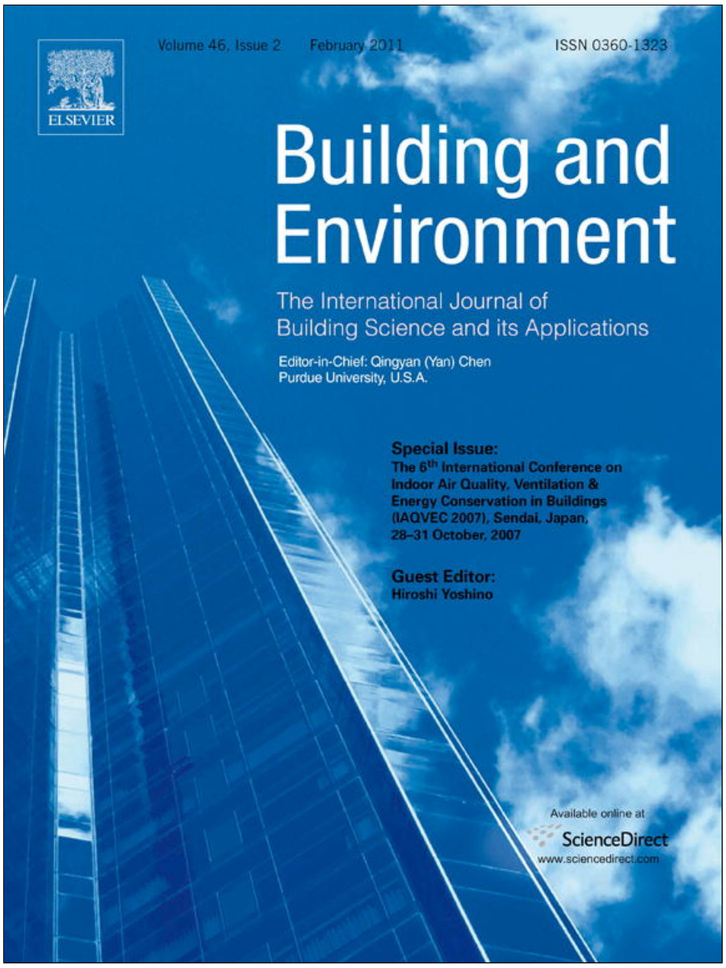
Volume 46, Issue 2 February 2011 ISSN 0360-1323
ELSEVIER
Building and Environment
The International Journal of Building Science and its Applications
Editor-in-Chief: Qingyan (Yan) Chen
Purdue University, U.S.A.
Special Issue:
The 6th International Conference on Indoor Air Quality, Ventilation & Energy Conservation in Buildings (IAQVEC 2007), Sendai, Japan, 28–31 October, 2007
Guest Editor:
Hiroshi Yoshino
Available online at
ScienceDirect
www.sciencedirect.com

Building and Environment 46 (2011) 428–437

Contents lists available at ScienceDirect

Building and Environment

journal homepage: www.elsevier.com/locate/buildenv

Model-predictive control of mixed-mode buildings with rule extraction☆

Peter May-Ostendorp [a,*], Gregor P. Henze [a], Charles D. Corbin [a], Balaji Rajagopalan [a], Clemens Felsmann [b]

[a] *Dept. of Civil, Environmental and Architectural Engineering, University of Colorado at Boulder, Boulder, USA*
[b] *Technical University of Dresden, Dresden, Germany*

ARTICLE INFO

Article history:
Received 9 June 2010
Received in revised form
11 August 2010
Accepted 13 August 2010

Keywords:
Mixed-mode buildings
Model-predictive control
Rule extraction
Logistic regression
Generalized linear model
Particle swarm optimization

ABSTRACT

A series of model-predictive control (MPC) techniques have been explored for optimizing control sequences for window operation in mixed-mode (MM) buildings using EnergyPlus, and results for a simplified MM office building have been presented. Initial results for a small office in Boulder, Colorado show the ability to save upwards of 40% of cooling energy through near-optimal night cooling strategies, even in existing facilities. Strategies can be tuned to avoid overcooling the space by introducing heating energy into the objective function used in the MPC process. A complementary statistical technique has been introduced that allows for the "extraction" of logistic decision models from the optimal control results. The process works best when some time-lagged information is present as a predictor variable to ensure that some process memory is preserved. A generalized linear model (GLM) in the form of a multi-logistic regression was able to mimic the general characteristics of the optimizer results, achieving 70–90% of optimizer energy savings, but at a small fraction of the computational expense. Given the simple mathematical formulation of the logistic regression, it would be possible to implement this sort of decision model into modern direct digital control systems to control MM buildings in a near-optimal manner in real time.

1. Introduction and motivation

Mixed-mode (MM) buildings represent a hybrid approach to space conditioning, employing a combination of natural ventilation and mechanical systems alongside each other and intelligently switching between systems to minimize energy use, while preserving the comfort and well being of occupants [1]. MM buildings have demonstrated reductions in cooling- and ventilation-related energy use from 20% to 50% over code buildings [2,3] and consistently outperform conventional buildings on thermal comfort and occupant satisfaction [4]. However, the performance gains and promise of MM buildings hinge to a large degree on their controls; this is what distinguishes MM buildings from a mere conventional building with operable windows. The effectiveness of the MM control strategy directly determines the extent to which natural ventilation is able to displace mechanical cooling and ventilation, systems that on average account for a quarter of commercial building energy use in the United States today [5].

1.1. Existing control schemes

MM building controls have generally been classified into three topologies. Under *zoned* control, natural ventilation and mechanical conditioning are allowed to occur simultaneously, but in different zones of the building. For example, perimeter offices may be naturally ventilated and core zones mechanically conditioned. In *concurrent* operation, natural ventilation and mechanical conditioning may operate in the same space at the same time. Finally, *changeover* control allows natural ventilation and mechanical conditioning in the same space, but never at the same time. Most MM buildings will not fall cleanly into one of these categories, mainly because at least some amount of zoning is required to provide dedicated mechanical cooling to certain high load spaces like server rooms [1].

In the U.S. design guidelines and best practices for MM buildings have not yet been codified by professional building services organizations. Pioneering research in Europe, such as the International Energy Agency's HybVent project [3], has helped propel MM more into the mainstream. For example, the Chartered Institute of Building Services Engineers (CIBSE) now publishes two application manuals related to MM and naturally ventilated buildings [6,7].

☆ Funding for this research has been provided by the United States Green Building Council.

* Corresponding author.
E-mail addresses: mayosten@colorado.edu (P. May-Ostendorp), gregor.henze@colorado.edu (G.P. Henze), Charles.Corbin@Colorado.edu (C.D. Corbin), balaji.rajagopalan@colorado.edu (B. Rajagopalan), clemens.felsmann@tu-dresden.de (C. Felsmann).

0360-1323/$ – see front matter © 2010 Elsevier Ltd. All rights reserved.
doi:10.1016/j.buildenv.2010.08.004

However, even in Europe, there is no consensus on best practices for MM controls. As such, engineers are left to "start from scratch" or rely on intuition in developing control sequences for these buildings. Algorithms usually involve a series of simple heuristics and if/then statements developed by an HVAC designer for the building's sequence of operations. For example, "if the outdoor temperature drops below 68 °F, open all automated windows and turn off mechanical cooling." It should be noted that most MM buildings are not fully automated, and occupants are usually responsible for operating windows in office spaces. This adaptive approach can reduce the complexity of the control system and has been shown to improve occupant thermal comfort by affording them greater latitude to adapt to thermal disturbances [4,9]. However, introduction of occupant windows can also undermine the energy savings of MM buildings, since people cannot be expected to operate their windows in an energy-efficient manner all the time. As a result, some MM buildings incorporate informational systems, such as notification lights, to signal to occupants when windows should be opened [8].

1.2. Model-predictive control

Model-predictive control (MPC) can be applied to the study of MM buildings to develop optimized control strategies and to provide a benchmark against which existing control strategies can be measured. MPC is a control methodology that seeks strategies through time that minimize an objective or cost function, based on the predictions of a building-level or system-level model — in this case, a building energy simulation model. In the context of building systems, MPC allows for the development of near-optimal operation strategies that minimize the energy use, carbon dioxide emissions, or dollar cost of a facility. Although MPC has been applied extensively in the HVAC engineering field in the past decade [10,11], it has only recently been applied to MM buildings by Spindler and Norford through the optimization of inverse models specifically trained on two unique buildings [12–14].

The present study expands on the work of Spindler and Norford in two important ways. Firstly, it serves as a proof-of-concept for MPC conducted on physical/white box MM building models rather than the inverse/gray box approach taken in the previous research. This allows us to use freely available and validated building energy simulation tools like EnergyPlus to develop control strategies [15]. Secondly, we seek to generalize our results for use in "typical" MM buildings, hence our building models have evolved out of the commercial building reference models developed by Deru et al. [16].

1.3. Rule extraction

In developing generalized control guidelines for MM buildings, we seek to use MPC results to extract simplified control rules that are implementable in buildings today. To achieve truly near-optimal results in real buildings, one could of course couple MPC to a building automation system (BAS) to direct the optimal control actions of the building in real time. However, this approach would require a set of specialized technologies to communicate MPC decisions to a BAS and, therefore, may not be practical in all facilities.

As an alternative, we attempt to use a novel data mining process to extract generalized control rules and decision models from the MPC results. These rules can then serve as the basis for control logic in actual MM buildings, achieving nearly the same energy performance as the actual MPC, but at a fraction of the computational cost. We employ generalized linear models (GLM) for decision modeling purposes, which has been attempted in various forms in the water management field [17,18], but never in the context of MPC in buildings.

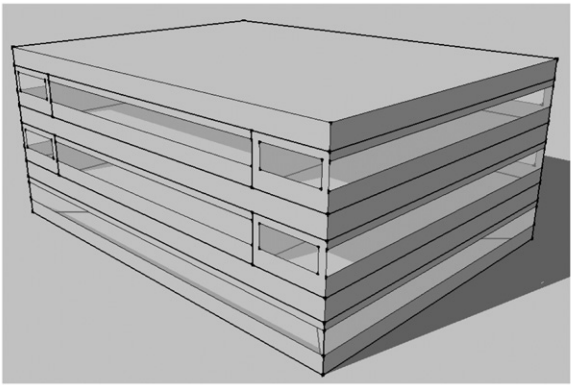

Fig. 1. 3D view of small MM office building model.

The results presented in this paper serve as a proof-of-concept for this suite of techniques, and we provide an outlook for future work, in which the procedures presented will be applied to a broader range of building types and climates and will then be validated on two real MM buildings in southern Germany.

2. Process development and research methodology

2.1. MM building energy model

The focus of initial studies is a prototypical, small (approximately 18,000 ft^2 or 1750 m^2), three-story office building located in Boulder, Colorado, USA, which has been modeled in EnergyPlus. The basic model — including surface geometries, materials, and systems — was adapted from the U.S. DOE reference commercial building models [16]. The floor plan was narrowed slightly to afford better cross ventilation opportunities, per general design rules of thumb presented in the trade press [8]. The building contains a total of 11 occupied thermal zones. The first floor employs standard core-perimeter zoning, whereas the second and third floors have a large open office and two perimeter office zones. An isometric view of the building is presented in Fig. 1. Future research will examine MM buildings with a more sophisticated mix of mechanical systems, but this simplified model is used to demonstrate the validity of the approach.

An issue of great importance in any MM building is occupant control of windows. Occupants have access to operable windows in all but the core zone on the ground floor, and the mean behavior of the occupants is dictated by an implementation of the "Humphreys Algorithm" enforced through an EnergyPlus Energy Management System object. The algorithm was developed based on field studies of occupant behavior in free running buildings by Rijal et al. [19,20], but has also been shown to adequately describe the behavior of occupants in some MM buildings as well [21]. The mean window position is estimated as a function of the current zone temperature in relationship to the occupants' comfort or neutral temperature. Each zone is allotted one instance of the behavior model, yielding the mean manually-controlled window position (between zero and 100%) during occupied periods (windows are assumed closed at other times). Airflow through the building is computed through EnergyPlus' nodal airflow model (i.e. Airflow Network). For the case in question, we employ a concurrent MM design in which packaged rooftop VAV air handlers can cool in conjunction with natural ventilation. This is a highly simplified and admittedly unsophisticated MM system arrangement, but again, the simplicity of

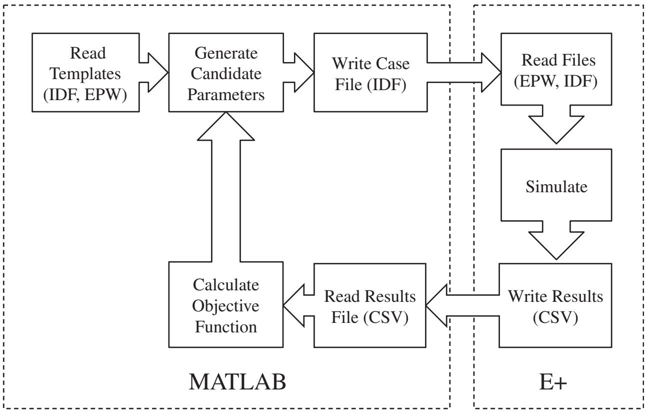

Fig. 2. The ME+ environment, coupling Matlab and EnergyPlus.

building's design can be used to illustrate that the MPC process can find intuitively low energy control strategies.

2.2. Model-predictive control

The goal of MPC in the context of MM buildings is to minimize the energy use, energy cost, or CO_2 emissions of a building (while preserving thermal comfort) by manipulating operable window positions. Naturally, one can simultaneously investigate optimal control of the building's mechanical systems as well, but for demonstration purposes, we limit the decision space to window openings. The problem can be mathematically formulated as the minimally constrained, integer optimization problem:

$$\text{Min } Z(\vec{x}_t) = E + P \tag{1}$$

$$\text{Subject to}: \vec{x}_t \in \{0, 1\}, \tag{2}$$

where $\vec{x}_t$ is a vector of binary decisions (in time) regarding window positions, E is either the energy use or cost over a planning horizon (determined through building energy simulation), and P is a general penalty term applied to discourage certain undesirable characteristics in the solutions, such as thermal discomfort or excessive switching between open and closed window states. The specific development of P and experimentation with different penalty functions is discussed later, but is introduced here in lieu of the set of additional constraints that customarily accompany the problem description.

The cost function that results does not lend itself well to traditional gradient or pattern search techniques because it can contain many local minima. As a result, a meta-heuristic search technique, particle swarm optimization (PSO), has been adopted to quickly and robustly search the decision space for near-optimal solutions. PSO combines simple rules with randomized weighting factors to generate complex search behavior in a population of "particles" evaluating the search space. The action is akin to the flocking behavior of birds and schooling behavior of fish. As with these organisms, information shared between individuals in the swarm affects the decisions of others, all of whom eventually converge on the best solution found by the group. The algorithm is non-deterministic and therefore the search pattern of any swarm is impossible to determine *a priori*. This characteristic of the algorithm decreases the likelihood that it will become stuck in local minima, at the expense of guaranteed convergence upon the true global minimum. The particular implementation of the PSO algorithm used in this study is a variant of the algorithm presented in the foundational work conducted by Kennedy and Eberhardt [22].

In Fig. 2 below, a block diagram schematic of the overall optimization environment is presented to demonstrate the general solution approach. Building models are read in and modified by the PSO algorithm in Matlab by manipulating schedules that control window openings. The resulting models are evaluated using the U.S. DOE's EnergyPlus simulation engine [15,23]. Results are read back into Matlab, where the cost function is computed, and the PSO algorithm decides how to proceed to the next decision vector. The algorithm recurses until a predetermined exit criteria is reached.

This process is used to optimize decisions over a 24-h planning horizon in "one-day-at-a-time" fashion. For the purposes of this simulation study, we assume perfect knowledge of both weather and internal loads during the planning horizon. Attempts to embed such an MPC scheme in a building automation system for supervisory control would likely require more frequent updating of simulation inputs, as well as the use of real (i.e. imperfect) forecast values.

For the cases under consideration in this paper, even though the time step of the building energy simulation is sub-hourly, the planning horizon is segmented into 2-h-long blocks of time called "modes" during which the optimizer is allowed one decision on a given variable. This temporal aggregation of decisions significantly reduces the size of the decision space and the computational expense of the optimization. When the near-optimal decision vector is found for the current planning horizon, the optimizer proceeds to the next day. The thermal history of the building is preserved between planning horizons by running the building through a "historical" period that captures one week of boundary

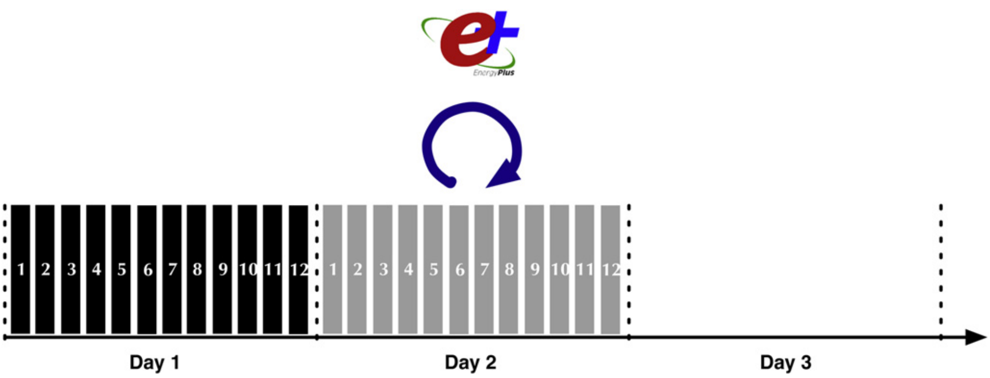

Fig. 3. Procession of Matlab/EnergyPlus MPC environment. Previous decisions (black) determine the thermal history of the building, thus impacting decisions under the current planning horizon (gray). A pre-conditioning period of 1 week is used in actual simulations, rather than 1 day.

conditions and decisions prior to the start of the current planning horizon. This has been shown to preserve the thermal history of the decisions implemented on the previous day(s). This concept is qualitatively illustrated in Fig. 3 below. At the end of a single-variable optimization, the near-optimal result is simply a vector *n*middot*m* hours long, where *n* is the length in days of the optimization period and *m* is the number of modes per day, in our case 12. As mentioned, all decisions are binary.

2.3. Rule extraction

The MPC results are simply a vector in time of binary decisions on window openings for the building. One would ideally like to determine 1) what logic, if any, is being exploited by the optimizer that is embedded in these results and 2) determine a streamlined process for extracting that logic which minimizes human subjectivity. The following sections describe the various techniques employed to extract a nested decision model from the offline MPC results. Models were trained on results from a summer-long optimization, and model parameters were pruned through a stepwise regression process. The robustness of the best models was examined through cross-validation, and the skill of the various models was evaluated using a probabilistic skill score. All of the statistical investigations were performed using the technical computing language R.

2.3.1. GLM model formulation

The generalized linear models (GLM) framework was used to build simplified decision models that mimic the general characteristics of the offline MPC optimal solutions. Because the MPC optimization examined here provides binary decisions, a statistical model was required that could offer predictions on binomially distributed data. Given recent attempts at modeling occupant control of windows through logistic regressions [19,20], the logistic link GLM (i.e. a multiple logistic regression) was chosen. This model was used to relate the optimizer window predictions, *y*, to a given predictor variable *x* through the logistic link function:

$$\theta(x) = \log\frac{p(x)}{1-p(x)}, \tag{3}$$

where $p(x)$ is the probability of a window opening signal being issued. In physical terms, $p(x)$ could also be interpreted as the fraction of windows open in the building or a uniform opening factor applied to each window in the building. For the general case, one might have a total of l observations on which to train the model, using a total of m possible predictor variables, and the GLM takes the form shown in Equation (4)

$$\widehat{\theta} = \widehat{\beta}\mathbf{X}', \tag{4}$$

where $\widehat{\theta}$ is a $1 \times l$ vector containing the predicted values of θ, $\widehat{\beta}$ is a $1 \times (m+1)$ vector of the estimated model parameters, and $\mathbf{X}'$ is a $(m+1) \times l$ augmented matrix of predictor variables (the leading row contains all ones). Unlike standard multiple linear regression, the model parameters cannot be determined in closed form, and an iterative process must be used instead. The common approach is to choose parameters that maximize the model's likelihood function, *L*. We do not discuss the underlying mathematical formulation here, but the purpose of this approach is to find a set of model parameters that maximize the likelihood of reproducing the data distribution of the training set. This algorithm is implemented by default through R's standard GLM libraries. The predicted probability of window opening, $\widehat{p}$, is found through the inverse logit function

Table 1
Predictor variables considered.

Variable	Description
T_{oa}	Outdoor dry bulb temperature
T_{dp}	Outdoor dew point temperature
v_{wind}	Wind speed
θ_{wind}	Wind direction
I_{dn}	Direct normal solar radiation
T_{core}	Core zone temperature (first floor only)
$T_{floor,zone}$	Mean temperature for a given floor and zone (total of 10)
y	Binary window state at a given point in time

$$\widehat{p} = \frac{e^{\widehat{\theta}}}{1+e^{\widehat{\theta}}}. \tag{5}$$

To further process the output of the GLM such that it can be used as a binary control signal, two additional post-processing techniques were applied. First, the probability "signal" was converted back to a binary by assigning a 1 or 0 to the GLM output based on exceeding a threshold. The optimal threshold was found through a sequential search operation that minimized the sum of squared errors between the binary signal and the original optimizer signal. In nearly all cases, the optimal threshold was approximately 0.5. To eliminate unwanted noise in the model's output, a form of hysteresis smoothing was applied to the binary signal such that state changes of less than 2 h in duration were ignored. In this way, we obtained a cleaner and more desirable signal from the GLM.

2.3.2. Model input pruning

We initially examined predictor variables that can be readily measured in today's buildings. For this reason, we employed outdoor dry bulb and dew point temperature, wind speed, wind direction, global horizontal solar irradiation, and zone temperatures as potential model inputs. Time-lagged input variables were also considered in an attempt to incorporate the process memory inherent in building heat transfer. For example, one might consider both the current temperature, *T*, and the previous hour's temperature, T_{t-1}, as predictors. To introduce autoregressive characteristics into the model, one can also use the prior window opening state, y_{t-1}, as a predictor variable. A summary of the variables considered and their nomenclature is presented in Table 1 below. If one includes all the predictors, their 1-h lagged values, and the previous window position state as predictors, this yields a total of 33 model inputs, only some of which may contribute significantly to the model's predictive power.

To "prune" the various predictor variables, a stepwise regression technique was employed that minimizes the Akaike's Information Criterion (AIC). The AIC is a statistical figure of merit that objectively measures the model's ability to reproduce the variance of the observations with the fewest model parameters [25]. It is given by

$$\mathrm{AIC} = 2k - 2\ln(L), \tag{6}$$

where *k* is the number of model parameters and *L* is the maximized value of the likelihood function for the model. The stepwise regression method employed attempts a "forward" and "backward" search through the potential list of model predictors and determines the combination of predictors that maximizes the AIC. An exhaustive search of predictor combinations was also attempted, but this provided little benefit over the stepwise search and resulted in a significant number of iterations for even small numbers of predictors.

2.3.3. Evaluation of model skill

An objective criterion for measuring the predictive power of the model is required. We can measure its performance against the original near-optimal control signal through a ranked probability score (RPS), demonstrating the degree to which the model predicted the original optimizer results (the RPS is quite nearly a squared error indicator for categorical predictions like binary window openings, in which a 0 value indicates a perfect score). However, this information is only partly useful, since the optimizer signal might just as easily be reproduced by a white noise process. We can therefore compare the RPS of the model against the RPS for a random process and determine in a relative sense which is more effective at reproducing optimizer results. This is done through the ranked probability skill score (RPSS), which has been used in various climatological contexts to compare model skill in predicting categorical rainfall and streamflow quantities [24]. The method has been described in detail by Wilks [25]. Note that when the RPSS is applied to two-category forecasts, it reduces to the Brier Skill Score (BSS), also described in Ref. [25]. We employ the more general RPSS here because it could later be adapted to multi-category forecasts, such as the predicted state of a three-position window controller (e.g. closed, half-open and open).

The RPSS compares the accuracy of model predictions against chance, but rather than simply compare our model against a 50–50 chance of a window opening, we compare against the probability of window openings we see in the optimizer results — a rather weighted coin. This provides for a more rigorous test of model performance. The RPSS is negative if model results are worse than chance, 0 if model results reproduce chance events, and positive if model results are closer to the original observations than chance.

The score is computed by dividing window opening predictions into j categories, in this case two because our optimizer produces two window states: open and closed. A vector of forecast probabilities, p_j, is constructed based on the GLM model predictions. Similarly, a vector of observed events, y_j, is constructed from the optimal results, in which y_1 is 1 for window closings and y_2 is 1 for window openings. We then take the cumulative density function of p_j and y_j, resulting in the 2-category vectors, p_{cdf} and y_{cdf}. Note that, in our case, the RPS is computed for each instance that a window opening is predicted, so over a given 24-h period, we would naturally have a total of 24 RPS values for each hour during the day when the GLM was employed. The RPSS is then computed by forming a ratio between the average RPS values of the model and chance as shown below.

$$\mathrm{RPS} = \sum_{i=1}^{j} \left(p_{\mathrm{cdf}} - y_{\mathrm{cdf}}\right)^2 \tag{7}$$

$$\mathrm{RPSS} = 1 - \frac{\overline{\mathrm{RPS}}_{\mathrm{model}}}{\overline{\mathrm{RPS}}_{\mathrm{chance}}}. \tag{8}$$

2.3.4. Cross-validation

A process of cross-validation was used to examine the robustness of several models. Cross-validation (CV) is often used to quantify model skill in predicting values that were not part of the training dataset. In typical CV implementations, a select subset of data is dropped from the original dataset of observations, and the statistical model is developed based on the remaining observations in the training set. The model is then used to predict the values of the dropped subset of points. The process is repeated for every subset in the dataset, and some figure of merit, such as AIC, is computed for each prediction to assess the robustness of the model in predicting values that it has not seen before.

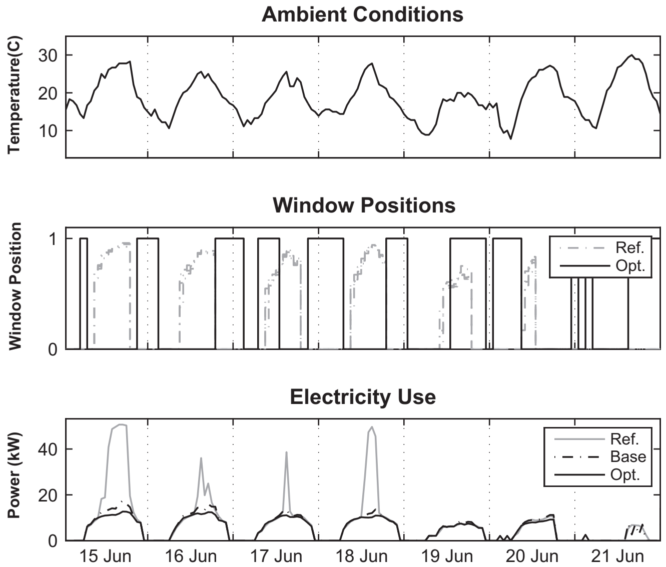

Fig. 4. Optimal solution and energy use profile for a typical week in mid-June. The optimal solution outperforms the base and reference cases through a night ventilation strategy. In the reference case, occupant window openings during peak solar gains and daytime temperatures result in a dramatic spike in fan and direct expansion cooling energy use.

In our case, we have chosen to cross-validate by dropping individual day-long sequences of data from the training set in a sequential fashion. Since our original training set is 11 weeks long, this provides 77 sets of data with which to cross-validate the model performance. The RPSS is then computed for each of the 77 CV datasets and can be compared to the RPSS values originally found across the entire summer period.

3. Results

3.1. Model-predictive control

MPC runs were conducted on the prototypical MM small office using publicly available TMY3 weather data for Boulder, CO [26]. The total window of the optimization spanned June 15 through August 30 — exactly 11 weeks. During each 24-h planning horizon, the optimizer manipulated a single binary decision vector for global window on/off position in 2-h blocks, for a total of 12 decisions per day. The only constraint on the optimization was the requirement of binary window opening decisions. Individual optimization runs for the 11-week period lasted approximately 12 h when run on a 2×2.8 GHz 4-core Intel Xeon server. A "base" and "reference" case were run over the same period for comparison. The base case is the standard DOE benchmark building, without any natural ventilation, whereas the reference case is the corresponding MM building, but with occupant window control per the modified Humphreys Algorithm.

The initial objective function used for optimization captures the electric energy use of the cooling equipment in the building (both fans and DX cooling equipment) and adds a penalty that scales linearly with the number of transitions between window states. The switching penalty has been limited to 5% of the total cooling energy use of the building, effectively stating that we are willing to forgo 5% energy savings for a cleaner solution.

$$Z(\overrightarrow{x}_t) = E\left(1 + 0.05P_{\mathrm{switching}}\right) \tag{9}$$

E is the cooling electricity use and $P_{\mathrm{switching}}$ is a penalty function derived from the sum of state transitions in the optimizer's current decision vector.

Table 2
Summer-long savings comparison.

	Objective Function	
	Cooling	Cooling & Heating
Base Case		
Electricity Use (kWh)	12,258	12,258
Gas Use (kWh)	54	54
Total Energy Use (kWh)	12,313	12,313
Reference Case		
Electricity Use (kWh)	19,681	19,681
Gas Use (kWh)	4,518	4,518
Total Energy Use (kWh)	24,199	24,199
Optimal Case		
Electricity Use (kWh)	10,628	10,735
Gas Use (kWh)	2905	335
Total Energy Use (kWh)	13,533	11,070
Optimizer Savings vs. Base Case		
Electricity (kWh)	1630	1523
Gas (kWh)	−2851	−281
Total (kWh)	−1221	1242
Total (%)	−10%	10%
Optimizer Savings vs. Reference Case		
Electricity (kWh)	9053	8946
Gas (kWh)	1613	4183
Total (kWh)	10,666	13,129
Total (%)	44%	54%

Fig. 4 illustrates some representative results from a "swing season" week in June. The upper graph shows ambient temperatures over the week. The middle graph shows the optimal solution alongside the mean window opening behavior in the reference model. The base case building is, by definition, sealed, so no window positions are shown. The hashed portions of the optimizer solution represent concurrent natural ventilation and mechanical cooling. Finally, the bottom graph illustrates electric power consumption for HVAC equipment for all three cases, showing time periods during which savings accrue.

The optimal solution found during the early summer manifests itself as a night ventilation strategy, with the optimizer opening windows during cooler nighttime periods. This form of passive thermal energy storage utilization allows the building to ride out some of the daytime cooling loads without the need for mechanical cooling. Load reductions are modest for the week shown, but provide double-digit percentage savings over the course of the season (see Table 2). The electricity savings accrue most noticeably on the cooling peak in the afternoon for all days except June 19, which is somewhat more mild day with lower cooling loads to begin with. The optimizer arrived at these solutions with minimal constraints and absolutely no expert knowledge about the decision space, yet the solution mimics a heuristic approach used in many mixed-mode buildings.

It is interesting to note that the simulated occupants in the reference case building frequently makes use of the operable windows during relatively warm periods of the day that also coincide with peak solar gains, resulting in increased cooling loads on the mechanical system. In short, mean occupant behavior in this building is highly suboptimal since the HVAC system is allowed to operate concurrently, and thus the reference case building performs poorly even when compared to the fully sealed base case building. More importantly, the occupants do not have the ability to recreate the night cooling strategy, so as a result, the reference building actually consumes *more* energy compared to both the optimal and base cases.

Unfortunately, we can see that the cooling energy savings occur at some expense both to heating energy use and thermal comfort.

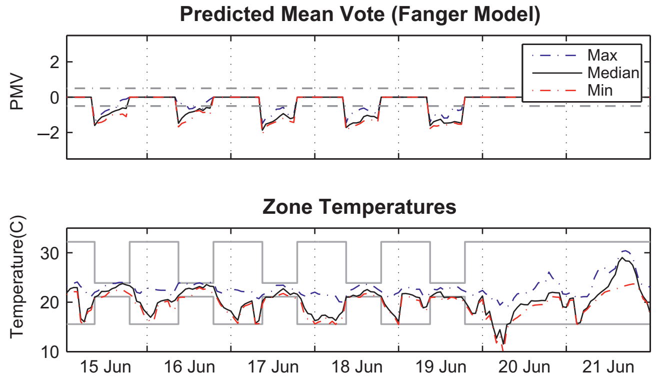

Fig. 5. Comfort summary for the same one-week period of the optimizer solution. Cooling and heating setpoints are denoted by the gray lines above and below the zone temperatures, respectively. Note that zone temperatures and predicted mean votes are low, particularly at the beginning of occupancy, due to night ventilation. PMV values have been zeroed for unoccupied periods.

Fig. 5 shows zone temperatures and thermal comfort during the same June week. By the end of many night ventilation periods, temperatures in most zones have dropped to the heating setpoint, and the heating system is activated periodically as a result. The low zone temperatures also result in low predicted mean votes or a "very cool" thermal sensation at the beginning of the occupied period. There are also serious consequences for energy use, as shown in Table 2. Note for the comparison with the base case building that the "optimal" solution actually uses 10% *more* energy than the base case when the algorithm is only seeking to reduce electricity consumption.

To counteract this overcooling, we can expand the scope of the objective function to include heating energy as well. Under the new objective function, energy use associated with cooling electric and heating natural gas consumption are included, in addition to the previous penalty term from Equation (9). The results for the same week-long period in June are shown in Figs. 6 and 7, with a comparison of HVAC electricity and gas usage in units of kWh. Note that night cooling is significantly shortened and only occurs during mild cool periods of the night, avoiding extremes in cold that might otherwise result in overcooling and, therefore, additional heating. This means, of course, that the magnitude of cooling savings is somewhat reduced for this week. However, when taken on a seasonal basis, there are ample opportunities for free cooling that do not result in drastically higher gas consumption (Table 2).

3.2. *Rule extraction*

3.2.1. *GLM development with stepwise parameter pruning*

Four models, summarized in Table 3, were formulated to compare different parameter sets and pruning approaches. For models 1 through 3, a stepwise regression approach was used to find the parameter set that minimized the model AIC. With each successive model, a larger number of lagged predictor variables were included to attempt to capture process memory. Model 1 utilized only current time step predictors (x_t), whereas model 2 included 1-h lagged predictors as well (x_{t-1}). Model 3 included the previous hour's optimal window state (y_t) as a predictor in addition.

The fourth model was manually pruned based on the results of models 1 through 3 in an attempt to eliminate some of the redundancy present in the zone temperature variables. The large number of parameters for models 1 through 3 result from the inclusion of multiple zone temperatures in the best model predictor set. Rather than selecting a full complement of zone

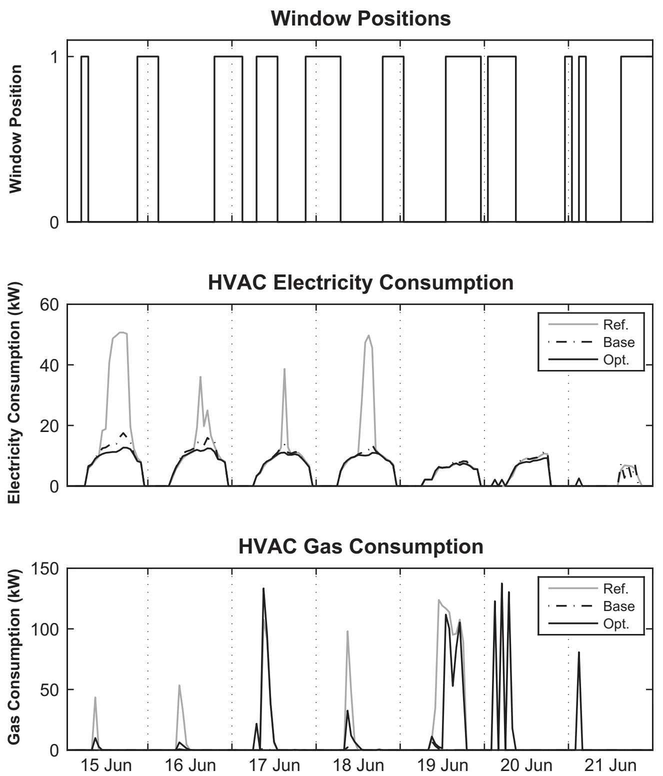

Fig. 6. The electricity-only objective function results in large heating spikes during the early morning hours, seen in the bottom gas consumption plot.

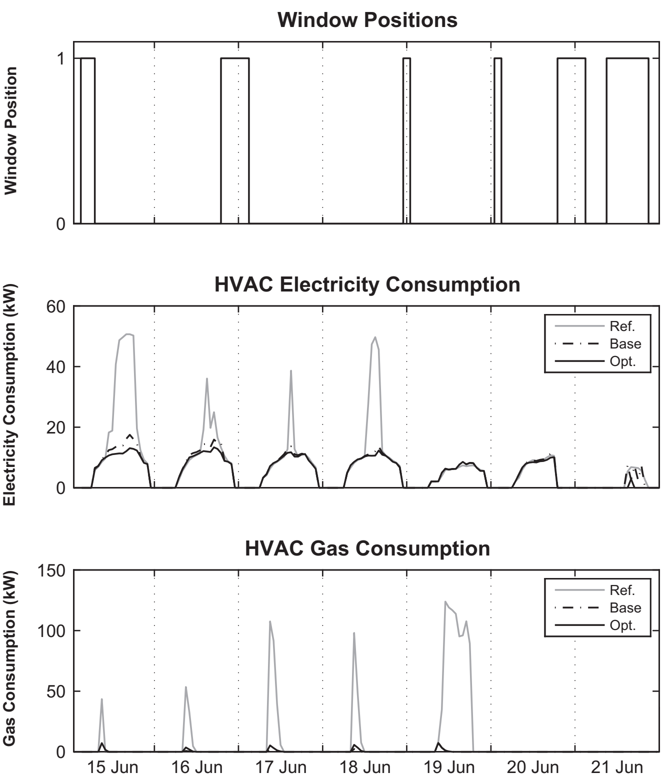

Fig. 7. The more holistic electricity and gas objective function discourages overcooling, resulting in shorter periods of night flush ventilation and only on the milder nights of this particular week. Spikes in heating are automatically eliminated.

Table 3
GLM model parameter summary.

Model	Predictor Types x_t	x_{t-1}	y_{t-1}	Predictors Considered	Predictors Used
1	x			16	11
2	x	x		32	24
3	x	x	x	33	24
4		x	x	33	14

temperatures for each floor of the building, one zone can be used to represent each floor, resulting in a model with a total of 14 predictors (including lagged terms). Although this model neither exhibits as low an AIC nor as high a RPSS as the best model (#3), it contains 10 fewer parameters mostly as a result of zone temperatures that have been eliminated.

The resulting model predictions and the original optimizer sequence are presented in Fig. 8 for the week of June 15 though 21, with probabilistic predictions as a dashed line. All models perform remarkably well when compared to the original data, tracking periods of opening and closing accurately, even very brief periods, such as the 2-h opening occurring on the first night of the week. However, note that models 1 and 2 also miss several long periods during which the optimizer windows were open, namely the beginning of the sixth and seventh nights. In the case of model 2, these openings are missed as a result of the hysteresis placed on the output signal.

3.3. Model cross-validation

A "remove-one-day" cross-validation of the models above demonstrated that the best models could be used to effectively predict sequences of window opening to which they were not exposed in the model fitting process. Fig. 9 shows sample CV predictions as a binary signal for the manually pruned model 4 during weeks 1 and 5 of the optimization period. The performance under CV is nearly as accurate as the results under non-CV conditions from Fig. 8.

The RPSS of the CV predictions were also examined for the entire 77-day cooling season period to diagnose any particular weaknesses in the models' predictive abilities. Fig. 10 provides a plot of those values for the two best models found, 3 and 4. The dashed lines represent the mean RPSS found for the entire 77-day period under non-CV conditions. Most obviously, we see that the manually tuned model 4 performs significantly worse under CV than model 3, with RPSS values proportionately lower than the corresponding values for model 3. It is possible that through the subjective process of manually tuning the model, some non-redundant zone temperatures were excluded from the predictor set.

A second interesting phenomenon plagues both models. Each GLM demonstrates skill in predicting optimal window openings for most weekdays; however, they often encounter difficulties predicting openings around weekends. This near-weekly pattern suggests that, in future iterations of the model, occupancy or general building use patterns might need to be taken into account as predictor variables, as these parameters are both week in nature and have significant impact on the energy use of the building systems. This would assist in capturing the periodicity of the process and would provide the GLM with information currently available to the optimizer (i.e. the impact of occupancy patterns on energy use), but effectively hidden in the statistical model formulation. In a practical controls implementation, sensing occupancy throughout the entire building could present challenges and considerable expense. The building's setpoint schedule could serve

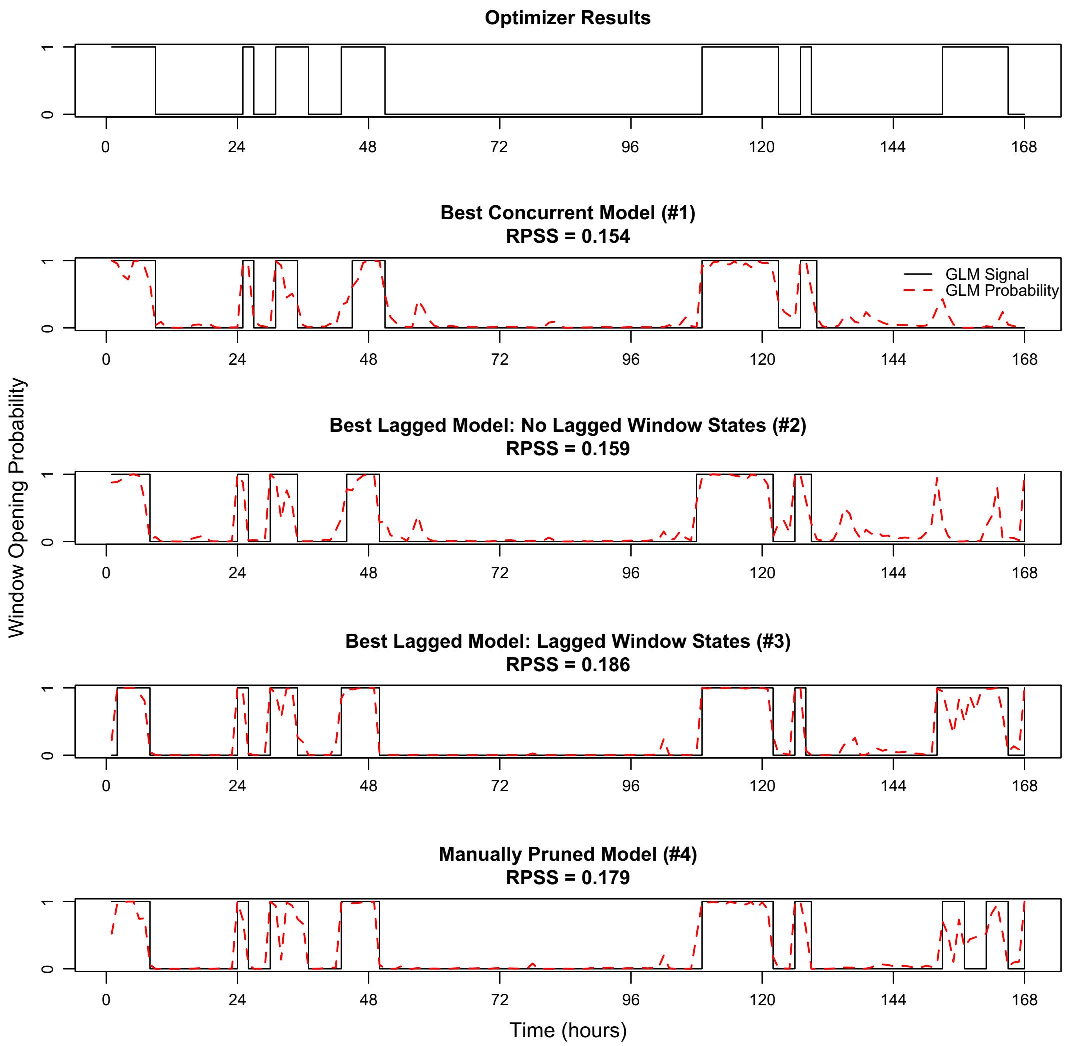

Fig. 8. Predictions of window position for models 1 through 4 are presented above alongside offline MPC optimal results for a typical week in mid-June. Models 3 and 4 demonstrate the overall best performance. Both models contain the autoregressive window state term.

as a proxy for occupancy, since these schedules are usually developed by building operators around the expected occupancy of a building. Alternately, a sinusoidally varying "clock" signal could be implemented as a predictor to provide a signal with a known weekly period, akin to the technique employed by Dodier and Henze in Ref. [27].

3.3.1. *Final model parameters*

Based on the series of diagnostics described above—particularly the CV results—model 3 was selected for testing in energy simulation. The final model parameter set and the salient parameters from an analysis of variance (ANOVA) are provided in Table 4 below. ANOVA is used here primarily to judge the statistical significance of the various model parameters, $\widehat{\beta}$, with the null hypothesis that the coefficients are zero. The most self-evident feature of the parameter set is the extremely high confidence (α) of the lagged window state term in comparison to all other model parameters. This is not surprising, given the prominent increase in predictive skill observed by adding the lagged window state term in models 3 and 4.

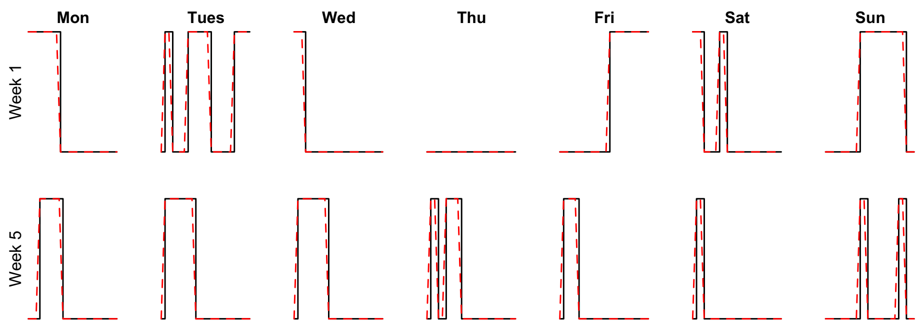

Fig. 9. The GLM binary window position signal (dashed) compared to optimizer decisions (solid) for the first and fifth weeks of the optimization period. Model predictions are a near-identical match to the original MPC results.

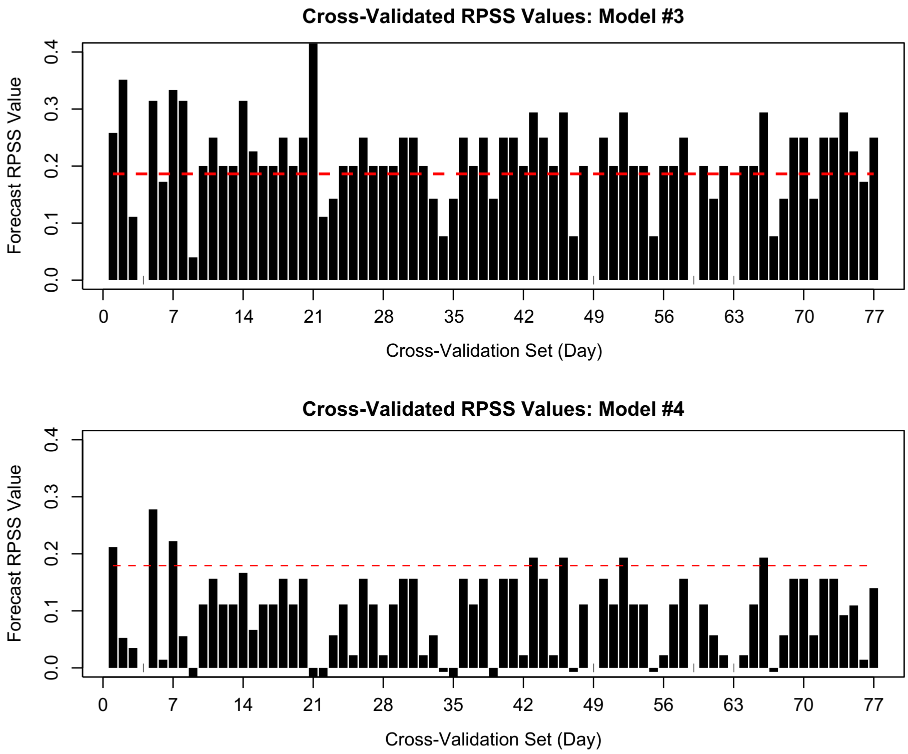

Fig. 10. RPSS values for each of the day-long cross-validation periods for the 6-variable model. The RPSS of the summer-long, non-CV model is shown by the dashed line.

3.4. Rule implementation in energy model

Once an appropriate decision model has been extracted, the model can be used as a form of controller to signal changes in window position. This process can be simulated in EnergyPlus to allow for a comparison with the optimizer and determine what portion of the optimizer savings we are able to capture through the nested decision model. In its preliminary implementation, this has been done by simply creating a window opening schedule based on the predictions of the GLM and re-running the building energy simulation with this schedule controlling the window positions. This implementation is not meant to function as a true controller, but rather to demonstrate that the GLM we have extracted from the optimal results and its predictions are nearly as effective at achieving energy savings as the optimizer itself and at a fraction of the computational cost.

Table 4
Model 3 parameters.

Model Parameter	$\hat{\beta}$	Standard Error	α-Value (%)[a]
Intercept	−15.23	5.24	99.640
T_{oa}	0.40	0.13	99.780
v_{wind}	−0.19	0.11	92.600
I_{dn}	0.0047	0.0010	99.999
T_{core}	1.13	0.29	99.998
$T_{bot,1}$	−0.71	0.31	97.600
$T_{bot,3}$	1.12	0.51	97.200
$T_{mid,1}$	−1.19	0.50	98.300
$T_{mid,2}$	6.57	1.59	99.999
$T_{mid,3}$	−6.52	1.55	99.999
$T_{top,2}$	−9.12	1.75	99.999
$T_{top,3}$	7.46	1.50	99.999
$T_{oa,t-1}$	−0.56	0.15	99.983
$T_{dp,\ t-1}$	−0.11	0.05	97.400
$I_{dn,\ t-1}$	−0.0020	0.0011	92.800
$T_{bot,2,\ t-1}$	0.78	0.55	82.000
$T_{bot,3,\ t-1}$	−1.68	0.65	99.020
$T_{bot,4,\ t-1}$	1.42	0.56	98.900
$T_{mid,1,\ t-1}$	−5.40	1.46	99.979
$T_{mid,2,t-1}$	−3.10	2.08	86.000
$T_{mid,3,\ t-1}$	6.16	1.94	99.850
$T_{top,1,\ t-1}$	6.36	1.55	99.995
$T_{top,2,\ t-1}$	4.32	2.06	96.400
$T_{top,3,\ t-1}$	−7.01	1.85	99.985
y_{t-1}	8.05	0.70	99.999

[a] All α-values greater than 99.999% presented as 99.999%.

Table 5 provides a simple summary of the energy savings achieved for the whole summer with the optimizer and using the best performing model (#3) described above. The GLM is able to achieve over 90% of the energy savings when compared to the reference case building, but only about 70% on the base case building. The timing of window openings for the reference case is somewhat less crucial to energy savings, so long as windows are not opened during the day, for it is this time that occupants will tend to open windows and often add dramatically to cooling loads. In the comparison against the base case building, the solution is more sensitive to the appropriate timing and duration of window openings because the base case building is already under tight control, so the optimizer outperforms the GLM. This is consistent with the CV results presented in Fig. 10, which shows how the GLM tends to poorly predict optimizer actions under certain circumstances.

There is a clear trade-off between the use of MPC and a simplified decision model like our GLM in the context of controlling real buildings. Indeed, the optimizer is able to achieve consistently higher energy savings, but at significant computational cost. The GLM only needed to be run once using simple linear algebra that takes fractions of a second to compute, as opposed to the optimizer, which required thousands of daily runs of a full EnergyPlus model to converge on a solution meeting our tolerance criteria (a 12-h process). Clearly the MPC approach as described here would not be

Table 5
Summer-long savings comparison: optimizer vs. GLM.

	Optimizer	GLM
Base Case		
Electricity Savings (kWh)	1630	1190
Electricity Savings (%)	13%	9.3%
Reference Case		
Electricity Savings (kWh)	9053	8613
Electricity Savings (%)	46%	43%

practicable in a real-time building control application with todays computing hardware and simulation capabilities. However, the GLM poses a viable, mathematically simple and computationally efficient alternative that could be implemented "in the loop" with direct digital control systems.

4. Conclusion and outlook

A series of MPC techniques have been explored for optimizing control sequences for window operation in MM buildings, and results for a simplified MM office building have been presented. Initial MPC studies on MM buildings have demonstrated the capability to optimize these types of buildings using a physical modeling approach, rather than the data-driven approaches attempted in previous research [Spindler and Norford (2008), Spindler and Norford (2008), Spindler (2004)]. The physical modeling approach provides several advantages, mainly that we can explore prototypical MM building designs in different climates and that we can begin to couple occupant behavior to these models through new simulation program features. Initial MPC results for a small office in Boulder, CO shows the ability to save upwards of 40% of cooling energy through near-optimal night cooling strategies, even in existing facilities. Strategies can be tuned to avoid overcooling the space by introducing heating energy into the objective function used in the MPC process.

A complementary statistical technique has been introduced that allows for the "extraction" of logistic decision models from the optimizer results, effectively allowing us to examine the logic embedded in the optimizer solution. The GLM technique presented works best when some time-lagged information is present as a predictor variable to ensure that some of the process memory is preserved. Even with a simplified model that does not incorporate this autoregressive component, we can very closely mimic the general characteristics of the optimizer results, achieving 70–90% of optimizer energy savings, but at a small fraction of the computational expense. Given the simple mathematical formulation of the GLM, it would be possible to implement this sort of decision model into modern direct digital control systems to control MM buildings in a near-optimal manner in real time.

This paper has served as an introduction to and proof-of-concept for the above-described techniques, but significant additional research is warranted and ongoing. With regards to MPC investigations, extensive parametric studies examining mechanical system complexity, climate and occupant comfort expectations are needed to fully explore the nuances of various MM system controls. Since the particle swarm optimization algorithm used is inherently stochastic in its exploration of the decision space, additional investigations are being conducted to determine any potential variance in the MPC solutions.

On the topic of rule extraction, significant follow-on research is underway to improve the GLM technique presented and its predictive skill. The issue of multicollinearity in predictor variables can complicate the process of selecting a parameter set for a model. Principal component analysis shows promise as a technique for eliminating multicollinearity issues. It will also be crucial to demonstrate that the GLM maintains its predictive skill in a real-world environment. In a real building or even simulated implementation, the GLM model output will influence the thermal state of the building, which in turn will influence the subsequent decisions of the GLM. It must be demonstrated that GLMs, when implemented as controllers, can maintain their model skill and stability when the coupling of model inputs and outputs is taken into account.

References

[1] Brager G, Borgeson S, Lee Y. Summary report: control strategies for mixed-mode buildings. Technical Report. University of California, Berkeley, CA: Center for the Built Environment; 2007.

[2] Torcellini P, Deru M, Griffith B, Pless S, Judkoff R, Crawley D, Lessons learned from field evaluation of six high-performance buildings. In: Proceedings of the ACEEE summer study on energy efficiency in buildings; 2004. p. 1–13.

[3] Heiselberg P. Principles of hybrid ventilation. Technical report. Aalborg, Denmark: Aalborg University; 2002.

[4] Brager G, Baker L. Occupant satisfaction in mixed-mode buildings. In: Proceedings of air conditioning and the low carbon cooling challenge. Windsor, UK; 2008.

[5] US DOE EIA, Table E6. Electricity consumption (kWh) intensities by end use for non-mall buildings, Technical Report, US Department of Energy – Energy Information Administration, 2003.

[6] Chartered Institution of Building Services Engineers, CIBSE AM10: natural ventilation in non-domestic buildings, Page Bros; 2007.

[7] Chartered Institution of Building Services Engineers, CIBSE AM13–2000: mixed mode ventilation, Page Bros., 2000.

[8] McConahey E. Mixed mode ventilation: finding the right mix. ASHRAE Journal 2008;50:36–48.

[9] DeDear R, Brager G. Developing an adaptive model of thermal comfort and preference. ASHRAE Transactions 1998;104:1–18.

[10] Henze G, Felsmann C, Knabe G. Evaluation of optimal control for active and passive building thermal storage. International Journal of Thermal Sciences; 2004:173–83.

[11] Henze G, Le T, Florita A, Felsmann C. Sensitivity analysis of optimal building thermal mass control. Journal of Solar Energy Engineering; 2007.

[12] Spindler H, Norford L. Naturally ventilated and mixed-mode buildings—part I: thermal modeling. Building and Environment; 2008:736–49.

[13] Spindler H, Norford L. Naturally ventilated and mixed-mode buildings—part II: optimal control. Building and Environment; 2008:750–61.

[14] Spindler H. System identification and optimal control for mixed-mode cooling, Ph.D. thesis. Massachusetts Institute of Technology; 2004.

[15] The Regents of the University of California, The Board of Trustees of the University of Illinois, EnergyPlus; 2009.

[16] Deru M, Griffith B, Torcellini P. Establishing Benchmarks for DOE Commercial Building R&D and Program Evaluation. In Proceedings of the ACEEE summer study on energy efficiency in buildings. Asilomar, CA; 2006. p. 1–12.

[17] Bessler FT, Savic DA, Walters GA. Water reservoir control with data mining. Journal of Water Resources Planning and Management 2003;129:26–34.

[18] Wei C, Hsu N. Optimal tree-based release rules for real-time flood control operations on a multipurpose multireservoir system. Journal of Hydrology 2009;365:213–24.

[19] Rijal H, Tuohy P, Humphreys M, Nicol J, Samuel A, Clarke J. Using results from field surveys to predict the effect of open windows on thermal comfort and energy use in buildings. Energy and Buildings 2007;39:823–36.

[20] Rijal H. Development of an adaptive window-opening algorithm to predict the thermal comfort, energy use and overheating in buildings. Journal of Building Performance Simulation; 2008:17–30.

[21] Rijal H, Humphreys M, Nicol J. How do the occupants control the temperature in mixed-mode buildings? Predicting the use of passive and active controls. In: Proceedings of air conditioning and the low carbon cooling challenge. Windsor, UK; 2008. p. 1–15.

[22] Kennedy J, Eberhart R. Particle swarm optimization. In: Proceedings of the IEEE international conference on neural networks, vol. 4, p. 39–43.

[23] The Mathworks Inc., MATLAB R2009b, 2009.

[24] Regonda S, Rajagopalan B, Clark M. A new method to produce categorical streamflow forecasts. Water Resources Research 2006;42.

[25] Wilks DS. Statistical methods in the atmospheric sciences. Academic Press; 1995.

[26] The Regents of the University of California, The Board of Trustees of the University of Illinois, EnergyPlus: Weather Data; 2008.

[27] Dodier R, Henze G. Statistical analysis of neural networks as applied to building energy prediction. Journal of Solar Energy Engineering 2004;126:592.