

Stochastic Model Predictive Controller for the Integration of Building Use and Temperature Regulation

Alie El-Din Mady¹, Gregory M. Provan¹, Conor Ryan² and Kenneth N. Brown²

¹Cork Complex Systems Lab and ²Cork Constraint Computation Centre,
Department of Computer Science, University College Cork, Ireland.

Abstract

The aim of a modern Building Automation System (BAS) is to enhance interactive control strategies for energy efficiency and user comfort. In this context, we develop a novel control algorithm that uses a stochastic building occupancy model to improve mean energy efficiency while minimizing expected discomfort. We compare by simulation our Stochastic Model Predictive Control (SMPC) strategy to the standard heating control method to empirically demonstrate a 4.3% reduction in energy use and 38.3% reduction in expected discomfort.

Introduction

The industrial and residential building sectors consume around 40% of total energy use in industrial societies, and account for nearly one-third of greenhouse gas emissions. Of that, approximately one-third can be attributed to the Heating, Ventilation and Air-Conditioning (HVAC) systems. There is now a significant effort being devoted to reducing these energy costs. Many buildings incorporate a Building Management System (BMS) to maintain a comfortable environment in an energy-efficient manner. A key task of HVAC control is to minimize occupant discomfort due to inappropriate indoor temperature and humidity. A typical BMS would provide a core functionality that keeps the building's climate within a specified range, automates the lighting and HVAC based on fixed schedules, and monitors system performance and device failures. There are many sources of energy inefficiency in BMSs, some of which arise due to the use of simple control algorithms (e.g., On/Off, P or PI controllers). One main inefficiency is the widespread use of schedule-based control, where buildings will be heated to 20°C during office hours, regardless of the actual occupancy of the building.

An intuitive approach to obtaining a desired temperature when it is needed is to anticipate the requirement and turn on the heating system in advance. This raises the following questions: (a) what is the optimal moment to trigger the system, and (b) what is the required actuation to reach the temperature reference? The optimal period of preheating can vary between minutes and several hours, depending on the

outdoor temperature, the desired internal temperature, and the building physics. This is further complicated by uncertainty in the time at which the temperature is required. In this paper we introduce a more sophisticated HVAC control approach that incorporates a Model-Predictive Control (MPC) algorithm augmented with stochastic occupancy information, which we assume is provided by a separate building information model. We compare this algorithm with a standard schedule-based controller, and show that it decreases user discomfort by an average of 38.3% while decreasing energy usage by 4.3%. MPC (Garcia, Prett, and Morari 1989) is a control algorithm that computes the current control action by solving, at each sampling instant, a finite-horizon open-loop optimal control problem, using the current state of the plant as the initial state. The optimization process generates an optimal control sequence, the first control value of which is applied to the plant.

Our contributions are as follows:

- we formulate a building HVAC control system as an optimal control problem where we minimize energy use and expected user discomfort, using occupancy predictions as inputs for control;
- we empirically compare a baseline approach of schedule-based control to solving the optimal control problem using fixed schedule MPC and stochastic MPC;
- we show empirically that our MPC approach improves user comfort by an average of 38.3% while decreasing energy usage by 4.3%.

Related Work

BMSs typically incorporate simple control mechanisms in practical applications (Chandan, Mishra, and Alleyne 2010; Kolokotsa et al. 2009). We study the use of MPC in such applications. There have been a few approaches that have used MPC for building systems, such as (Oldewurtel et al. 2010; Morosan et al. 2010; Chandan, Mishra, and Alleyne 2010; Siroky et al. 2011). The novelty of our approach in comparison to this prior work is that we incorporate stochastic occupancy models within the control loop, whereas prior work assumes fixed deterministic schedules.

The concept of predicting occupant movements to help improve building efficiency is now being addressed. Hoes et

al. (2009) demonstrate that human activities can have a significant effect on efficiency, particularly in passive buildings, although this work is focused on the design stage rather than on control. Dodier et al. (2006) consider Bayesian Networks for predicting occupant behaviour based on sensor data. Perhaps the most comprehensive model is by (Page et al. 2008), based on Markov chains, with added features to account for divergences from normal behaviour. However, few of these papers consider the link back to active control systems, nor demonstrate the tradeoff between increased energy savings and enhanced occupant comfort.

Application Domain

HVAC systems provide a specified ambient environment for occupants with comfortable temperature, humidity, etc. Several control strategies have been introduced to control the temperature regulation, where an operating scheduling is pre-defined. In this context, standard PI control algorithms are adequate for the control of HVAC processes (Donnis and Caraiscos 2009), so it is considered as one of our baseline control algorithms. In addition to the PI-Controller, we have considered an MPC controller that uses a fixed operating scheduler as an additional baseline in order to show the impact of integrating the stochastic occupancy model.

In our approach, we integrate a heating controller with a stochastic building occupancy model that uses predicted occupant profiles for improving HVAC control. Fig. 1 compares the room temperature generated from three controllers, which use PI, MPC and stochastic MPC (SMPC) algorithms, given a room that is occupied between 0800 and 1200 and again between 1300 and 1700, and is empty between 1200 and 1300. The figure shows that the PI approach overshoots the 20° C setpoint S , and fails to converge totally; the MPC algorithm reaches an equilibrium near S quickly, but does not account for the absence of the occupant; the SMPC algorithm reaches equilibrium quickly, but also can save energy by reducing the room temperature when the occupant is absent for lunch. This plot shows the potential for the SMPC approach to save energy relative to the other algorithms through its fast convergence and energy savings during periods of no occupancy.

We describe in this section the evaluation metrics for the control effectiveness and the design specifications for the model used to evaluate the developed controller.

Performance Metrics

We regulate the temperature (for maintaining thermal comfort) with respect to two performance metrics: energy consumption and user discomfort.

- The energy consumed for indoor heating, in kWh:

$$M_E = \sum_{k=k_s}^{k_f} u(k) Pwr_{max} T_s \quad (1)$$

Where Pwr_{max} is the maximum power consumed by the actuator, $u(k)$ is the actuation variable at time-step k (where k varies from k_s to k_f), and T_s is the sampling period.

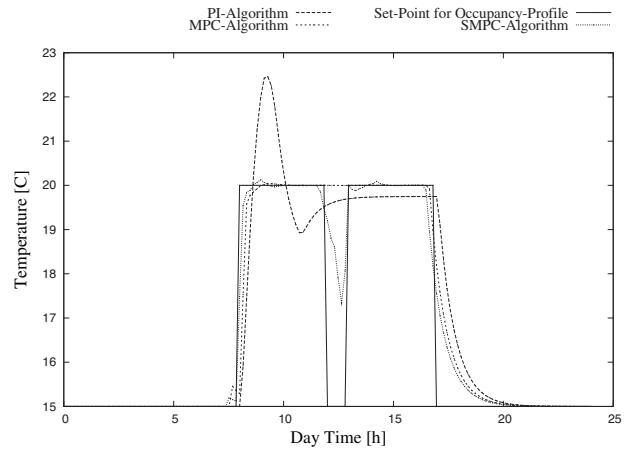


Figure 1: Comparison of PI, MPC and SMPC algorithms in controlling a room to a setpoint of 20° C

- The expected discomfort metric penalizes the difference between the measured indoor air temperature $y(k)$ and the reference temperature $r(k)$ weighted by the probability $p(k)$ of occupancy. We use the weighted Mean Square Error (wMSE) to reflect the indoor temperature variation around the reference temperature in $^{\circ}C^2$:

$$M_C = \sum_{k=k_s}^{k_f} \frac{(y(k) - r(k))^2}{k_f} p(k) \quad (2)$$

Design Specification

In order to demonstrate our control approach, we have modelled a single room that contains two main groups of components, as shown in Fig. 2: (a) environment components reflect the plant physics (e.g., walls, windows); (b) control components show the control/sensing algorithm that used to modify/monitor the environment. Hybrid Systems (HS) (Heninger 1996) are used to create our model, where both discrete (e.g. presence detection) and continuous (e.g. heat dissipation) dynamics are represented. Here, the continuous dynamic is represented using differential-equation and hence the Linear Hybrid Automata (LHA) become a suitable HS candidate for this model.

Environment Models Three variables are identified to evaluate the model behaviour: external temperature T_e , set-point temperature T_{sp} , and consistency parameter h . Moreover, four environment components have been used; Wall, Window, Radiator and Indoor Air model as follows (for more details, we refer the reader to (Mady, Boubekour, and Provan 2009)):

1. *Wall Model*: One of the room walls is facing the building façade, which implies heat exchanges between the outdoor and indoor environments. In general, a wall can be modelled using several layers, where greater fidelity is obtained with increased layers in the wall model. In our case, four layers have been considered to reflect sufficient fidelity (Yu and van Paassen 2004), using the following differential equation; Eq. 3:

$$\rho_{wall} V_{wall} c_{wall} \frac{dT_{wall}}{dt} = \alpha_{wall} A_{wall} (T_e - T_{wall}) \quad (3)$$

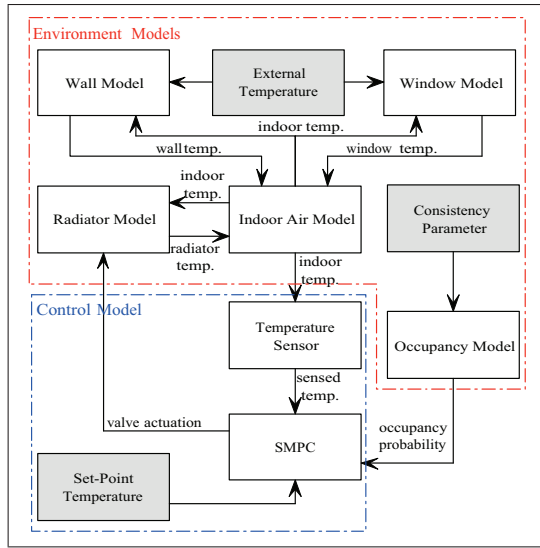


Figure 2: Single Room Model

Where, ρ_{wall} is wall density [kg/m^3], V_{wall} is wall geometric volume [m^3], c_{wall} is wall specific heat capacity [$J/kg.K$], T_{wall} is wall temperature [$^{\circ}C$], α_{wall} is wall thermal conductance [$W/(m^2.K)$], A_{wall} is wall geometric area [m^2] and T_e is outdoor temperature [$^{\circ}C$].

2. **Radiator Model:** The radiator model uses the temperature difference between the water-in and water-out in order to heat the room. In this case, the temperature is controlled through the radiator water flow using the valve occlusion (called actuation variable u). Moreover the radiator exchanges temperature with its environment. Here we assume that the radiator is fixed on a wall that does not exchange temperature and hence it has negligible effect on the radiator, therefore the indoor air is the only effective component on the radiator as shown in equations: Eq. 4 and Eq. 5.

$$M_{wtr}c_{wtr} \frac{dT_{rad}}{dt} = m_{wtr}c_{wtr}(T_{wtrin} - T_{wtrout}) - Q \quad (4)$$

$$Q = Q_{air} = \alpha_{air}A_{rad}(T_{rad} - T_{air}) \quad (5)$$

Where, M_{wtr} is water mass [kg] and m_{wtr} is water mass flow rate through the radiator valve [kg/s].

3. **Indoor Air Model:** In order to model the indoor temperature T_{air} propagation, all HVAC components have to be considered as they exchange heat with the air inside the controlled room following equations 6 through 9.

$$\rho_{air}V_{air}c_{air} \frac{dT_{air}}{dt} = Q_{air} + Q_{wall} + Q_{window} \quad (6)$$

$$Q_{wall} = \alpha_{air}A_{wall}(T_{wall} - T_{air}) \quad (7)$$

$$Q_{air} = \alpha_{air}A_{radiator}(T_{radiator} - T_{air}) \quad (8)$$

$$Q_{window} = \alpha_{air}A_{window}(T_{window} - T_{air}) \quad (9)$$

4. **Window Model:** A window has been modelled to calculate the effects of glass on the indoor environment. Since the glass capacity is very small, the window has been modelled as an algebraic equation, Eq. 10, that calculates the heat transfer at the window node.

$$\alpha_{air}(T_e - T_{window}) + \alpha_{air}(T_{air} - T_{window}) = 0 \quad (10)$$

Occupancy Model The occupancy model assumes that the probability of occupancy at any given time slot is conditional on the state of the occupant in the previous time slot. This probability will be different for different times of day, to model arrival, departure, lunch breaks, etc.. This suggests a time-inhomogeneous markov chain over two states $\{in, out\}$. However, we assume the occupant's behaviour is more subtle. Most days, the occupant will conform to a standard pattern with high probability. But occasionally the standard pattern is broken, due to visitors, meetings or external events. Averaging over all of those cases, we have a different probability of occupancy, closer to 0.5, but still conditional on the previous occupancy value. We therefore introduce n consistency states $\{1, 2, \dots, n\}$, representing an estimate of how closely the behaviour is conforming to the standard pattern. This gives a total of $2n$ states for the markov chain. There are $(k_f - k_s)$ time slots. For each time slot k , we have a vector P_k of $2n$ probabilities that the occupant is *in* or *out* in a particular consistency state. The probability of the occupant being in at time k , as used by the controller, is

$$p(k) = P(O_k = in) = \sum_{i=1}^n P(X_k = (in, i)) \quad (11)$$

We have a $2n \times 2n$ matrix M_k , representing the transition probabilities from time k to $k + 1$, where each element is of the form $P(X_{k+1} = (occ, c) | X_k = (occ', c'))$. P_0 is a fixed vector of probabilities for the initial occupancy and consistency states. The Markov chain is defined by the relations $P_{k+1} = P_k * M_k$. At any time k , if we are given a definite occupancy value and consistency value h , we revise the matrix P_k so that that entry is set to 1, and then we revise our computation of P_{k+1} . To simulate an occupancy sequence for a single day, we generate one sample from the Markov chain. We first sample from P_0 , send it to the controller, and obtain a value for X_0 . We then update P_1 , send it to the controller, sample it, and obtain a value for X_1 , and so on. The transition probabilities are set so that high initial values of the consistency parameter tend to produce an initial set of probabilities which are close to 1.0 or to 0.0, which represents a prediction with little uncertainty, while low initial values produce a set of probabilities which are close to 0.5, and thus indicate close to random behaviour.

Control Model A temperature sensor and SMPC are used to monitor and control the environment model. The temperature sensor samples the indoor temperature ($T_s=1$ min) and sends the sampled value to the controller. The SMPC receives the sampled temperature and the occupancy probability $p(k)$ generated from the occupancy model. Assuming that predicted input, and a control horizon no greater than an hour, the control model can use expected occupancy probabilities for the next hour, as computed from the current occupancy data.

Control Algorithm

For a given reference signal $r(k)$ at sample time k and within a prediction horizon N_p , the objective of the predictive control system J is to minimize the energy consumption and to bring the predicted output $Y = [y(k +$

$1|k) \dots y(k + N_p|k)]^T$ as close as possible to the reference signal R_s , where we assume that the reference signal remains constant in the optimization window, such that $R_s^T = [11 \dots 1]r(k)$, where $[11 \dots 1]$ is a $1 \times N_p$ identity vector. Our objective is to find the “best” actuation vector $\Delta U = [\Delta u(k) \dots \Delta u(k + N_c - 1)]^T$ that minimizes J .

The cost function J reflects the control objective described in Eq. 12, where the first term is linked to the objective of minimizing the errors between the predicted output and the reference signal based on the occupancy probability $P = [p(k + 1) \dots p(k + N_p)]$. The second term reflects the power used at each sample time started from $(k + 1)$ to $(k + N_c)$, where N_c is the control horizon. Finally, the last term considers dampening the actuation ΔU . \bar{R} is a diagonal matrix with the form $\bar{R} = r_w I_{N_c \times N_c}$, where r_w is used as a tuning parameter for the desired closed-loop performance. Limiting ΔU is standard in many control applications that because the true dynamics contain inherent non-linearities that are not expressed in the dynamics, however it is not reflected in the metrics as it is a secondary-goal.

$$J = P(R_s - Y)^T(R_s - Y) + (u(k) + \Delta U)^T(u(k) + \Delta U) + \bar{R}\Delta U^T\Delta U \quad (12)$$

The system model for the heating process can be described by a linear model, represented in a discrete state space as in Eq. 13 and Eq. 14:

$$x_m(k + 1) = A_m x_m(k) + B_m u(k) \quad (13)$$

$$y(k) = C_m x_m(k), \quad (14)$$

where $u \in \mathbb{R}$ is the actuation variable (valve actuation), $y \in \mathbb{R}$ is the process output (sampled indoor temperature), and $x_m \in \mathbb{R}^{n \times 1}$ is the state variable.

Due to the principle of receding horizon control, where a current information of the plant is required for prediction and control, we have assumed that the input $u(k)$ cannot affect the output $y(k)$ at the same time.

We denote the difference of the state variable by $\Delta x_m(k + 1) = x_m(k + 1) - x_m(k)$ and the difference of the actuation variable by $\Delta u(k) = u(k) - u(k - 1)$, therefore:

$$\begin{bmatrix} \overbrace{\Delta x_m(k + 1)}^{x(k+1)} \\ \overbrace{y(k + 1)}^{x(k+1)} \end{bmatrix} = \begin{bmatrix} \overbrace{A_m \quad o_m^T}^A \\ \overbrace{C_m A_m \quad 1}^A \end{bmatrix} \begin{bmatrix} \overbrace{\Delta x_m(k)}^{x(k)} \\ \overbrace{y(k)}^{x(k)} \end{bmatrix} + \begin{bmatrix} \overbrace{B_m}^B \\ \overbrace{C_m B_m}^B \end{bmatrix} \Delta u(k) \quad (15)$$

$$y(k) = \begin{bmatrix} \overbrace{o_m \quad 1}^C \end{bmatrix} \begin{bmatrix} \overbrace{\Delta x_m(k)}^{x(k)} \\ \overbrace{y(k)}^{x(k)} \end{bmatrix} \quad (16)$$

where $o_m = [00 \dots 0]$ and the triplet (A, B, C) is called the augmented model.

Therefore, a compact matrix form can be concluded from Eq. 15 and Eq. 16 as:

$$Y = Fx(k) + \Phi\Delta U \quad (17)$$

$$\text{Where, } F = \begin{bmatrix} CA \\ CA^2 \\ \vdots \\ CA^{N_p} \end{bmatrix}; \quad \Phi = \begin{bmatrix} CB & 0 & 0 & \dots & 0 \\ CAB & CB & 0 & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ CA^{N_p-1}B & CA^{N_p-2}B & CA^{N_p-3}B & \dots & CA^{N_p-N_c}B \end{bmatrix}$$

In order to evaluate the optimal ΔU that minimizes J , using Eq. 17, J is expressed as in Eq. 18.

$$J = P(R_s - Fx(k))^T(R_s - Fx(k)) + u^2(k) - 2\Delta U^T(P\Phi^T(R_s - Fx(k)) - u(k)) + \Delta U^T\Delta U(P\Phi^T\Phi + \bar{R} + 1) \quad (18)$$

We have used quadratic programming to optimize the objective function J , constraining ΔU to the interval $\Delta U \in [\Delta U^{max}, \Delta U^{min}]$, where ΔU^{max} and ΔU^{min} are the maximum and minimum ΔU variation, respectively.

Experimental Design

In our experimental design, we vary external temperature T_e , set-point temperature T_{sp} and consistency parameter h to evaluate the control efficiency, as measured using the M_E and M_C metrics. These experiments consider the cross product space $T_e \times T_{sp} \times h$, where some constraints are identified for the variables' search space to eliminate the physically unrealizable solutions. We consider two baseline models that use schedule-based strategies: (a) a PI-Controller is selected to show the improvement compared to a standard heating controller, and (b) MPC is used to show the improvement possible from integrating occupancy models. Note that the baseline algorithms need only the control variables T_e and T_{sp} . As an example, if we have input triple $(10, 23, 0)$ we will measure outputs $M_E = 90$ and $M_C = 5$.

Prior to running experiments, we tune the control parameters (N_p, N_c, r_w) in order to evaluate the “best” operating points (based on M_E, M_C) for the proposed control algorithm, where N_p and N_c are measured in minutes. Therefore, a schedule-base MPC algorithm has been considered with typical values for $T_e=10$ and $T_{sp}=23$. The search space for each parameter is identified as following: $N_p \in \{5, 10, \dots, 60\}$, $N_c \in \{5, 10, \dots, 60\}$, and $r_w \in \{0.1, 0.2, \dots, 0.5\}$, where $N_p \geq N_c$.

Using a pareto-frontier technique (Zitzler and Thiele 1998), we identify a set P_f of preferred operating points, where each $P \in P_f$ has lower energy consumption than all other points with a similar expected discomfort level. In order to select one operating point among the pareto points, we have assumed an expected discomfort tolerance equal to $1 \text{ } ^\circ C^2$. Using this constraint, we identify the point with minimum energy consumption as the “best” operating point ($N_p = 45, N_c = 30, r_w = 0.1$).

Experimental Results

Our main aim is to examine the tradeoff between energy use and expected occupant discomfort. We vary T_e, T_{sp} and h as follows: $T_e \in \{5, 10, 15, 20\}$, $T_{sp} \in \{18, 20, 23\}$ and

$h \in \{-0.2, -0.1, \dots, 0.2\}$, where $T_e < T_{sp}$, $|T_e - T_{sp}| \leq T_{max}$ and T_{max} is the maximum allowable temperature difference ($T_{max}=13^\circ C$). The value of T_e is selected based on real data collected at our experimental site¹, and T_{sp} is based on the ASHRAE standard (Healy 2008).

Fig. 3 shows the expected discomfort and the energy consumption for the PI and MPC controllers with respect to (T_e, T_{sp}) . We see that the MPC algorithm is more efficient in both expected discomfort and energy consumption: the PI controller initially overshoots the temperature set point and requires a long time to reach equilibrium, while MPC quickly approaches and settles on it. MPC reduces energy consumption by 2.6% and reduces expected discomfort by 34.5% compared to the PI-Controller (Table 1).

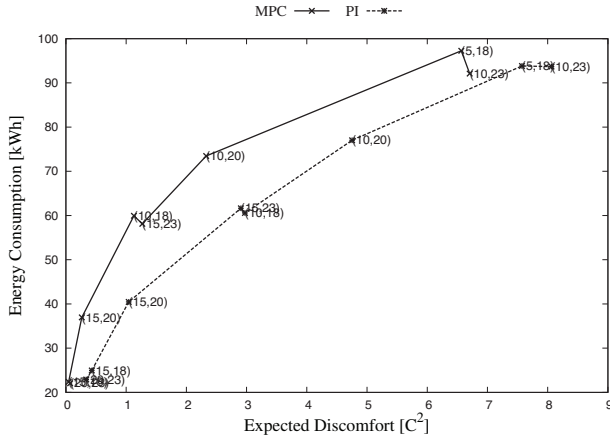


Figure 3: MPC vs. PI-Controller

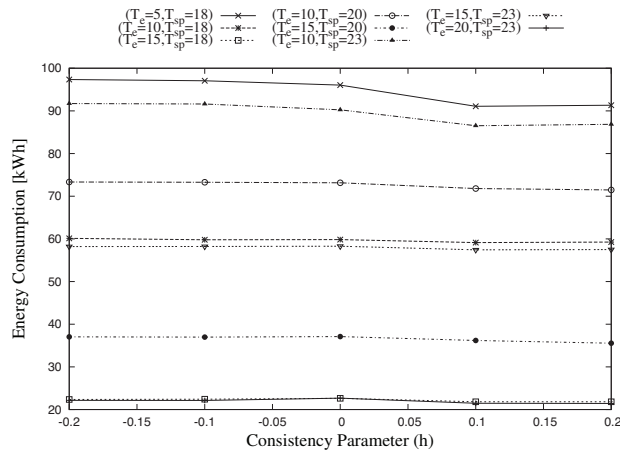


Figure 4: SMPC Energy Consumption Evaluation

We evaluate SMPC (stochastic MPC) for different triples (T_e, T_{sp}, h) , each averaged over 5 samples from the occupancy model. For samples with high certainty (Fig. 4, high h values), the energy consumption is reduced, with a clearer improvement when the temperature difference between T_e and T_{sp} is high. The controller is taking advantage of near-certain absences to reduce the heating. However, this is at the expense of expected discomfort (Fig. 5). Our performance metric accumulates discomfort over time, and since our occupancy model predicts high occupancy, high h values mean

¹<http://www.ucc.ie/en/ERI/>

that the occupant is expected to be present in most time slots, leading to persistent discomfort, while low h values tend to produce many absences and thus less total discomfort. This effect is magnified in scenarios where there is a large gap between T_e and T_{sp} , when the controller trades-off energy use for expected discomfort.

Table 1 summarises the SMPC, MPC and PI controllers in terms of energy consumption and expected discomfort. Both MPC and SMPC show significant improvement in expected discomfort over PI, with moderate energy gains. Compared to MPC, the stochastic version allows a little more expected discomfort, but reduces energy consumption.

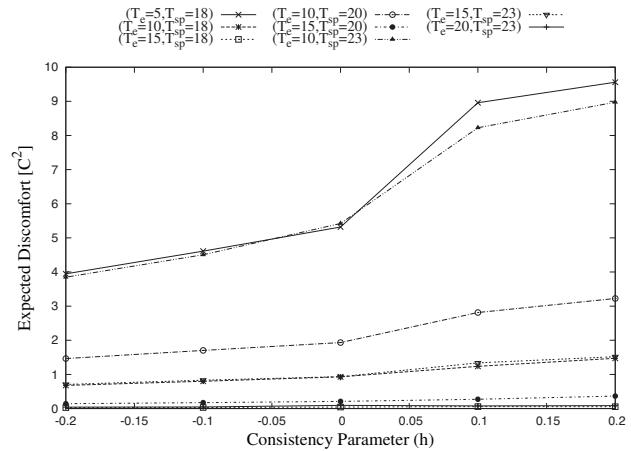


Figure 5: SMPC Expected Discomfort Evaluation

Table 1: Controller Performance Improvement

	Energy Consumption	Expected Discomfort
PI Baseline		
MPC	2.6 %	34.5 %
SMPC	4.3 %	38.3 %
MPC Baseline		
SMPC	2 %	6 %

Although our main aim is to control for *expected* discomfort, we also evaluated the performance of the controller given actual occupancy. We did this by computing *actual* discomfort, where in the performance metric we substitute for $p(k)$ the 0/1 values representing observed occupancy generated by the samples from the model, as shown in Fig. 6. The results here are mixed, reflecting three issues: (i) random variations from the small numbers of samples, (ii) the fact that our stochastic controller operates only on predicted occupancy in the next timeslot, with no role for actual occupancy in the current time slot, and (iii) the bias in our occupancy models towards full occupancy. We will explore the impact of issues (i) and (ii) on energy use and discomfort levels in future work.

To start addressing the third issue, we consider three simple occupancy profiles which include expected absences: (a) an expected absence in the morning ($p(k) \approx 0.1, 9 < k < 12$); (b) an expected absence in the afternoon ($p(k) \approx 0.1, 13 < k < 17$), and (c) expected absence all day ($p(k) \approx 0.1$). The expected discomfort is determined by the number of times the model predicts an absence but the occupant appears, and by the difference between T_e and T_{sp} , and

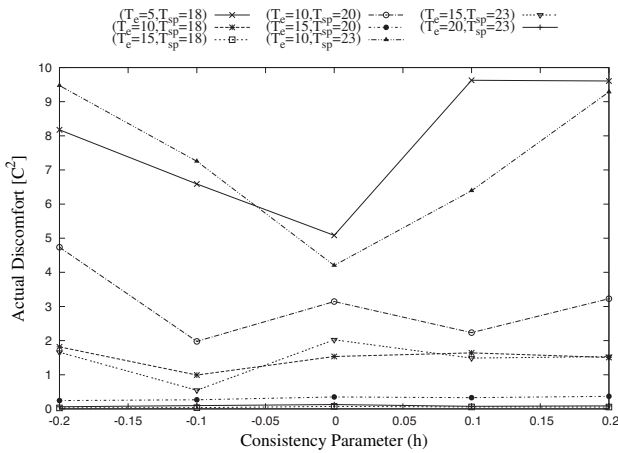


Figure 6: SMPC Actual Discomfort Evaluation

is bounded by $0.1 \times (T_{sp} - T_e)$. The energy consumption, however, is significantly reduced, as shown in Table 2. A complete analysis of this trade-off between energy use and discomfort, and the influence of the certainty of the occupancy profile, will be addressed in future work.

Table 2: Energy Saving

	PI Baseline	MPC Baseline
Profile a	38 %	37.5 %
Profile b	39.9 %	39.4 %
Profile c	37.9 %	37.4 %

Conclusion and Future Work

In this article, we have proposed a Stochastic Model Predictive Control (SMPC) algorithm to regulate temperature in building heating applications². We empirically compared our SMPC strategy for building heating control with the standard PI method, and achieved a 4.3% reduction in energy use and 38.3% reduction in expected user discomfort. We focus on user comfort since it is a cornerstone in the building automation, as it significantly affects the employee's productivity (Fisk 2003). Our experiments compared SMPC with PI and MPC (with fixed occupancy schedules). We followed a two-phase approach, where we first optimized the parameters of the control algorithm, and then compared the effectiveness of the SMPC controller against the other control strategies. We have adopted a stochastic occupancy model as input to the SMPC algorithm.

In future work we plan to extend this work to consider a distributed control system, and to accept input data from occupancy sensors. Moreover, the number of the occupants for meeting rooms has a considerable effect on the heating system. Therefore, we plan to adapt our SMPC approach to take into account the stochastic prediction for the number of occupants and their preferences (e.g. temperature, lighting).

References

Chandan, V.; Mishra, S.; and Alleyne, A. 2010. Predictive control of complex hydronic systems. In *American Control Conference (ACC), 2010*, 5112–5117. IEEE.

²This work is funded by Science Foundation Ireland under grant no. 07/SRC/I1170

Dodier, R.; Henze, G.; Tiller, D.; and Guo, X. 2006. Building occupancy detection through sensor belief networks. *Energy and Buildings* 38(9):1033–1043.

Dounis, A., and Caraiscos, C. 2009. Advanced control systems engineering for energy and comfort management in a building environment—a review. *Proc. of Renewable and Sustainable Energy Reviews* 1246–1261.

Fisk, W. 2003. Health and productivity gains from better indoor environments and their relationship with building energy efficiency. *Annual Review of Energy and the Environment*.

Garcia, C.; Prett, D.; and Morari, M. 1989. Model predictive control: Theory and practice—A survey. *Automatica* 25(3):335–348.

Healy, W. 2008. Effect of temperature stratification near heating elements on the measured energy factors of electric water heaters. *ASHRAE Transactions* 114.

Henzinger, T. 1996. The theory of hybrid automata. *Proc. 11th Annual IEEE Symposium on Logic in Computer Science (LICS 96)* 278–292.

Hoes, P.; Hensen, J.; Loomans, M.; de Vries, B.; and Bourgeois, D. 2009. User behavior in whole building simulation. *Energy and Buildings* 41:295–302.

Kolokotsa, D.; Pouliezios, A.; Stavrakakis, G.; and Lazos, C. 2009. Predictive control techniques for energy and indoor environmental quality management in buildings. *Building and Environment* 44(9):1850–1863.

Mady, A.; Boubekeur, M.; and Provan, G. 2009. Towards integrated hybrid modelling and simulation platform for building automation systems; first models for a simple hvac system. *Proc. 9th Information Technology and Telecommunication Conference (9th IT&T)* 191–199.

Morosan, P.; Bourdais, R.; Dumur, D.; and Buisson, J. 2010. A dynamic horizon distributed predictive control approach for temperature regulation in multi-zone buildings. *18th Mediterranean Conference on Control & Automation (MED)* 622–627.

Oldewurtel, F.; Parisio, A.; Jones, C.; Morari, M.; Gyalistras, D.; Gwerder, M.; Stauch, V.; Lehmann, B.; and Wirth, K. 2010. Energy efficient building climate control using stochastic model predictive control and weather predictions. *American Control Conference (ACC)* 5100–5105.

Page, J.; Robinson, D.; Morel, N.; and Scartezzini, J.-L. 2008. A generalised stochastic model for the simulation of occupant presence. *Energy and Buildings* 40(2):83 – 98.

Siroky, J.; Oldewurtel, F.; Cigler, J.; and Privara, S. 2011. Experimental analysis of model predictive control for an energy efficient building heating system. *Applied Energy*.

Yu, B., and van Paassen, A. 2004. Simulink and bond graph modeling of an air-conditioned room. *Simulation Modelling Practice and Theory, Elsevier*.

Zitzler, E., and Thiele, L. 1998. An evolutionary algorithm for multiobjective optimization: The strength pareto approach. *Swiss Federal Institute of Technology, TIK-Report*.