Energy-Comfort Optimization using Discomfort History and Probabilistic Occupancy Prediction

Abhinandan Majumdar*, Jason L. Setter*, Justin R. Dobbs†, Brandon M. Hencey† and David H. Albonesi*
*Computer Systems Laboratory, Cornell University, Ithaca, New York 14853, USA
†Mechanical and Aerospace Engineering, Cornell University, Ithaca, New York 14853, USA
*Email: {am2352, jls548, dha7}@cornell.edu
†Email: {jrd288, bmh78}@cornell.edu

***Abstract*—Heating ventilation and air-conditioning (HVAC) systems consume a significant portion of the energy within buildings. Current HVAC control systems use simple fixed occupant schedules, while proposed energy optimization schemes do not consider past discomfort in making future energy optimization decisions.**

We propose a Model-based predictive control (MPC) algorithm that adaptively balances energy and comfort while the system is in operation. The algorithm combines occupancy prediction with the history of occupant discomfort to constrain expected discomfort to an allowed budget. Our approach saves energy by dynamically shifting discomfort over time based on its real time performance. The system adapts its behavior according to the past discomfort and thus plays the dual role of saving energy when discomfort is smaller than the target budget, and maintaining comfort when the discomfort margin is small. Simulation results using synthetic benchmarks and occupancy traces demonstrate considerable energy savings over a smart reactive approach while meeting occupant comfort objectives.

Keywords*—*Smart Buildings; Energy-Comfort Optimization; Occupancy Prediction

I. Introduction

Indoor environmental quality systems, such as heating ventilation and air-conditioning (HVAC) and lighting, consume a majority of the energy within buildings [1]. In contrast to replacing systems, better control algorithms present the most practical means of reducing energy consumption and enhancing Indoor Environmental Quality (IEQ). In current practice, algorithms react to changes in the occupancy while using simple fixed schedules to forecast occupancy. More advanced energy optimization schemes do not consider past discomfort in making future energy optimization decisions. This paper presents an approach that optimizes energy with respect to past and future discomfort while meeting a cumulative discomfort goal.

Although many predictive occupancy models have been proposed [2], [3], [4], [5], [6], the ramifications of prediction inaccuracy on future control decisions are not well addressed, especially for systems operating in real time. Furthermore, Model Predictive Control (MPC) has been used to constrain instantaneous discomfort while minimizing energy [7], [8], [9], [10], [11], [12], [13], [14], [15], [16], [17]. Since discomfort is felt over time, constraining instantaneous discomfort does not include occupants' memory of discomfort, which reduces energy savings and perceived occupancy satisfaction. Such systems also fail to dynamically adapt their future energy optimal decisions based on its past discomfort performance. To the best of our knowledge, adaptive energy optimization using memory of past discomfort has not been published before.

In this paper, we propose an MPC-based algorithm that adaptively balances energy and comfort while the system is in operation. Our method is novel because it uses the history of occupant discomfort in combination with occupancy prediction to constrain the expected discomfort to an allowed budget. In principle, our method saves energy by dynamically shifting discomfort over time based on its real time performance in a similar spirit to MPC that shifts loads for economical energy consumption [18], [19]. The system adapts its behavior according to the past discomfort and thus plays the dual role of saving energy when discomfort is smaller than the target budget, and maintaining comfort when the discomfort margin is small. If the accumulated past discomfort exceeds the allowed limit due to occupancy mispredictions, the algorithm automatically corrects the situation in real time while still attempting to optimize for energy.

We evaluate our approach using several synthetic occupancy benchmarks and real occupancy datasets in comparison to three baselines: 1) an energy-efficient reactive scheme that heats up the room whenever the room becomes occupied, 2) a smart reactive algorithm that reacts to the occupancy changes while using a simple fixed schedule to forecast occupancy, and 3) an oracle scheme with perfect knowledge of future occupancy. For predictable occupancy patterns, our algorithm operates close to the perfect prediction scheme with 4-10% energy savings over the smart reactive policy while meeting the allowed discomfort budget. For the irregular occupancy patterns, our method meets the discomfort goal while consuming only 2% more energy than the smart reactive policy.

The rest of this paper is organized as follows. Section II presents a list of variables used in this paper to describe our energy-comfort optimization framework. The overall optimization approach is introduced in Section III followed by the problem formulation in Section IV, and the system architecture of the predictive algorithm in Section V. The experiment setup is described in Section VI and the results in Section VII. Lastly, we present related work in Section VIII and conclude in Section IX.

II. Nomenclature

E	instantaneous energy
Occ	occupancy
d	discomfort density
T_s	set temperature
T_r	room temperature
T_o	outdoor temperature
T_g	ground temperature
T_{upper}	Upper set temperature
T_{lower}	Lower set temperature
T_{int}	Intermediate set temperature
T_h	$\|T_{upper} - T_r\|$ at which $d = 1$, when occupied
T_l	Smallest $\|T_{upper} - T_r\|$ when occupied, below which $d = 0$
$MADD$	Moving Average Discomfort Density
Φ_{MD}	Maximum allowed average discomfort density
$\triangleq$	equal by definition

III. Overall Approach

The overall objective of our proposed approach is to minimize cumulative energy consumption (E) while meeting a maximum discomfort goal over a sliding window of fixed length timesteps. Throughout the paper, we consider the heating season and assume that the controller must set the set temperature, T_s, to a specified comfortable value, T_{upper}, for every timestep that the room is occupied. To save energy, the controller may use a lower specified set temperature, T_{lower}, when the room is unoccupied[1]. The room temperature (T_r) dynamics are a function of the current T_r, ground temperature (T_g), outdoor temperature (T_o) and the current set temperature (T_s). The discomfort at any given timestep is calculated according to the difference $T_{upper} - T_r$ using the model defined in the next section.

At each timestep, the system calculates a Moving Average Discomfort Density ($MADD$) over the last M occupied timesteps. It tries to maximize energy savings while ensuring that the $MADD$ does not exceed a maximum discomfort goal (Φ_{MD}). Intuitively, at any given timestep, the aggressiveness by which the system attempts to save energy in future timesteps depends on recent discomfort as well as future occupancy probabilities. If past discomfort is low and future occupancy is projected to also be low, then the system may try to aggressively reduce energy consumption. If discomfort has been high and future occupancy is also expected to be high, then the system will operate conservatively.

IV. Problem Formulation

Our objective is to minimize cumulative energy while meeting the maximum discomfort constraint at every timestep i. The objective function and constraints are given by Equation 1.

[1]One of our baseline schemes also uses an intermediate temperature, T_{int}.

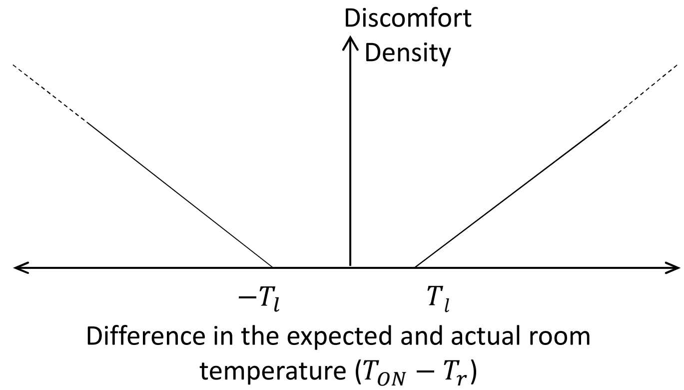

Fig. 1: Discomfort density

$$
\begin{aligned}
\min_{\vec{T_s}} \quad & \sum_{i=1}^{N} E(T_{r,i}, T_{o,i}, T_{g,i}, T_{s,i}) \\
\text{subject to} \quad & \\
& \sum_{j \in \mathcal{J}(i)} d(T_{r,j}) \leq \Phi_{MD} \times \sum_{j \in \mathcal{J}(i)} Occ_j \\
& T_{r,i+1} = f(T_{r,i}, T_{o,i}, T_{g,i}, T_{s,i}) \\
& T_{lower} \leq T_{s,i} \leq T_{upper} \\
& T_{s,i} = T_{lower} \quad \text{if } Occ_i = 1 \\
& \qquad \forall 1 \leq i \leq N
\end{aligned} \tag{1}
$$

$\mathcal{J}$ is the set of occupied periods, and $\mathcal{J}(i)$ is the set of occupied indices containing M occupied periods prior to index i. Occ_i at timestep i has a value of 1 when occupied for any time during the timestep and 0 when unoccupied[2].

The discomfort density, $d(T_{r,j})$, is a function of the difference in the expected room temperature (T_{upper}) and actual room temperature (T_r) when the room is *occupied*[3]. Our discomfort model is inspired from the work of Putta et al. [20], and from similar violation-based discomfort models used in the past [13], [14], [15], [16], [17]. This is shown in Figure 1 and given by Equation 2. Whenever the temperature difference is less than T_l, the discomfort density is zero, else it scales linearly. At the temperature difference of T_h, the discomfort density is one.

$$
d(T_{r,i}) \triangleq \begin{cases} 0, & \text{if } |T_{upper} - T_{r,i}| \leq T_l \\ \frac{(|T_{upper} - T_{r,i}| - T_l)}{(T_h - T_l)}, & \text{otherwise} \end{cases} \tag{2}
$$

V. System Architecture

Figure 2 shows the overall system architecture, which consists of our supervisory Predictive Control coupled to conventional HVAC Control. The Predictive Control attempts to

[2]As discussed later, our system optimizes over a horizon of past timesteps of known occupancies (for which Occ_i is 1 or 0) and future timesteps of unknown occupancies. For these future timesteps, we use expected occupancy values for Occ_i instead of 1 or 0 values.

[3]Our discomfort model could be extended to include other factors such as humidity and indoor air quality (IAQ) violations.

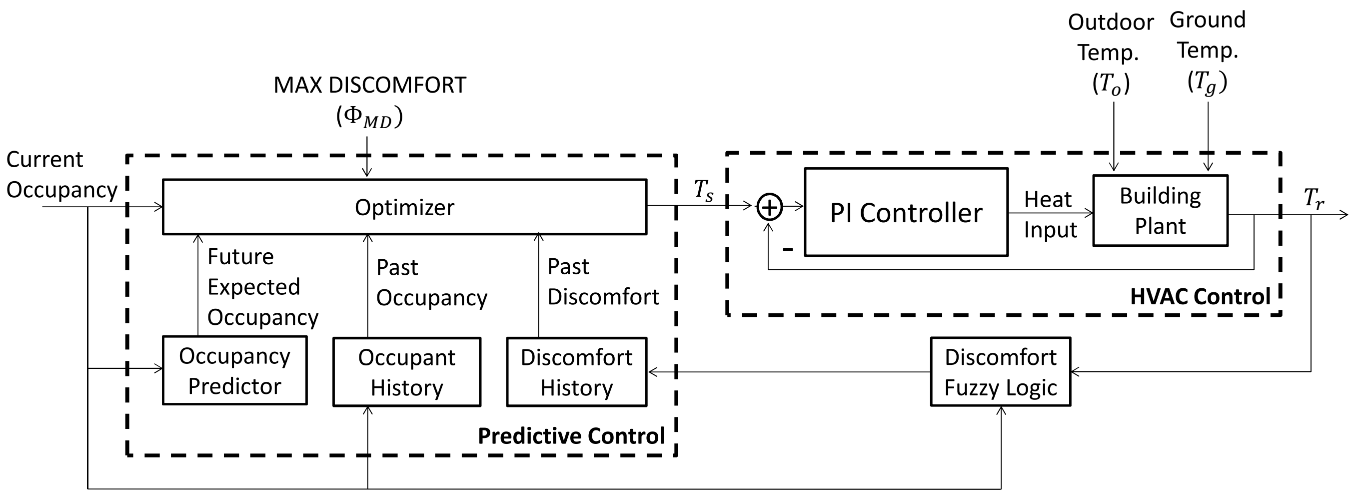

Fig. 2: System architecture

produce an optimum T_s for the HVAC Control. The optimum T_s is the reference input for the HVAC Control to maintain the thermal state of the building. We describe these modules in detail below.

A. HVAC Control

The thermal state of the building zone is maintained by a PI controller. Based on the difference between T_s and the room temperature T_r, the PI controller generates a heat input to the building plant. This is the heating power injected to the building to meet the reference T_s requirement. The building thermal model is a linear time-invariant state-space dynamical system

$$\begin{aligned} x_{k+1} &= Ax_k + Bu_k \\ y_k &= Cx_k \end{aligned} \tag{3}$$

Here, x_k is the temperature state vector containing T_r, and the input vector u_k encompasses the outdoor and ground temperature and the injected heating power. The output vector y_k contains the room temperature T_r. The output matrix C appropriately selects T_r from x_k. The system matrix A contains the building thermal model, while the input matrix B contains the building's response to the applied heat input and weather disturbance.

B. Predictive Control

The Predictive Control runs an optimization algorithm to generate T_s based on the historical discomfort behavior, current and past actual occupancies, expected occupancies for future timesteps, and Φ_{MD}. It comprises an *Occupancy Predictor* to produce future occupancy probabilities and an *Optimizer* to run the optimization.

1) Occupancy Predictor: The Occupancy Predictor uses the past occupancy to predict the expected occupancy at each given timestep in the horizon. The prediction is done on a timestep-by-timestep basis, with the expectation at each timestep computed from past occupancy data from days of the week for which similar occupancy patterns are likely.

For offices and labs, we use two separate occupancy models: one for weekdays and the other for weekends. The sample interval is 15 minutes. Thus, a weekday comprises 96 different expected occupancies. Each of these is calculated as an average of some number of past occupancies, with 1 representing an occupied timestep and 0 an unoccupied one. A weekend day has 96 similarly derived timestep expected occupancies. For meeting rooms, we average data over individual weekdays, under the assumption that occupancy patterns will differ among days of the week. Thus, for an office or lab, the expected occupancy for Monday at 10am is identical to Tuesday at the same time, while these could differ for a meeting room.

2) Optimizer: The Optimizer runs the MPC algorithm for the optimal building HVAC control over a prediction horizon H during the system operation. The MPC algorithm generates a sequence of set points $\overrightarrow{T_s}$ based on the current occupancy, future expected occupancies, and previous discomfort and past occupancies. The optimizer then selects the first element of $\overrightarrow{T_s}$ as a reference for HVAC control. The optimizer assumes accurate outside and ground temperature prediction and uses the building state-space model to compute energy and thermal responses.

At the time index k with the optimizer looking over a horizon H, Equation 1, can be written as

$$\begin{aligned} \min_{\overrightarrow{T_s}} \sum_{i=1}^{k+H} E(T_{r,i}, T_{g,i}, T_{o,i}, T_{s,i}) = & \sum_{i=1}^{k-1} E(T_{r,i}, T_{g,i}, T_{o,i}, T_{s,i}) \\ & + \min_{\overrightarrow{T_s}} \sum_{i=k}^{k+H} E(T_{r,i}, T_{g,i}, T_{o,i}, T_{s,i}) \end{aligned} \tag{4}$$

The optimizer cannot affect the past energy, but can attempt to minimize the cumulative future energy over the horizon k to $k + H$. When the room is occupied, the optimizer does not invoke MPC but simply forces $T_{s,k} = T_{upper}$. However, when $Occ_k = 0$, the optimizer attempts to minimize the cumulative energy while keeping the $MADD$ less than Φ_{MD} at every time index h between k and $k + H$. Equation 5 shows the optimization problem, where $\mathcal{J}(k+h)$ is the set of occupancy indices with M total occupied periods split between the expected occupied periods from $k + 1$ to $k + h$ and the

TABLE I: Simulation parameters

Parameters	Values
T_l	±2°C
T_h	±6°C
T_{upper}	21°C
T_{int}	19°C
T_{lower}	15.6°C
K_p	900
K_i	750
Weather	Winter (January)
Location	Elmira, NY
Simulation Timestep	15 minutes
Horizon Length (H)	4 timesteps
Φ_{MD}	10%
Past Occupied Period (M)	40 timesteps

past occupancy before time k.

$$\min_{\vec{T_s}} \quad \sum_{i=k}^{k+H} E(T_{r,i}, T_{o,i}, T_{g,i}, T_{s,i})$$

subject to

$$\sum_{j \in \mathcal{J}(k+h)} \mathbb{E}[d(T_{r,j})] \leq \Phi_{MD} \times \sum_{j \in \mathcal{J}(k+h)} \mathbb{E}[Occ_j] \quad \forall 1 \leq h \leq H \tag{5}$$

The left-hand side of Equation 5 is the sum of the expected discomfort, which includes actual past discomfort and the predicted future discomfort over the horizon. The right-hand side is the product of Φ_{MD} and the sum of the expected occupancy, which includes the actual past occupancy and the future expected occupancy over the horizon. If the $MADD$ at time index k is very close to Φ_{MD}, the right-hand side is tightened and the optimizer has little room to minimize energy. However, if the gap between the current $MADD$ and Φ_{MD} is large, the right-hand side is relaxed, giving more opportunity for energy minimization. If the inequality is not met, the optimization becomes infeasible and the optimizer sets $T_{s,k} = T_{upper}$.

VI. Experimental Setup

We use the simulation parameters shown in Table I and the building model of Figure 3. We construct the building model using Google Sketchup [21] using realistic materials: a brick exterior, foam-insulated roofing, an insulated concrete slab floor, and double-pane windows. The building model is then converted directly from the CAD geometry and material data to a resistor-capacitor (RC) network using the Sustain framework [22].

Sustain generates a 41-state model encompassing convective and conductive transfer and assumes that the interior air volume is well-mixed. The model does not include radiation.

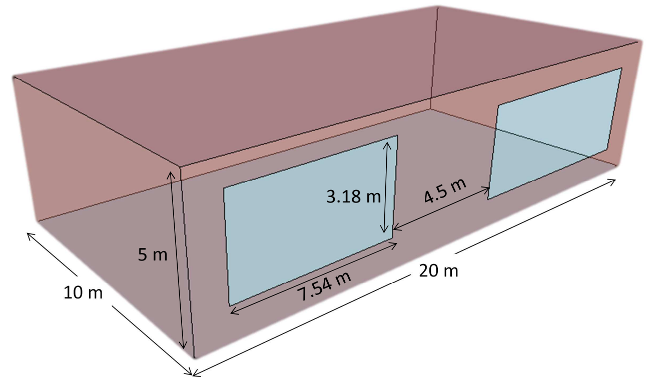

Fig. 3: Single-zone building model

TABLE II: Synthetic occupancy benchmarks

Benchmarks	Periods of Potential Occupancy
No Occupancy	None
One Hour	9:00-10:00, 13:00-14:00, 16:00-17:00, 18:00-19:00
Two Hours	8:00-10:00, 11:00-13:00, 14:00-16:00, 17:00-19:00
Office	8:00-12:00, 13:00-17:00

The exterior walls and floor slab are tied to ambient air and ground temperatures which, during simulation, are obtained from an EnergyPlus weather file.

A. Optimization Software

We use CVX [23] to solve the optimization problem of Equation 5. The discomfort model is a piecewise-affine function. To mimic the optimization formulation of Equation 5, we implement a scalarized multi-objective optimization of energy and discomfort. With discomfort numerically smaller than energy, energy is implicitly optimized with discomfort as a constraint. Therefore, our implementation does not require any scalar parameter to prioritize energy versus discomfort. Furthermore, if CVX is unable to find a feasible solution, the optimizer conservatively sets $T_s = T_{upper}$. This occurs less than 0.3% of the time in our simulations for regular benchmarks and around 10% for irregular data.

B. Occupancy Benchmarks

We evaluate our approach using real occupancy data as well as synthetic benchmarks (Table II). For the latter, we create time periods of potential occupancy, during which the probability of occupancy is 90% for each timestep. The actual occupancy data includes a Graduate Student Office in Duffield Hall at Cornell University and a lab within the Cornell Nanofabrication Laboratory (CNF). The occupancy data for these spaces are recorded by motion and CO_2 sensors, which we convert to 15 minute timesteps denoting whether the room is occupied or not. We then gather this data for a three month period and calculate a occupancy probability for each timestep

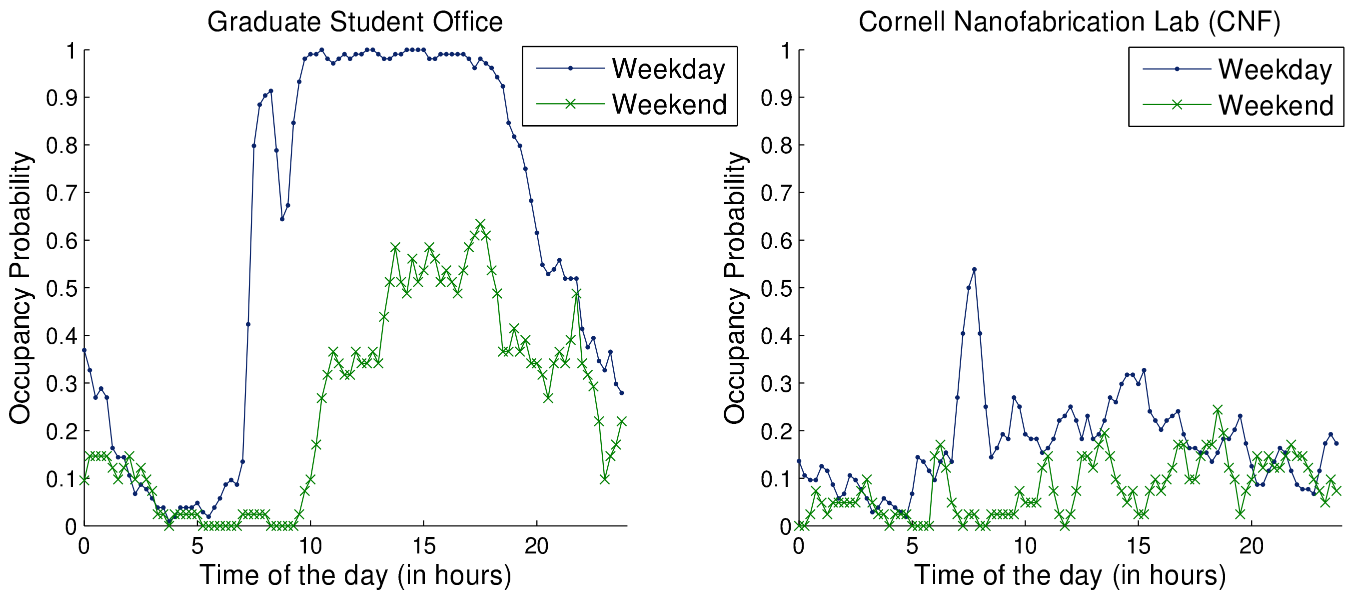

Fig. 4: Duffield occupancy probabilities

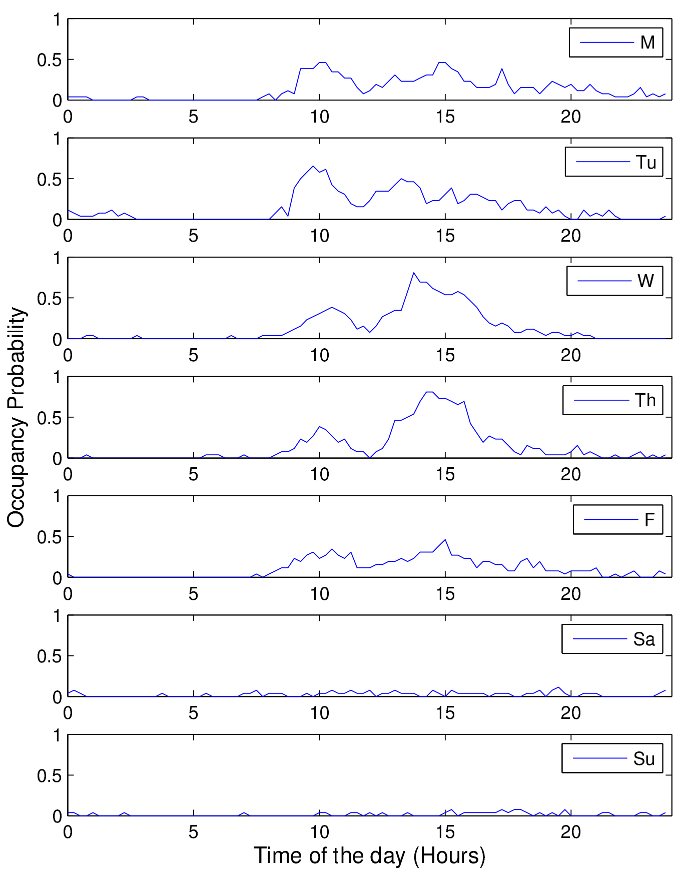

Fig. 5: MERL occupancy probabilities

on both weekdays and weekends. Our assumption for these spaces is that the expected occupancy at a given time of the day, e.g., 10am, will not significantly differ among different weekdays, but could differ considerably between weekdays and weekends.

Figure 4 shows the occupancy probabilities of these two spaces over a 24 hour period. The weekday probabilities of the office are high between 7am to 8pm. Weekends are more irregular with occupancy probabilities reaching around 50% during that time period. The CNF occupancy data is far more irregular, which makes optimization more challenging.

We also use occupancy data from the Mitshubishi Electric Research Laboratory [24] 8-North Conference Room. Motion sensors operate asynchronously; we convert these readings to 15 minute occupancy timesteps. However, since this is a conference room, we calculate occupancy probabilities on a weekday basis. That is, we assume that the expected occupancy at a given time of the day could vary significantly among different weekdays, which is borne out by Figure 5, which shows the probabilities generated from the MERL data over a six month period. For instance, the meeting room has a much higher probability of being occupied on Wednesdays and Thursdays compared to Fridays. As expected, weekend probabilities are low.

C. Baseline Control Schemes

We compare our $predictive$ scheme to three baseline control policies. All policies, including those we propose, immediately set the set temperature to T_{upper} whenever the room becomes occupied. The $reactive$ policy sets the set temperature to T_{upper} whenever the room is occupied and to T_{lower} when it becomes unoccupied. As expected, this approach saves energy but at the cost of an unacceptably high $MADD$. To address this shortcoming, the *smart reactive (SR)* policy sets T_s to an interim temperature T_{int} beginning at 6am in anticipation of impending occupancy, and shifts to

T_{upper} when occupancy is detected. When the room becomes vacant for 30 minutes, it changes the set temperature back to T_{int}. Beginning at 10pm, SR is like $reactive$. The choice of these specific timings is motivated by a typical office and university occupancy schedule that prioritizes comfort over energy in contrast to the reactive scheme. The SR policy improves $MADD$ over $reactive$ but at higher energy cost. Finally, the *perfect prediction (PP)* policy is identical to our scheme with the exception of using perfect knowledge of future occupancy (actual occupancy values from our benchmarks) instead of expected occupancies.

VII. Results

In this section, we present the energy savings and discomfort of our $predictive$ scheme compared to the baselines for the real occupancy and the synthetic benchmarks. Figure 6 illustrates the energy and discomfort of the different control schemes for a representative day within the `Office` benchmark. Beginning at 6am, SR transitions to the intermediate set point T_{int} and then reacts to the first occupant two hours later. After the latest occupancy period, it transitions to T_{int} and eventually to T_{lower}. $Reactive$ reacts similarly to the first occupant at 8AM, but from the T_{lower} set point, thereby impacting comfort. It reacts in a similarly ineffective manner throughout the day, and often exceeds Φ_{MD}. $Predictive$ and PP react more smoothly to occupant activities, keeping within Φ_{MD} by proactively conditioning the room before an occupant arrives, and transitioning to T_{lower} during periods of unoccupancy. The cumulative energy consumption over the course of the day of $predictive$ is 10.3% lower than SR, and is within 0.2% of PP. The energy of $predictive$ is only 1% higher than the $reactive$ scheme that frequently violates the discomfort goal.

Figure 7(a) shows the percent energy savings over SR for the synthetic and real occupancy benchmarks over a period of 25 days, while Figure 7(b) shows the maximum $MADD$. The $reactive$ scheme has the largest energy savings but its $MADD$ often exceeds Φ_{MD}. Over all of the synthetic benchmarks with the exception of `No Occupancy`, the $predictive$ scheme achieves 7-10% energy savings, which is almost identical to PP. For `Graduate Office`, $predictive$ saves 4.5% energy over SR and is within 0.3% of PP.

The results for `MERL` and `CNF` illustrate the limitations of our approach, in particular our occupancy predictor. These benchmarks lack regularity, with `MERL` being less regular and `CNF` highly irregular. The $predictive$ scheme saves only 0.3% energy over SR for `MERL`, and expends 2.3% more energy than SR for `CNF`. PP with its perfect prediction achieves over 10% savings for both benchmarks. Our occupancy prediction averaging scheme at times leads the controller to anticipate occupancy during unoccupied periods, and to wrongly predict unoccupancy, thereby violating the discomfort constraint. In the latter situation, $predictive$ corrects the situation by acting more conservatively in future timesteps, which wastes energy.

This behavior is illustrated in Figure 8 for the benchmark `MERL`. When the $MADD$ is smaller than Φ_{MD} (on the 11^{th} day for instance), the $predictive$ scheme reduces energy consumption. However, due to the irregularity of the benchmark, the $MADD$ increases beyond Φ_{MD} on the 13^{th}

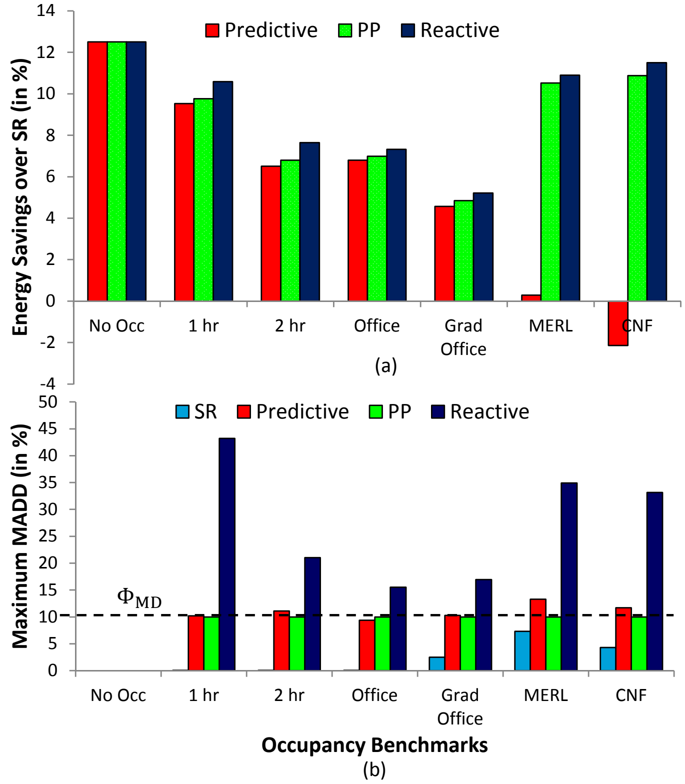

Fig. 7: Energy and $MADD$ for all benchmarks over 25 days

day. In this situation, the $predictive$ scheme automatically corrects the situation and switches its role to maintain comfort by conservatively spending more energy to bring the $MADD$ below Φ_{MD}. The maximum $MADD$ for `MERL` goes upto 0.13 for the $predictive$ scheme when simulated for 25 days.

Figure 9 shows the adaptive behavior of the $predictive$ scheme for the `CNF` benchmark. `CNF` is highly irregular and the $MADD$ rarely goes below the Φ_{MD} limit. Thus the $predictive$ scheme rarely saves energy because it must keep the $MADD$ below the Φ_{MD} limit, and thereby consumes more energy than the baseline SR. For `CNF`, the $MADD$ reaches as high as 0.12 for the $predictive$ when simulated for 25 days. These results highlight the need for more accurate occupancy prediction, the subject of our future work.

A. Daily Performance

Figure 10 shows a histogram of the energy savings over SR for `Office` and `Graduate Student Office`. For both benchmarks, the daily energy profile of $predictive$ closely resembles that of PP and $reactive$. For `Graduate Student Office`, all policies–even PP–consume more energy (negative energy savings) over SR for around 9 days. For these days, the room is occupied from early morning to late night, a near perfect fit to the SR schedule. However, there are few short periods of vacancy for which all schemes try to optimize, resulting in dropping the T_s and thus expending more energy when the heat must be turned up again. Future work includes exploring longer horizons and heuristic solutions to these short vacancy periods.

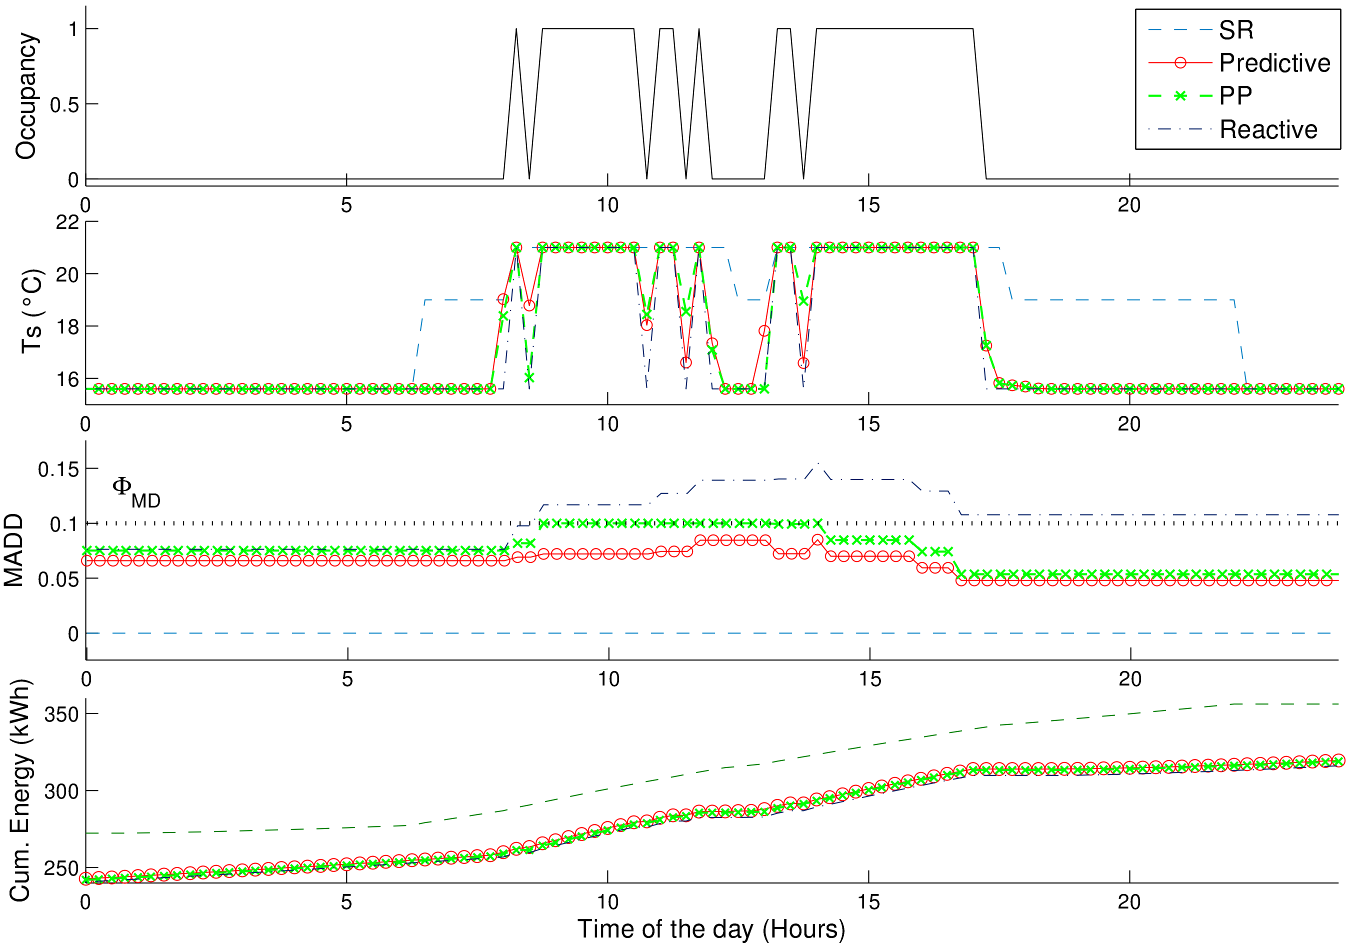

Fig. 6: Energy and discomfort of different algorithms for `office` occupancy data

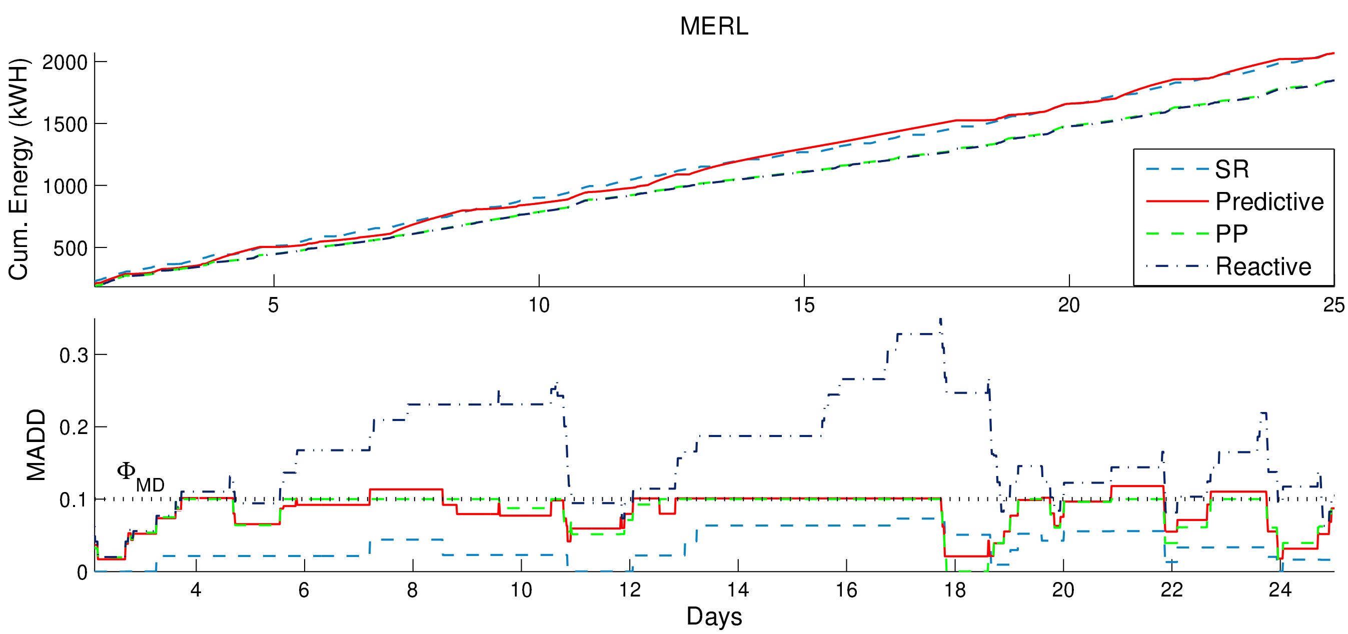

Fig. 8: Runtime adaptation of energy and discomfort for `MERL`

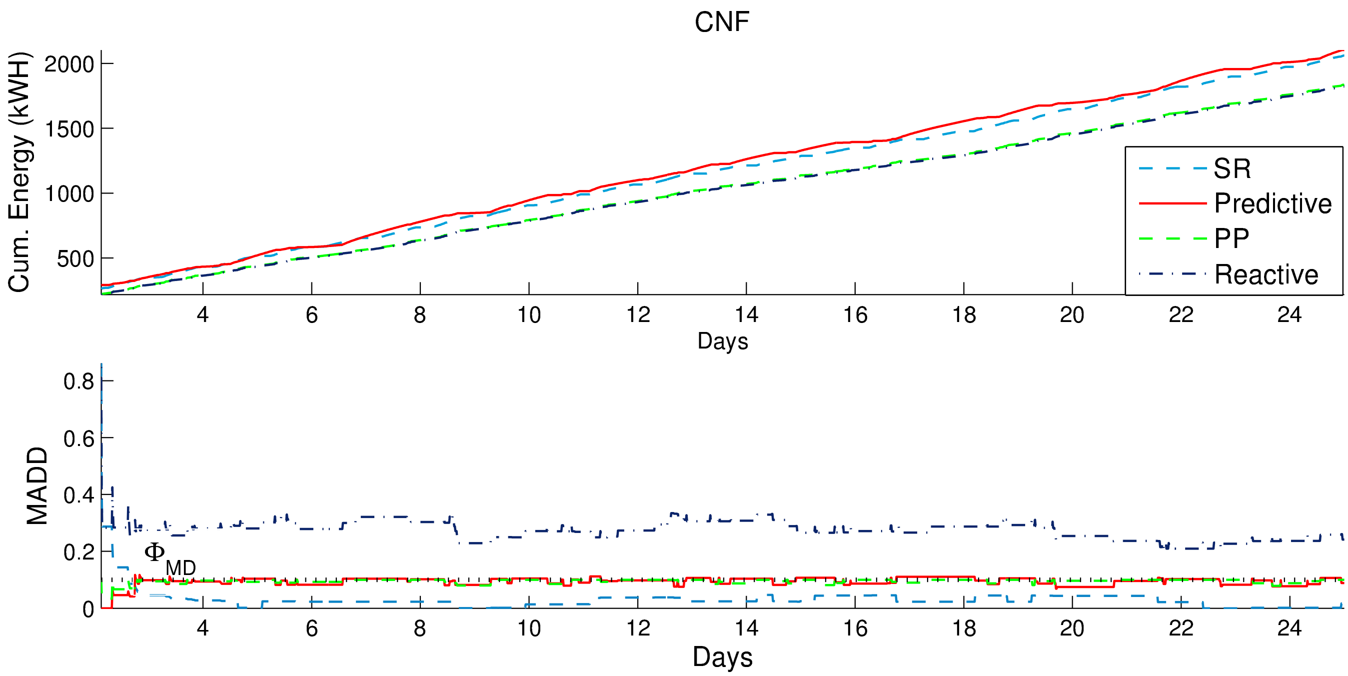

Fig. 9: Runtime adaptation of energy and discomfort for `CNF`

VIII. Related Work

HVAC energy optimization using the high-accuracy prediction models developed from the sensor data such as [24], [25] have been proposed before [3], [16], [17], [26], [27], [28], [29], [30], [31], [32], [33], [34], [35], [36], [37]. However, none of these works consider the ramification of occupancy mispredictions on energy and discomfort during actual system operation. In this paper, we examine the performance of our method for both regular and irregular occupancy captured from commercial buildings. We demonstrate that our scheme continuously corrects itself for the irregular occupancy behavior by automatically switching to the comfort maintenance mode whenever the discomfort margin is small.

Several other works use instantaneous discomfort values to balance energy and discomfort. For instance, [7], [17], [20] use MPC algorithms to jointly optimize energy and comfort by using a scalar parameter to prioritize energy versus comfort, while [12], [38], [39] use heuristic algorithms. Both approaches require operators to manually tune the parameters at the time of synthesis and the parameters, once fixed, cannot adapt to the dynamic behavior of the system during operation. Our method does not rely on any parameter tuning and dynamically adapts its behavior based on the past discomfort performance. Furthermore, minimizing energy by constraining the instantaneous discomfort to a certain limit is an alternative approach considered in [10], [11], [13], [40], [41], [42]. Since discomfort is felt over time, constraining instantaneous discomfort is not a true representation of occupant comfort. This also presents fewer opportunities to make future energy-efficient decisions based on the past energy-discomfort performance. To the best of our knowledge, adaptive energy optimization using memory of past discomfort has not been published before. Our method optimizes energy based on the history of occupant discomfort, and dynamically switches roles between energy-saving and comfort maintenance.

IX. Conclusion

We present a MPC-based algorithm that uses occupancy prediction and past occupant discomfort to meet a discomfort objective while optimizing energy efficiency. The system dynamically adjusts how aggressively it attempts to save energy in future timesteps based on recent discomfort as well as expected future occupancy. Our results show the potential for large energy savings over a smart reactive approach while meeting occupant comfort goals.

References

[1] DOE, "2010 buildings energy data book, table 3.1.5, US Department of Energy." http://buildingsdatabook.eren.doe.gov/docs/DataBooks/2010_BEDB.pdf, Mar. 2011. Accessed: 03/15/2012.

[2] C. Liao, Y. Lin, and P. Barooah, "Agent-based and graphical modeling of building occupancy," *Journal of Building Performance Simulation*, vol. 5, pp. 5–25, Jan. 2012.

[3] V. L. Erickson, Y. Lin, A. Kamthe, R. Brahme, A. Surana, A. E. Cerpa, M. D. Sohn, and S. Narayanan, "Energy efficient building environment control strategies using real-time occupancy measurements," in *Proceedings of the First ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings*, BuildSys '09, pp. 19–24, ACM, 2009.

[4] S. Meyn, A. Surana, Y. Lin, S. Oggianu, S. Narayanan, and T. A. Frewen, "A sensor-utility-network method for estimation of occupancy in buildings," in *Proceedings of the 48th IEEE Conference on Decision and Control, 2009 held jointly with the 2009 28th Chinese Control Conference. CDC/CCC 2009.*, pp. 1494–1500, 2009.

[5] J. Page, D. Robinson, N. Morel, and J.-L. Scartezzini, "A generalised stochastic model for the simulation of occupant presence," *Energy and Buildings*, vol. 40, no. 2, pp. 83 – 98, 2008.

[6] D. Wang, C. C. Federspiel, and F. Rubinstein, "Modeling occupancy in single person offices," *Energy and Buildings*, vol. 37, no. 2, pp. 121–126, 2005.

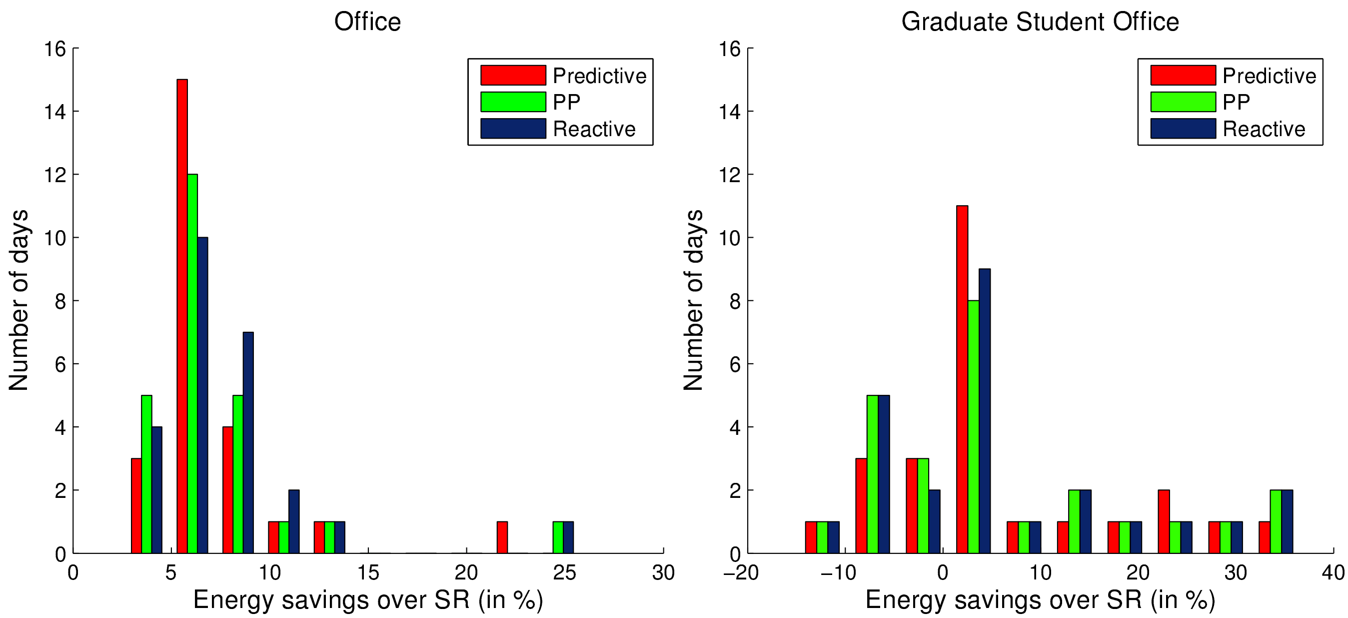

Fig. 10: Histogram of daily energy savings for `Office` and `Graduate Student Office` benchmarks

[7] R. Z. Freire, G. H. Oliveira, and N. Mendes, "Predictive controllers for thermal comfort optimization and energy savings," *Energy and Buildings*, vol. 40, no. 7, pp. 1353 – 1365, 2008.

[8] P. X. Gao and S. Keshav, "SPOT: A smart personalized office thermal control system," in *ACM e-Energy13*, 2013.

[9] L. Klein, G. Kavulya, F. Jazizadeh, J.-y. Kwak, B. Becerik-Gerber, P. Varakantham, and M. Tambe, "Towards optimization of building energy and occupant comfort using multi-agent simulation," in *International Symposium on Automation and Robotics in Construction*, 2011.

[10] M. Nowak and A. Urbaniak, "Utilization of intelligent control algorithms for thermal comfort optimization and energy saving," in *2011 12th International Carpathian Control Conference (ICCC)*, pp. 270–274, 2011.

[11] S. Ari, I. A. Cosden, H. Khalifa, J. Dannenhoffer, P. Wilcoxen, and C. Isik, "Constrained fuzzy logic approximation for indoor comfort and energy optimization," in *Fuzzy Information Processing Society, NAFIPS*, pp. 500–504, 2005.

[12] N. Nassif, S. Kajl, and R. Sabourin, "Evolutionary algorithms for multi-objective optimization in HVAC system control strategy," in *Fuzzy Information, 2004*, vol. 1, pp. 51–56, 2004.

[13] S. Goyal, H. A. Ingley, and P. Barooah, "Occupancy-based zone-climate control for energy-efficient buildings: Complexity vs. performance," *Applied Energy*, vol. 106, no. C, pp. 209 – 221, 2013.

[14] A. Kusiak, F. Tang, and G. Xu, "Multi-objective optimization of HVAC system with an evolutionary computation algorithm," *Energy*, vol. 36, no. 5, pp. 2440 – 2449, 2011.

[15] J. Scott, A. Bernheim Brush, J. Krumm, B. Meyers, M. Hazas, S. Hodges, and N. Villar, "Preheat: Controlling home heating using occupancy prediction," in *Proceedings of the 13th International Conference on Ubiquitous Computing*, UbiComp '11, (New York, NY, USA), pp. 281–290, ACM, 2011.

[16] J. Lu, T. Sookoor, V. Srinivasan, G. Gao, B. Holben, J. Stankovic, E. Field, and K. Whitehouse, "The smart thermostat: Using occupancy sensors to save energy in homes," in *Proceedings of the 8th ACM Conference on Embedded Networked Sensor Systems*, SenSys '10, (New York, NY, USA), pp. 211–224, ACM, 2010.

[17] J. R. Dobbs and B. M. Hencey, "Predictive HVAC control using a Markov occupancy model," in *American Control Conference*, 2014.

[18] Y. Zong, L. Mihet-Popa, D. Kullmann, A. Thavlov, O. Gehrke, and H. Bindner, "Model predictive controller for active demand side management with PV self-consumption in an intelligent building," in *2012 3rd IEEE PES International Conference and Exhibition on Innovative Smart Grid Technologies (ISGT Europe)*, pp. 1–8, Oct 2012.

[19] E. Vrettos, K. Lai, F. Oldewurtel, and G. Andersson, "Predictive control of buildings for demand response with dynamic day-ahead and real-time prices," in *2013 European Control Conference (ECC)*, pp. 2527–2534, July 2013.

[20] V. Putta, G. Zhu, D. Kim, J. Hu, and J. E. Braun, "A distributed approach to efficient model predictive control of building HVAC systems," in *International High Performance Buildings Conference*, (Purdue University, IN, USA), Purdue e-Pubs, 2012.

[21] "Trimble Sketchup." http://sketchup.google.com.

[22] D. Greenberg, K. Pratt, B. Hencey, N. Jones, L. Schumann, J. Dobbs, Z. Dong, D. Bosworth, and B. Walter, "Sustain: An experimental test bed for building energy simulation," *Energy and Buildings*, vol. 58, no. 3, pp. 44 – 57, 2013.

[23] I. CVX Research, "CVX: Matlab software for disciplined convex programming, version 2.0 beta." http://cvxr.com/cvx, Sept. 2012.

[24] C. Wren, Y. Ivanov, D. Leigh, and J. Westhues, "The MERL motion detector dataset: 2007 workshop on massive datasets," tech. rep., 2007.

[25] Y. Agarwal, B. Balaji, R. Gupta, J. Lyles, M. Wei, and T. Weng, "Occupancy-driven energy management for smart building automation," in *Proceedings of the 2Nd ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Building*, BuildSys '10, (New York, NY, USA), pp. 1–6, ACM, 2010.

[26] C. Liao, Y. Lin, and P. Barooah, "Agent-based and graphical modelling of building occupancy," *Journal of Building Performance Simulation*, vol. 5, no. 1, pp. 5–25, 2012.

[27] S. Meyn, A. Surana, Y. Lin, S. M. Oggianu, S. Narayanan, and T. A. Frewen, "A sensor-utility-network method for estimation of occupancy in buildings," in *Decision and Control, 2009 held jointly with the 2009 28th Chinese Control Conference. CDC/CCC 2009. Proceedings of the 48th IEEE Conference on*, pp. 1494–1500, IEEE, 2009.

[28] L. Vu, Q. Do, and K. Nahrstedt, "Jyotish: Constructive approach for context predictions of people movement from joint wifi/bluetooth trace," *Pervasive Mob. Comput.*, vol. 7, pp. 690–704, Dec. 2011.

[29] R. H. Dodier, G. P. Henze, D. K. Tiller, and X. Guo, "Building occupancy detection through sensor belief networks," *Energy and buildings*, vol. 38, no. 9, pp. 1033–1043, 2006.

[30] V. L. Erickson and A. E. Cerpa, "Occupancy based demand response HVAC control strategy," in *Proceedings of the Second ACM Workshop*

on Embedded Sensing Systems for Energy-Efficiency in Buildings, pp. 7–12, ACM, 2010.

[31] J. Page, D. Robinson, N. Morel, and J.-L. Scartezzini, "A generalised stochastic model for the simulation of occupant presence," *Energy and Buildings*, vol. 40, no. 2, pp. 83–98, 2008.

[32] K. P. Lam, M. Höynck, B. Dong, B. Andrews, Y. Chiou, R. Zhang, D. Benitez, J. Choi, *et al.*, "Occupancy detection through an extensive environmental sensor network in an open-plan office building," *IBPSA Building Simulation*, pp. 1452–1459, 2009.

[33] J. Hutchins, A. Ihler, and P. Smyth, "Modeling count data from multiple sensors: A building occupancy model," in *Computational Advances in Multi-Sensor Adaptive Processing, 2007. CAMPSAP 2007. 2nd IEEE International Workshop on*, pp. 241–244, Dec 2007.

[34] S. Mamidi, Y.-H. Chang, and R. Maheswaran, "Improving building energy efficiency with a network of sensing, learning and prediction agents," in *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems - Volume 1*, AAMAS '12, (Richland, SC), pp. 45–52, 2012.

[35] Z. Yang, N. Li, B. Becerik-Gerber, and M. Orosz, "A systematic approach to occupancy modeling in ambient sensor-rich buildings," *SIMULATION*, 2013.

[36] J. Du, W. Guo, and R. Wang, "Tourism room occupancy rate prediction based on neural network," in *Advances in Neural Networks ISNN 2007*, vol. 4493 of *Lecture Notes in Computer Science*, pp. 80–84, Springer Berlin Heidelberg, 2007.

[37] Z. Yang, N. Li, B. Becerik-Gerber, and M. Orosz, "A multi-sensor based occupancy estimation model for supporting demand driven HVAC operations," in *Proceedings of the 2012 Symposium on Simulation for Architecture and Urban Design*, SimAUD '12, (San Diego, CA, USA), pp. 2:1–2:8, Society for Computer Simulation International, 2012.

[38] V. L. Erickson, M. Carreira-Perpiñán, and A. E. Cerpa, "OBSERVE: Occupancy-based system for efficient reduction of HVAC energy," in *Information Processing in Sensor Networks (IPSN), 2011 10th International Conference on*, pp. 258–269, IEEE, 2011.

[39] A. Kusiak, G. Xu, and F. Tang, "Optimization of an HVAC system with a strength multi-objective particle-swarm algorithm," *Energy*, vol. 36, no. 10, pp. 5935 – 5943, 2011.

[40] J. Xu, P. B. Luh, P. William, E. Blankson, R. Jerdonek, and K. Shaikh, "An optimization-based approach for facility energy management with uncertainties," in *Proceedings of IFAC 2002*, 2002.

[41] F. Oldewurtel, A. Parisio, C. N. Jones, D. Gyalistras, M. Gwerder, V. Stauch, B. Lehmann, and M. Morari, "Use of model predictive control and weather forecasts for energy efficient building climate control," *Energy and Buildings*, vol. 45, no. 0, pp. 15 – 27, 2012.

[42] B. Dong and K. Lam, "A real-time model predictive control for building heating and cooling systems based on the occupancy behavior pattern detection and local weather forecasting," *Building Simulation*, vol. 7, no. 1, pp. 89–106, 2014.