — to better use the thermal mass of building fabric — is a robust method, even with the limitations/challenges discussed in the current paper. This was attributed to the fact that the optimal control decisions made by the adapting agent relies overly on the first a few hours. This short period of time can lend itself for the linearization of the input states, as long as the existing sensor and meter resolution are adequate.

4 Conclusions

A self-adaptive MPC method is put forward. The availability and resolution of the sensors, meters, and forecasts for this method to undertake optimal operational decisions were systematically varied. For each of these variations, this method was contrasted to a classical reactive controller. It was noted that the learning abilities of the controller for this hypothetical south-facing perimeter office space is strongly dependent on the existence of sensors to estimate the amount of solar gains. It was also noted that the capabilities of a first order controller model to predict the temperature response two to six hours ahead is restricted with the availability and the resolution of the forecasts. However, considering the diminishing roles of these predictions further away in time in controls decision making and the increase in error spread due to limitations related with the availability/resolution of forecasts, increasing prediction horizon may become both computationally tedious and detrimental in energy efficiency. Thus, a future work is planned to study the influence of length of the prediction horizon on the performance of the MPC method for buildings of varying thermal mass and fenestration. In summary, the self-adaptive MPC method —based on the results of the current

simulation study— resembles a robust approach, even with the discussed limitations/challenges related with the availability/resolution of the sensors, meters, and forecasts.

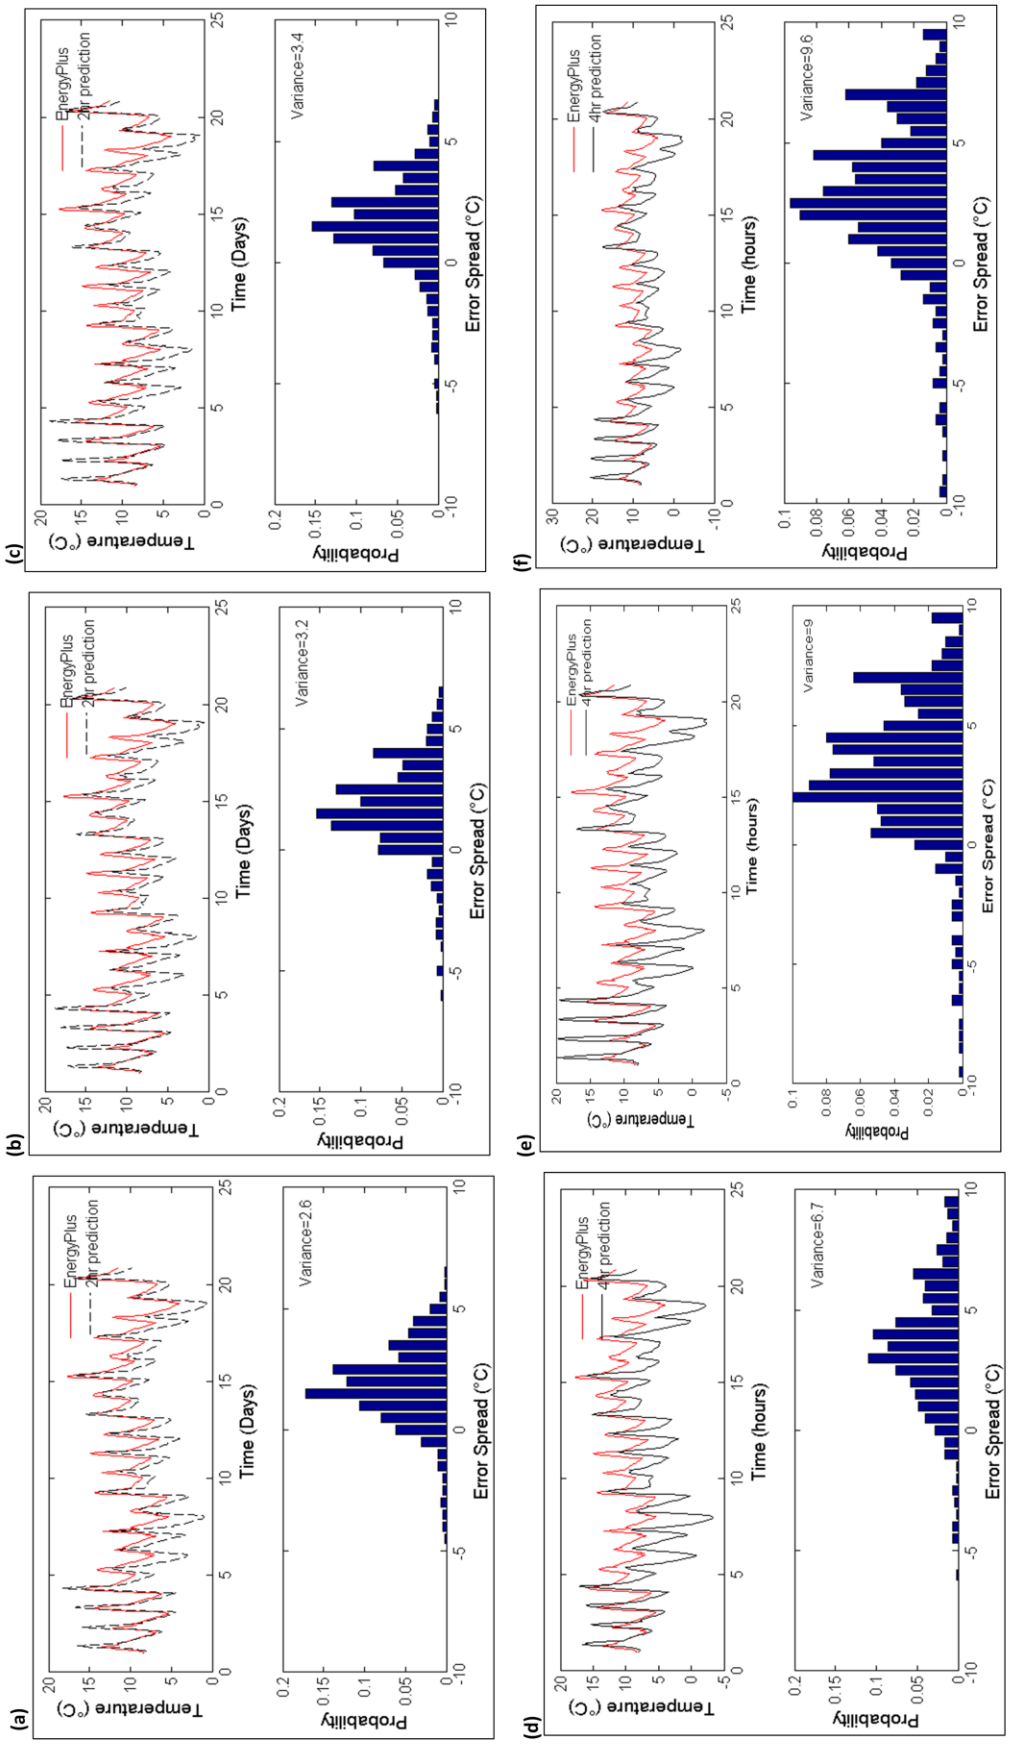

Figure 7: 2hr predictions and error spread in temperature response for the run cases —described in Table 2— (a) 1P (b) 2P (c) 3P and 4hr predictions and error spread for run cases (d) 1P (e) 2P (f) 3P.

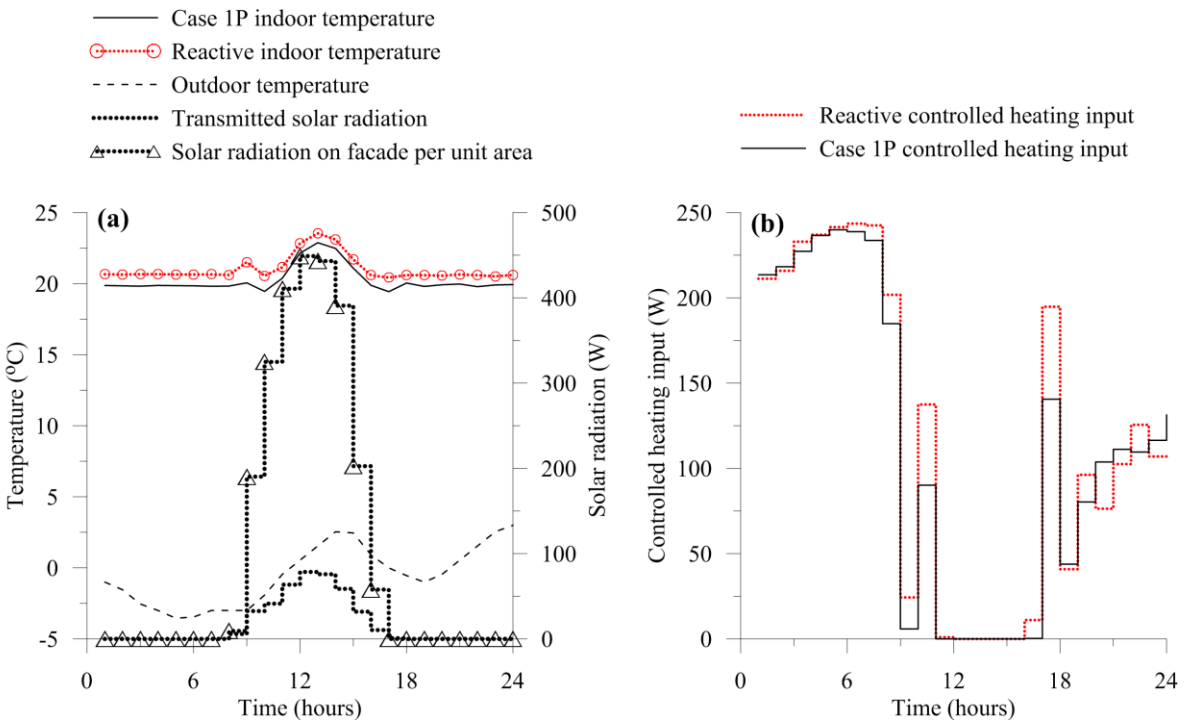

Figure 8: (a) Major disturbances in the controlled system and the indoor temperatures in a sunny winter day and **(b)** the comparison between the reactive and predictive control outputs.

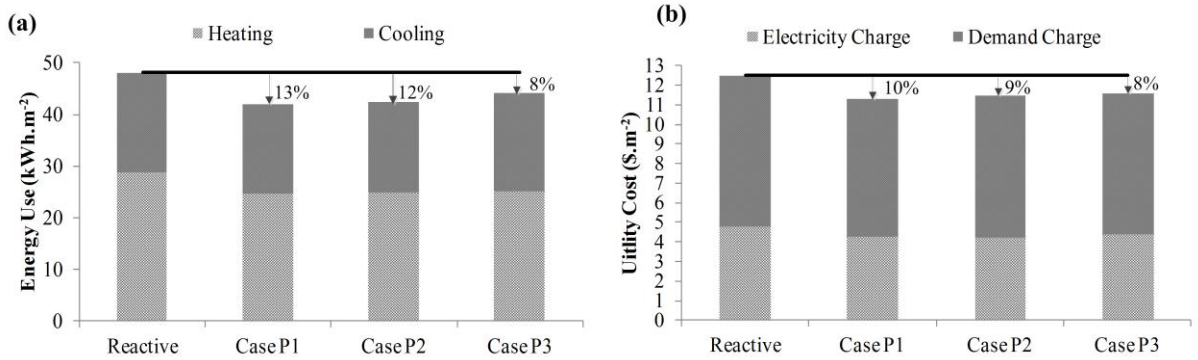

Figure 9: Comparison of the reactive run case and predictive run Cases P1, P2, and P3 in terms of (a) the heating and cooling energy use and (b) the electricity and demand charge per unit area.

Acknowledgements

This research was funded by a research grant provided by Natural Sciences and Engineering Research Council (NSERC) of Canada which is gratefully acknowledged.

References

Arulampalam, M. S., Maskell, S., Gordon, N. & Clapp, T. 2002. A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking. *Signal Processing, IEEE Transactions on,* 50**,** 174-188.

Ashrae 2010. Standard 90.1. Atlanta: American Society for Heating Refrigeration and Air-conditioning.

Candanedo, J. A. & Athienitis, A. K. 2010. Simplified Linear Models for Predictive Control of Advanced Solar Homes with Passive and Active Thermal Storage.

Candanedo, J. A., Dehkordi, V. R. & Lopez, P. A control-oriented simplified building modelling strategy. 13th Conference of international Building Performance Simulation Association, 2013 France. IBPSA.

Chen, T. Y. & Athienitis, A. K. 2003. Investigation of practical issues in building thermal parameter estimation. *Building and Environment,* 38**,** 1027-1038.

Doe, U. 2010. EnergyPlus Engineering Reference. *The Reference to EnergyPlus Calculations*.

Doucet, A., Godsill, S. & Andrieu, C. 2000. On sequential Monte Carlo sampling methods for Bayesian filtering. *Statistics and Computing,* 10**,** 197-208.

Fraisse, G., Viardot, C., Lafabrie, O. & Achard, G. 2002. Development of a simplified and accurate building model based on electrical analogy. *Energy and Buildings,* 34**,** 1017-1031.

Fux, S. F., Ashouri, A., Benz, M. J. & Guzzella, L. 2014. EKF based self-adaptive thermal model for a passive house. *Energy and Buildings,* 68, Part C**,** 811-817.

Haldi, F. & Robinson, D. 2010. Adaptive actions on shading devices in response to local visual stimuli. *Journal of Building Performance Simulation,* 3**,** 135-153.

Haldi, F. & Robinson, D. 2011. The impact of occupants' behaviour on building energy demand. *Journal of Building Performance Simulation,* 4**,** 323-338.

Henze, G. P., Kalz, D. E., Liu, S. & Felsmann, C. 2005. Experimental Analysis of Model-Based Predictive Optimal Control for Active and Passive Building Thermal Storage Inventory. *HVAC&R Research,* 11**,** 189-213.

Liu, Y., Pan, Y. & Huang, Z. 2013. Simulation-based receding-horizon supervisory control of HVAC system. *IBPSA 2013.* France: IBPSA.

Maasoumy, M., Moridian, B., Razmara, M., Shahbakhti, M. & Sangiovanni-Vincentelli, A. Online Simultaneous State Estimation and Parameter Adaptation for Building Predictive Control. 2012. Dynamic System and Control Conference (DSCC 2013), Stanford, CA.

Nghiem, T. X. & Pappas, G. J. Receding-horizon supervisory control of green buildings. American Control Conference (ACC), 2011, 2011. IEEE, 4416-4421.

Oldewurtel, F., Parisio, A., Jones, C. N., Gyalistras, D., Gwerder, M., Stauch, V., Lehmann, B. & Morari, M. 2012. Use of model predictive control and weather forecasts for energy efficient building climate control. *Energy and Buildings,* 45**,** 15-27.

Prívara, S., Cigler, J., Váňa, Z., Oldewurtel, F., Sagerschnig, C. & Žáčeková, E. 2013. Building modeling as a crucial part for building predictive control. *Energy and Buildings,* 56**,** 8-22.

Radecki, P. & Hencey, B. Online building thermal parameter estimation via unscented kalman filtering. American Control Conference (ACC), 2012, 2012. IEEE, 3056-3062.

Sourbron, M., Verhelst, C. & Helsen, L. 2012. Building models for model predictive control of office buildings with concrete core activation. *Journal of Building Performance Simulation,* 6**,** 175-198.

Wang, L. 2009. *Model Predictive Control System Design and Implementation Using MATLAB®*, Springer.

Wang, S. & Ma, Z. 2008. Supervisory and Optimal Control of Building HVAC Systems: A Review. *HVAC&R Research,* 14**,** 3-32.