

Automatic HVAC control with real-time occupancy recognition and simulation-guided model predictive control in low-cost embedded system



Muhammad Aftab, Chien Chen, Chi-Kin Chau*, Talal Rahwan

Khalifa University of Science and Technology, Masdar Institute, United Arab Emirates

ARTICLE INFO

Article history:

Received 22 February 2017
Received in revised form 6 June 2017
Accepted 27 July 2017
Available online 16 August 2017

Keywords:

Automatic HVAC control
Embedded system
Occupancy recognition
Model predictive control

ABSTRACT

Intelligent building automation systems can reduce the energy consumption of heating, ventilation and air-conditioning (HVAC) units by sensing the comfort requirements automatically and scheduling the HVAC operations dynamically. Traditional building automation systems rely on fairly inaccurate occupancy sensors and basic predictive control using oversimplified building thermal response models, all of which prevent such systems from reaching their full potential. Such limitations can now be avoided due to the recent developments in embedded system technologies, which provide viable low-cost computing platforms with powerful processors and sizeable memory storage in a small footprint. As a result, building automation systems can now efficiently execute highly sophisticated computational tasks, such as real-time video processing and accurate thermal-response simulations. With this in mind, we designed and implemented an occupancy-predictive HVAC control system in a low-cost yet powerful embedded system (using Raspberry Pi 3) to demonstrate the following key features for building automation: (1) real-time occupancy recognition using video-processing and machine-learning techniques, (2) dynamic analysis and prediction of occupancy patterns, and (3) model predictive control for HVAC operations guided by real-time building thermal response simulations (using an on-board EnergyPlus simulator). We deployed and evaluated our system for providing automatic HVAC control in the large public indoor space of a mosque, thereby achieving significant energy savings.

© 2017 Elsevier B.V. All rights reserved.

1. Introduction

Heating, ventilation, and air-conditioning (HVAC) units, which are a primary target of building automation, make up almost 50% of the energy consumed in both residential and commercial buildings [1]. In general, building automation systems aim to intelligently control building facilities in response to dynamic environmental factors, while maintaining satisfactory performance in energy consumption and comfort. The primary functions of a building automation system include: (1) sensing of the environmental factors by measurements, and (2) optimizing control strategies based on the current and predictive states of building and occupancy. These tasks require an integrated process of sensing, computation, and control.

Traditional building automation systems rely on fairly inaccurate occupancy sensors, which hinder the responsiveness of automation systems. For example, passive infrared and ultra-sound occupancy sensors produce poor accuracy, because they are unable to determine the occupancy state adequately when occupants remain stationary for a prolonged period of time. They also have a limited range which hinders their performance, especially in a large area. More accurate sensing technology, such as cameras that use visible or infra-red lights, can significantly improve the accuracy of occupancy recognition.

On the other hand, model predictive control, by which the future thermal response and external environmental factors are anticipated to make control decisions accordingly, has been considered in a number of studies [2–7] which are shown to be more effective than classical PID and hysteresis controllers that do not consider anticipated events. However, these studies are often based on time-invariant, first-principle linear models (also known as lumped element resistance-capacitance (RC) models [8]), considering only simple building geometry and single-zone in near-future time horizon. Although these linear models are easier for calibration (e.g.,

* Corresponding author.

E-mail addresses: muhaftab@masdar.ac.ae (M. Aftab), cchen@masdar.ac.ae (C. Chen), ckchau@masdar.ac.ae (C.-K. Chau), trahwan@masdar.ac.ae (T. Rahwan).

using frequency domain decomposition, or subspace system identification methods [8–10]), the error accumulates considerably when a longer time horizon is considered in model predictive control. While non-linear models are rather complicated and impractical, other alternatives based on physical models of building thermal response can provide a feasible solution.

Recently, there have been remarkable advances in embedded system technologies, which provide low-cost platforms with powerful processors and sizeable memory storage in a small footprint. In particular, the emergence of *system-on-a-chip* technology [11], which integrates all major components of a computer into a single chip, can provide versatile computing platforms with low-power consumption and mobile network connectivity in a cost-effective manner for mass production. As a result, smartphones are able to rapidly evolve from single-core to multi-core processors with a low incremental production cost. Notably, the *Raspberry Pi* project [12], which originally aimed to provide affordable solutions for the teaching of computer science, has rapidly evolved for a wide range of advanced scientific projects. Therefore, there are plenty of opportunities to harness recent embedded system technologies in intelligent building automation systems. Particularly, sophisticated computational tasks can be conducted on these embedded systems efficiently, such as real-time video processing and accurate building thermal response simulation (e.g., [13]).

With this in mind, we designed and implemented an occupancy-predictive HVAC control system in a low-cost yet powerful embedded system (using Raspberry Pi 3) for building automation. The rest of the paper is organized as follows. In Section 2, we present the background information and literature review. In the remaining sections, we highlight three key features of our system.

Real-time video-based occupancy recognition (Section 3). We apply advanced video-processing techniques to analyze the features of occupants from video cameras, and automatically classify and infer the states of occupancy. Moreover, we consider privacy enhancement using a frosted lens. Our system achieves 80–90% accuracy for occupancy recognition by real-time video processing. Furthermore, we improve the performance of our occupancy recognition by using Machine Learning for considerably crowded settings.

Dynamic occupancy prediction (Section 4). We employ various linear and non-linear regression models to capture and predict occupancy trends according to different day-of-week, seasonal patterns, etc. We present general as well as domain-specific approaches for occupancy prediction. Our models are able to identify the future occupancy trends considering a variety of dynamic usage patterns.

Simulation-guided model predictive control (Sections 5 and 6). We employ *EnergyPlus* simulator [14] for real-time HVAC control. We ported *EnergyPlus* simulator to the Raspberry Pi embedded system platform for simulation-guided model predictive control. A co-simulation framework is utilized to provide accurate building thermal response simulation under proper calibration. Noteworthy, we also release our Raspberry Pi version of *EnergyPlus* publicly [15] to enable other researchers to take advantage of our work for future building automation projects.

Our automatic HVAC control system is intended for public indoor spaces, such as corridors, libraries, or communal areas. Unlike private spaces such as homes, these public indoor spaces are not controlled by a particular occupant and can be affected by a diverse set of occupancy patterns. Such patterns tend to vary more dynamically in public spaces compared to private ones, posing challenges for effective occupancy sensing and prediction systems.

In particular, our system is deployed and evaluated for providing automatic HVAC control in the large public indoor space of a mosque (see Fig. 1), which is the worship place for followers of Islam. Typically, mosques have large public spaces, and are open

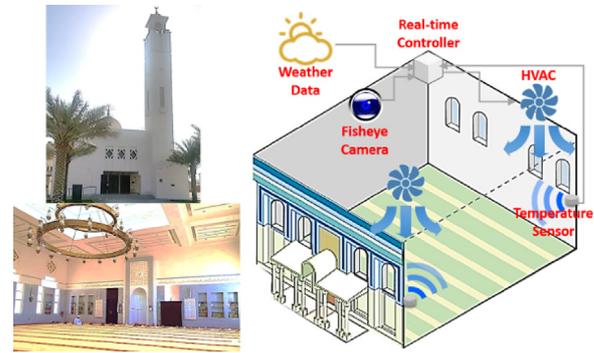


Fig. 1. A large public indoor space of a mosque is used as a testbed for our automatic HVAC control system. Fisheye video camera, temperature and humidity sensors, as well as real-time controller for HVAC have been deployed in our testbed.

24-h a day and 7-days a week. There are nearly 5000 mosques in the UAE [16], and over 55,000 mosques in Saudi Arabia [17]. Due to the hot climate in this region, HVAC is required on a regular basis. The results obtained from our testbed implementation in Section 7 demonstrate the significant energy savings that can be achieved by using automatic HVAC control systems in public indoor spaces.

2. Background and related work

In this section, we present the background and related work of our system. In particular, this section consists of two subsections. The first provides a brief review of occupancy recognition and prediction. The second subsection compares three main approaches of model predictive HVAC control.

2.1. Occupancy recognition and prediction

Any object with a temperature higher than perfect zero emits heat in the form of radiation. The conventional passive infra-red (PIR) sensors can be used to detect a certain wavelength when a person is near the sensor. Previous papers [18–20] demonstrated the possibility of applying this to control lighting systems. However, the disadvantage of the PIR sensor becomes evident when the occupant remains stationary for a certain period of time. In particular, PIR sensors are designed to detect changes in the movement, and if a person remains stationary in front of the sensor, then as far as the sensor is concerned, there will be no change in the movement, leading to an erroneous observation.

Video-based occupancy-detection algorithms can provide better accuracy than their PIR counterparts. One such algorithm was proposed in [21], whereby certain features (such as edges, textures, etc.) are extracted from the video and used in a regression model to estimate the number of occupants. Unfortunately, this algorithm is computationally intensive, rendering it unsuitable for implementations on embedded systems such as Raspberry Pi. In contrast, our algorithm can detect occupancy in real-time, even when implemented on an embedded system, as is the case with our testbed. An alternative algorithm was proposed in [22]. Although this algorithm is computationally less intensive than the previous one, it nevertheless relies heavily on identifying the heads of the occupants. This head-detection process is particularly challenging in our testbed, as it is common for a typical occupant to have a head cover indistinguishable from his or her outfit. Since our system does not require head detection, it is insensitive to whether or not the occupant is wearing a head cover.

In addition to occupancy recognition, occupancy prediction is required to anticipate building usage and control pre-cooling in advance. A variety of techniques to predict occupancy have been proposed in the literature, including statistical analysis, machine

learning, and stochastic modeling. A comprehensive review of occupancy prediction techniques is provided in [23,24].

2.2. Model predictive control

There are three major approaches to model predictive control in the literature:

- *LTI model predictive control*: This uses a linear time-invariant (LTI) mathematical model, which is a simplified thermal dynamics model considering only a near-future short time horizon. A common approach is to use an RC model to capture the first-order heat transfer dynamics. It is suitable for simple settings, such as a single zone with simple building geometry. There are usually a small number of parameters in the model. LTI model predictive control is explored in [8–10,25,26].
- *Non-linear model predictive control*: Many real-world systems exhibit rich non-linearity. There are many general non-linear mathematical models of system dynamics, such as Volterra series, neural networks and NARMAX models. However, most non-linear models require a large parameter space, which is difficult to calibrate from measurements. The use of non-linear models for predictive control is investigated in [27,28].
- *Simulation-guided model predictive control*: In this approach, model predictive control is guided by real-time physical model simulators considering future anticipated events. For buildings, there are a number of simulators, such as *EnergyPlus* and *TRNSYS*, that are much more accurate than LTI models, and also easier to calibrate than general non-linear models. Reviews of different building simulators and their merits are presented in [5,29].

For the above control approaches to be effective, it is crucial to calibrate the model parameters so that the model response is consistent with the empirical data. Several model calibration methods have been proposed in the literature. These methods can be broadly categorized as *manual* or *automated*. In particular, *manual* approaches require the modeler to intervene repeatedly and make adjustments, whereas *automated* approaches use mathematical and statistical models to automate the calibration process. A review of model calibration can be found in [30].

Several previous studies used simulation programs to facilitate model predictive control. Specifically, in [5,7,31], the authors employed co-simulation for model predictive control (MPC) with *EnergyPlus*. However, these studies did not implement the MPC models in real-world HVAC systems and were limited to simulations. Other studies [32,33] tested the MPC models with real-world HVAC systems, but relied on powerful desktop computers running costly numerical computation software such as MATLAB.

There are other studies using occupancy prediction for MPC in HVAC systems. For example [34,25,35,26], used machine learning to predict occupancy patterns based on environmental sensor data. More specifically [34,35] used predicted occupancy patterns to simulate HVAC control in *EnergyPlus*, whereas [25,26] developed LTI MPC algorithms for HVAC control in real buildings. Furthermore [36], investigated the potential benefits of occupancy information for HVAC control, but their investigation was only limited to simulations.

The differences between our study and the aforementioned studies are: (1) we implemented our MPC algorithm in a real-world testbed using free-software platforms and low-cost embedded systems; (2) we employed *EnergyPlus* for real-time simulation-guided MPC of HVAC systems; (3) we developed a comprehensive solution that integrates occupancy recognition, occupancy prediction and simulation-guided MPC, thereby demonstrating the viability of automatic HVAC control using low-cost embedded systems.

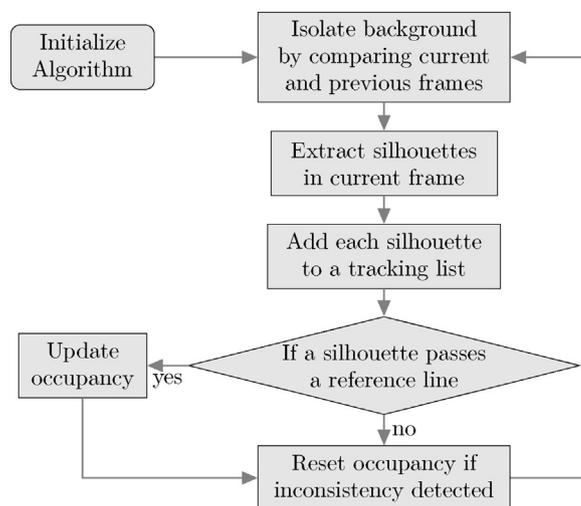


Fig. 2. Flowchart of our occupancy-recognition algorithm.

3. Occupancy recognition

Occupancy information is crucial for many applications, such as building management and human behavior studies. We develop an occupancy recognition system based on real-time video processing of a video stream to infer the occupancy patterns dynamically. This raises a number of challenges. First, structures differ from one building to another, leading to the possibility of occupants being obscured by various obstacles, such as pillars. Hence, we need to track the movements of occupants to determine the occupancy more accurately. Second, our algorithm is executed on an embedded system (e.g., Raspberry Pi), which has limited processing power and memory space compared to a typical desktop computer; this is particularly challenging since the typical video-processing algorithms are computationally demanding.

The basic idea of our occupancy-recognition algorithm is to count the number of people crossing a virtual reference line in the video, captured by a fisheye camera. Objects are identified as moving blobs (i.e., a set of connected points whose position is changing during the video stream); every such blob is interpreted as a person. Whenever a moving blob is detected, the algorithm keeps track of its movement to determine whether it passes the virtual reference line, and then updates the number of occupants accordingly. In particular, we position the line near the entrance of the space. Whenever the moving blob passes inward, the total number of occupants is increased by 1, and whenever the moving blob passes outward, the total number of occupants is decreased by 1.

Our algorithm consists of the following five steps: (1) background isolation, (2) silhouette detection, (3) object tracking, (4) inward/outward logging, and (5) inconsistency resolution. The flowchart of our algorithm is presented in Fig. 2. In our implementation, two open-source projects were used: *OpenCV* [37] and *openFramework* [38]. Next, we will explain each step in details.

3.1. Background isolation

The purpose of background isolation is to identify a background image in the current video frame. Note that the appearance of the background may vary over the course of the video, depending on the time-of-day (e.g., turning on the lights at night may significantly alter the appearance of the background compared to natural light). The shadows of occupants must also be taken into consideration during the background isolation. To overcome these challenges, we employ the Gaussian mixture-based background/foreground

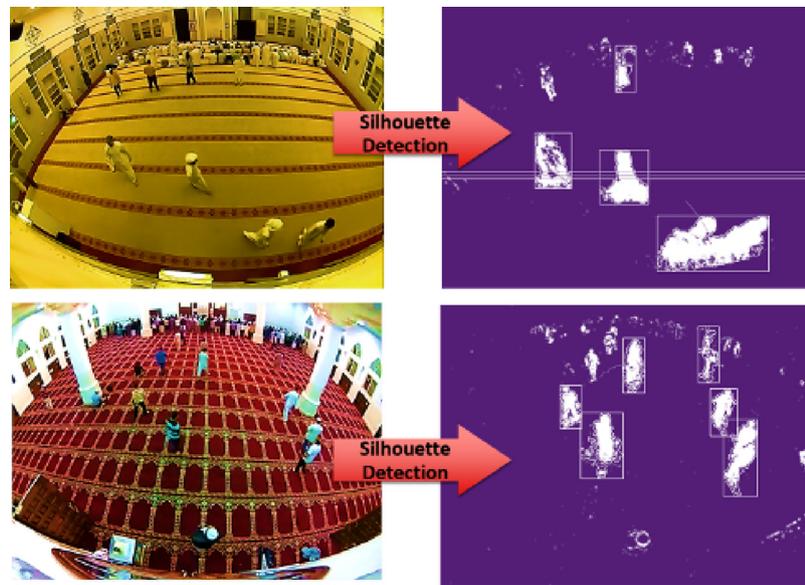


Fig. 3. Examples of silhouette detection, taken from different buildings in which our system is deployed.

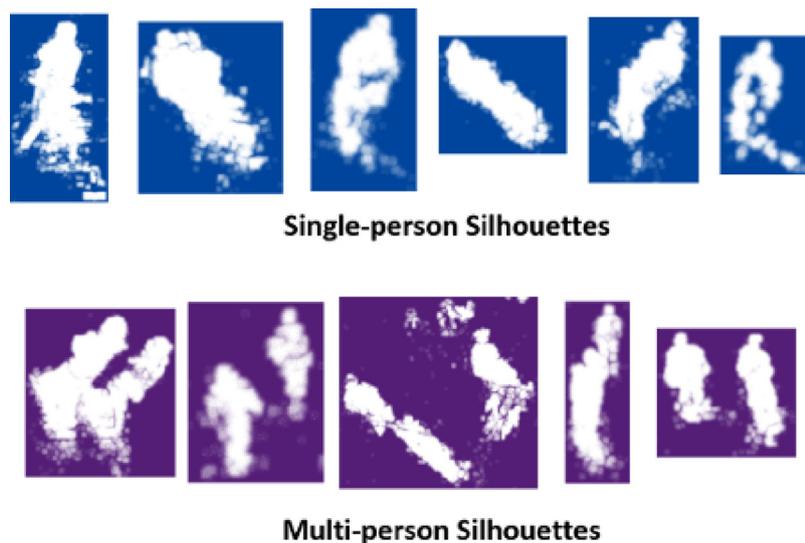


Fig. 4. Sample images from the training dataset used for machine learning and image classification.

segmentation algorithm proposed in [39,40], and implemented on OpenCV.

3.2. Silhouette detection

In our setting, the term “silhouette” is used to refer to the border of a set of continuous points. Silhouette detection is based on the algorithm proposed in [41]. The silhouettes of moving objects are extracted by comparing the current frame with the previous one; see Fig. 3 for some examples. Here, only the silhouettes larger than a certain threshold area, A , are considered. The threshold is adjusted depending on the viewpoint and orientation of the camera. Furthermore, we use different values of A for different parts of the space, to reflect the fact that individuals appear smaller as they move farther away from the camera. Importantly, the silhouettes of two or more people may overlap, and may, therefore, be interpreted as only one individual. To overcome this challenge, we use

a machine-learning technique which will be explained later on in Section 3.6.

3.3. Object tracking

After obtaining the silhouettes of occupants, their bounding boxes are logged for tracking. To this end, the list T_a of silhouettes from the last frame is maintained. The algorithm then obtains the list T_b of silhouettes from the current frame. For each bounding box B_b in T_b , the algorithm determines whether there exists a box B_a in T_a that is within a certain distance, d , from B_b . If so, then B_b is assumed to be the new location of B_a . Consequently, B_a is updated in T_a , and its location is set to be that of B_b , in preparation for the next iteration of the algorithm. Note that the silhouettes are detected when they are moving. However, in cases where people might stop for a few seconds and then continue to walk, those objects will be

removed from T_a when they are missed for a certain number of seconds.

3.4. Inward/outward logging

Given the collection of moving objects and their locations, a virtual reference line is used to count occupants. Specifically, whenever the locations are updated, the algorithm checks every object to determine whether that object has crossed from one side of the line to the other. If so, then the number of occupants is updated according to the direction of the movement. For inward movement, the number of occupants increases by 1, otherwise it decreases by 1. Furthermore, since the silhouettes of any two moving objects may overlap, the width of the bounding box can be used to infer the number of occupants contained in each object.

3.5. Inconsistency resolution

With all the techniques described thus far, the performance of the algorithm may not be satisfactory, due to one major challenge: *the silhouettes of different occupants may overlap*. In this case, some of the overlapping occupants may go undetected by the algorithm. This problem becomes even more evident when the movement patterns are affected by whether the occupants are entering or leaving the building, e.g., due to the fact that occupants arrive one by one, but leave all at once. For instance, in our application domain, the pace at which people leave is significantly faster than the pace at which they enter. Consequently, when all occupants leave at once, the number of overlapping silhouettes increases, leading to a larger number of people going undetected by the algorithm. As a result, the algorithm on average misses more occupants leaving than entering the building, leading to the erroneous conclusion that there are still occupants in the building when in fact there are none.

To resolve such inconsistency, two simple techniques are used. First, if the number of occupants becomes negative, the occupant counter is frozen until another object crosses the reference line inward. Second, if no moving object is detected for a certain period of time, the occupant counter is reset to zero. While these simple techniques reduce the error, they are clearly insufficient. In Section 3.6, we propose a dedicated technique to address this issue.

3.6. Improvement by machine learning

As mentioned earlier, one of the major challenges that we have encountered while deploying our occupancy-detection algorithm is to resolve the problem of overlapping silhouettes. One solution is based on the width of bounding box; the wider the box, the more occupants it contains. However, we observe that such a solution is insufficient, especially when an occupant happens to be directly in front of another. To resolve this issue, we employ machine-learning and image-classification techniques, coupled with randomized principal component analysis (PCA) [42], which is implemented in [43].

In more detail, we collected 13,000 blobs from video footages spanning a period of one week and segmented each such blob as a separate image to create our training dataset. The number of occupants in each image was manually identified; see Fig. 4 for some examples. The images were then transformed to grayscale in order to reduce the computational load. After that, the images were rescaled to the same size, 30×15 pixels, thereby making the size of each image 450 pixels. Second, randomized PCA is used to project the pixels in the original array to a smaller array that preserves the characteristics of the images, as a set of 25 features for each image in this case. The projection aims to further reduce the computational complexity. Moreover, we added the original width, height and the ratio of black pixels to the set of features. After obtaining

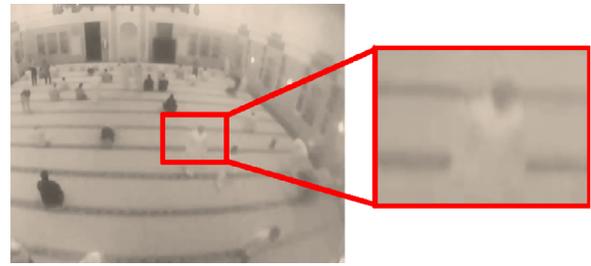


Fig. 5. Enhanced privacy by using frosted lens on camera.

Table 1

Results of occupancy recognition given different durations and different numbers of occupants.

Video length	Total no. of occupants	AccuracyRate
40 min	127	90%
40 min	154	90%
20 min	407	84%

the features, Gaussian Naive Bayes is used to classify the blobs with respect to the number of occupants.

3.7. Privacy enhancement

In our system, the videos and images are discarded immediately after processing. Still, privacy invasion is always a concern for video based detection methods. With this in mind, we introduced a number of techniques to preserve the privacy of the occupants. First of all, to prevent any potential hacker from accessing the video feed, we ensured that the camera stream can only be accessed by a single process at a time. Consequently, when our system is running, all the other programs are blocked from accessing the camera. In addition, a hardware-based solution has been tested. Specifically, we created a frosted lens by overlaying a semi-transparent layer in front of the camera, so as to blur the camera feed. In so doing, the faces of occupants are no longer recognizable; see Fig. 5 for an illustration.

3.8. Evaluation results

In this section, we evaluate the accuracy of our occupancy recognition. To this end, we collected video footages from our testbed, each with a frame size of 800×600 pixels, and a sample rate of 30 frames per second. The true occupancy was obtained by manually counting the number of occupants in each frame. Due to this laborious process, we were only able to obtain the true occupancy in a small number of videos, which are used in the evaluations.

We start by defining the accuracy rate as follows:

$$\text{AccuracyRate} \triangleq 1 - \frac{\text{Missing}_{\text{in}} + \text{Missing}_{\text{out}}}{\text{Total}_{\text{in}} + \text{Total}_{\text{out}}} \quad (1)$$

where Total_{in} is the number of individuals who entered the building and $\text{Missing}_{\text{in}}$ is the number of such individuals who were undetected by the algorithm. Conversely, $\text{Total}_{\text{out}}$ is the number of individuals who exited the building and $\text{Missing}_{\text{out}}$ is the number of those who exited without being detected by our algorithm.

Table 1 presents the system accuracy given different numbers of occupants over different durations. As can be seen, the algorithm is able to recognize the number of occupants with high accuracy, even when over 400 individuals enter the building in just 20 min. Naturally, given a greater rate at which occupants enter the building, the accuracy of the system decreases because there are more cases in which the occupants overlap (in many such cases, it was hard, even for a human, to determine the exact number of occu-

Table 2
Results of occupancy recognition comparing weekend, weekday, and Friday.

Type	Video length	AccuracyRate
Weekend	1 day	88%
Weekday	1 day	86%
Friday	1 day	81%

Table 3
Results of occupancy recognition comparing normal lens and frosted lens.

Type	Video length	Total no. of occupants	AccuracyRate
Normal lens	160 min	322	87.8%
Frosted lens	160 min	339	80.2%

Table 4
Results of occupancy recognition comparing with and without machine-learning technique.

Type	Precision	Recall	F1-Score
Without machine learning	0.97	0.27	0.42
With machine learning	0.69	0.78	0.73

pants, and the video had to be replayed several times before the human was able to determine the true occupancy with certainty).

Next, we evaluate our system over one-day periods during the weekend, weekday and Friday. Note that in our testbed Friday is the busiest day in the week due to the Friday sermon, which takes place only once a week, and typically attracts a large audience. The results are presented in Table 2. As expected, the system accuracy correlates with the occupancy rates. In particular, the system is least accurate on Friday (which is the busiest day in our experiment), and most accurate during the weekend (which is the least busy in our experiment). Importantly, the accuracy seems sufficiently high throughout the week to provide a reasonable approximation of the actual occupancy state in the building, which is arguably sufficient for our purpose of HVAC control.

We now turn our attention to quantifying the loss in accuracy that occurs when using our privacy-preserving frosted lens from Section 3.7. The results of this evaluation are presented in Table 3. As can be seen, the frosted lens only reduces the accuracy slightly, and that is despite the fact that the video footage is considerably blurred, as we have shown earlier in Fig. 5.

Next, we evaluate the effectiveness of our machine-learning technique from Section 3.6. To this end, we use three performance measures that are widely used in the machine-learning literature. The first is Precision, which is defined as the fraction of occupants that were correctly classified out of all those who were classified as either moving inward or as moving outward. The second measure is Recall, defined as the fraction of occupants that were correctly classified out of all those who actually entered the building or exited it. Since it is often possible to have a naive classifier that has a high Precision but a low Recall or vice versa, a better metric called F1-Score has been utilized in [44,45], which is basically a harmonic mean of Precision and Recall. Based on those three measures, Table 4 shows the results before and after applying our machine-learning technique. The evaluation is carried out using 10-fold cross-validation of our dataset of 13,000 blobs. As can be seen, according to the most important measure, namely F1-Score, our machine-learning technique from Section 3.6 significantly improves the performance of our system.

Finally, to better understand how occupancy changes during the daytime in our application, we plotted in Fig. 6 the actual occupancy, as well as the occupancy detected by our algorithm, given a typical Friday, and a typical Saturday. As can be seen, the general occupancy trend is clearly captured by the algorithm.

4. Occupancy prediction

This section describes our methodology for predicting the future occupancy of the building. Such predictions can be very helpful when optimizing the HVAC control, especially in an application like ours, where occupants arrive in large numbers over short periods of time (see Fig. 6). For example, being able to predict the arrival of a large number of people allows the system to pre-cool the building in anticipation of their arrival. Likewise, by predicting that all the occupants will shortly be leaving the building (e.g., if a certain social event was coming to an end), the system can turn off the HVAC system before the occupants even start departing.

In Section 4.1 we propose a general-purpose approach to occupancy prediction. After that, in Section 4.2, we further develop our occupancy prediction to produce a domain-specific approach, tailored to our application. Lastly, in Section 4.3 we evaluate our domain-specific prediction by quantifying the impact that it makes on the overall performance of our system.

4.1. General approach

As a starting point, let us consider a linear regression model, trained using all past occupancy data; let us denote such an approach by $\widehat{LR}_{AllData}$. To take this simple approach one step further, let us extend the linear regression model by incorporating any “special events” that may cause an abrupt change in the occupancy trend. Any such special event may be recurrent, with a slightly flexible timing and duration. For instance, consider a hall that can be booked in advance for social activities; here every such activity can be thought of as a *special event*. Information about a special event (such as the average timing and duration, for example) may be available *a priori* or may be inferred from the past occupancy data. The resulting approach, whereby special events are explicitly modeled, will be denoted by \widehat{LR}_{SpEv} , where SpEv stands for *Special Event*. We suggest using this approach with the following linear regression model:

$$y^{(t)} = \beta_0 + \beta_1 x_{pastEvent}^{(t)} + \hat{\beta}_2 x_{nextEvent}^{(t)} + \beta_3 x_{specialCase}^{(t)} \quad (2)$$

where $\hat{y}^{(t)}$ is the target variable (i.e., the predicted occupancy at time t); $\beta_i: i \in \{0, 1, 2, 3\}$ are the parameters of the model; $x_{pastEvent}^{(t)}$ is the difference in time between t and the timing of the *past event*; and likewise $x_{nextEvent}^{(t)}$ is the difference between t and the timing of the *next event*; $x_{specialCase}^{(t)}$ is a binary variable that takes a value of 1 if t coincides with a special case and takes a value of 0 otherwise (an example of such a binary variable is $x_{newYearEve}^{(t)}$, which indicates whether t coincides with new year’s eve). Of course, Eq. (2) is only meant as an example of the possible models that one could choose in order to explicitly model the special events that affect the occupancy. One may instead use other features, depending on the application at hand.

4.2. Domain-specific approach

In this section, we tailor our general-purpose approach from Section 4.1 to our testbed, i.e., the mosque. In this particular application, the *special events* are the five daily prayers in Islam, the timing of which depends on the position of the sun in the sky. More specifically, the daily prayers are held (1) at dawn; (2) at midday right after the sun passes its highest; (3) at the late part of the afternoon; (4) just after sunset; and (5) between sunset and midnight. As such, the prayer times vary over the course of the year, because days tend to be longer in summer and shorter in winter. Importantly, the “Friday prayer” (which takes place every Friday at midday) is preceded

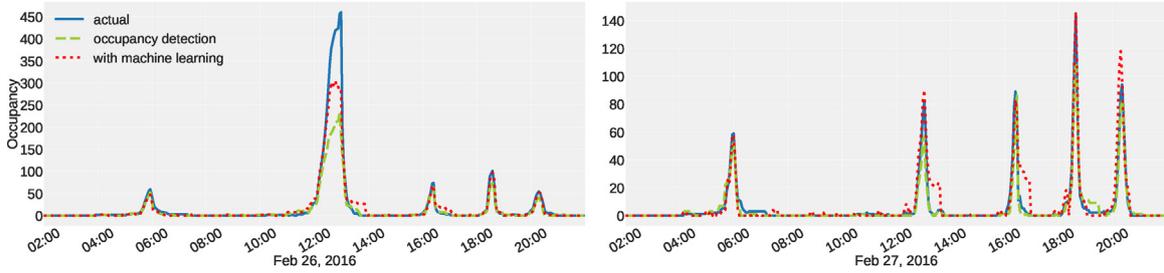


Fig. 6. How occupancy changes during a typical day. Each peak represents one of the five daily prayers in Islam. We distinguish between Friday and other days of the week due to the Friday sermon, which precedes the midday prayer and usually attracts many more worshippers compared to any other prayer throughout the week (note that the scale of the y-axis is greater in the left plot than in the right plot).

by a sermon, the attendance of which is obligatory for Muslims. As such, the number of occupants increases significantly compared to any other prayer throughout the week.

With this in mind, we now modify the general-purpose approach from Section 4.1 by incorporating our domain-specific knowledge of the application at hand. To this end, we consider (i) the difference in minutes between the current time and the *past* prayer, (ii) the difference in minutes between the current time and the *next* prayer, (iii) the current day of the week, and (iv) whether or not it is a public holiday. Then, to predict the occupancy at any given time, \bar{t} , we consider historical data for which the timestamp is within a certain threshold from \bar{t} . This threshold varies according to time of the year, to reflect the constantly shifting prayer times. Based on these features, we propose the following linear regression model:

$$\hat{y}^{(t)} = \beta_0 + \beta_1 x_{\text{pastPrayer}}^{(t)} + \beta_2 x_{\text{nextPrayer}}^{(t)} + \beta_3 x_{\text{holiday}}^{(t)} + \beta_4 x_{\text{dayOfWeek}}^{(t)} \quad (3)$$

where $x_{\text{pastPrayer}}^{(t)}$ is the difference between t and the timing of the *past* prayer; $x_{\text{nextPrayer}}^{(t)}$ is the difference between t and the timing of the *next* prayer; $x_{\text{holiday}}^{(t)}$ is a binary variable indicating whether or not t coincides with a public holiday; and $x_{\text{dayOfWeek}}^{(t)}$ is the day of the week at time t . The resulting approach will be denoted by $\hat{\text{LR}}_{\text{DomSp}}$, where *DomSp* stands for *Domain Specific*. In addition to this linear regression model, we also experimented with a polynomial regression model, denoted by $\hat{\text{PR}}_{\text{DomSp}}$, which uses the same features as those used in $\hat{\text{LR}}_{\text{DomSp}}$.

4.3. Evaluation results

We use two standard metrics to evaluate the effectiveness of the proposed prediction models: *R*-squared and *RMSE* (Root Mean Squared Error). Both of these metrics are always between 0 and 1. Importantly, with *R*-squared the *greater* the value the better the model, whereas with *RMSE* the *smaller* the value the better the model.

As a baseline model, we use a naive approach, denoted by *LastWeek*, whereby the current week's occupancy is assumed to be identical to last week's occupancy (e.g., the occupancy pattern in, say, the coming Monday is assumed to be identical to that of last Monday). The results are shown in Fig. 7, where all models are evaluated with a forecasting horizon of 24-h ahead. As can be seen, even *LastWeek* (the most naive approach) outperforms $\hat{\text{LR}}_{\text{AllData}}$ (the general-purpose linear regression model trained using all past occupancy data).

In contrast, the performance of $\hat{\text{LR}}_{\text{SpEv}}$ (the linear regression model whereby the special events are explicitly modeled) appears

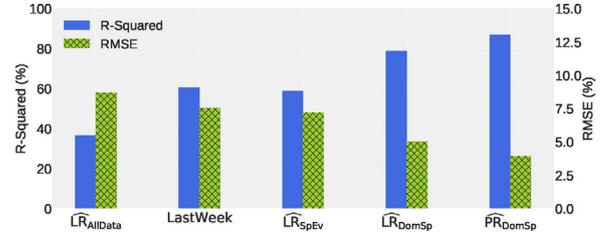


Fig. 7. Results of occupancy prediction.

to be on par with that of *LastWeek*. As for the domain-specific approaches, namely $\hat{\text{LR}}_{\text{DomSp}}$ and $\hat{\text{PR}}_{\text{DomSp}}$, they outperform the other alternatives due to three main reasons: (i) they take into account the day of the week, which is particularly important since the weekly sermon takes place every Friday; (ii) they are able to recognize public holidays, during which the occupancy typically increases in residential areas and decreases in commercial areas; (iii) they are trained using data from only the past 30 days (as opposed to using all past occupancy data), which is important since the prayer times are constantly shifting throughout the year; this shift is usually negligible over a 30-day period, but can be significant over the course of a year. Perhaps not surprisingly, between the two domain-specific approaches, $\hat{\text{PR}}_{\text{DomSp}}$ outperforms $\hat{\text{LR}}_{\text{DomSp}}$. Based on this evaluation, we adopt $\hat{\text{PR}}_{\text{DomSp}}$ as our model of choice. With an *R*-squared value of about 0.87 and an *RMSE* value of about 0.03, this model appears to capture the overall occupancy trend to a satisfactory degree for the purpose of HVAC control.

5. Building thermal response simulation

Traditional building modeling and control is often based on simplified mathematical building models due to their tractability and ease of use. However, these simplified models (e.g., first-principle linear models) can capture only limited aspects of the dynamic nature of buildings and systems. On the contrary, sophisticated building simulation software can use more realistic physical models, which in turn provide accurate modeling of the thermal behavior of buildings. One prominent example is *EnergyPlus*, a popular building energy simulation program that can create detailed building envelope and system models. *EnergyPlus* can perform extensive conductivity analyses, and can even consider local weather conditions, by either incorporating user-supplied information or by using *EnergyPlus Weather Files* [46]. Furthermore, it provides interfaces with which external tools and programs can inter-operate. This process is known as *co-simulation*, whereby different subsystems are integrated to carry out the simulation and calibration process simultaneously. With *co-simulation*, multiple simulators and software tools can be coupled in such a way that the strengths of each tool are exploited while overcoming their

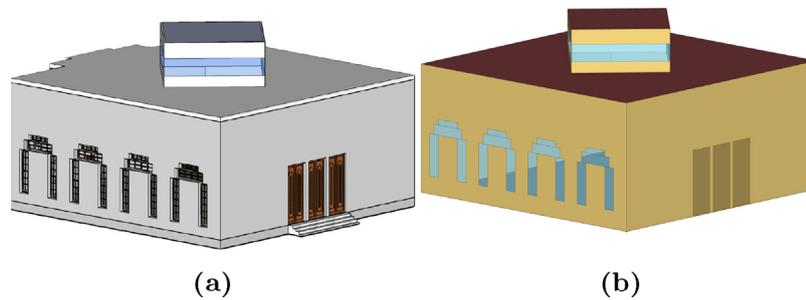


Fig. 8. Basic building models. (a) 3D SketchUp building model and (b) the corresponding OpenStudio model.

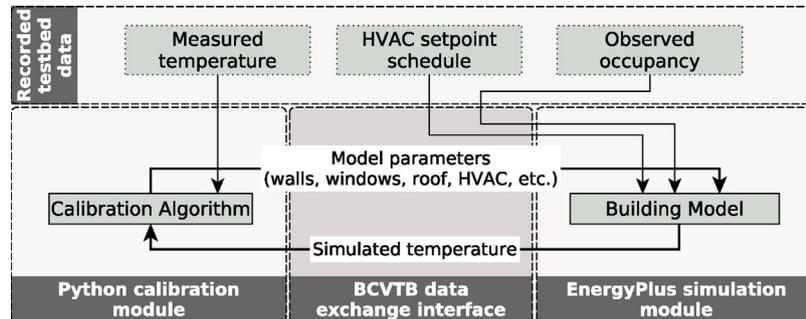


Fig. 9. Co-simulation framework using BCVTB.

Table 5
Description of the testbed building.

Dimensions	18 m × 18 m × 7 m	
Walls	Outer layer	Stucco
	Middle layer	Concrete
	Inner layer	Gypsum
Doors	Wood stile and rail	
Windows	Glass with vinyl framing	
Packaged rooftop HVAC	Cooling capacity	175,840 W
	Air flow rate	7.56 m ³ /s

individual weaknesses. Currently, EnergyPlus is commonly used for planning and energy auditing. Using EnergyPlus for real-time HVAC control presents both opportunities and challenges. In this section, we utilize a framework of co-simulation in order to adapt EnergyPlus for real-world HVAC control.

5.1. Basic building geometry

First, the basic geometry of the building is created for EnergyPlus. In our testbed, we consider the large indoor space of a mosque installed with packaged rooftop HVAC units with ceiling-based cool air distribution. The detailed descriptions of the testbed, including dimensions and HVAC system specifications are provided in Table 5.

We employ SketchUp for 3D modeling with the default material properties for walls, windows, and doors. Then, the 3D model is imported to OpenStudio, where the parameters of the HVAC system are incorporated into the model, based on the vendor's specifications and datasheet. The SketchUp building model and the corresponding OpenStudio model are visualized in Fig. 8. Apart from the basic geometry and properties of the building, many parameters still need to be calibrated according to the available data; this will be the focus on the next subsection.

5.2. Building model calibration

To obtain an accurate simulation of the thermal response of the building, we must first calibrate the parameters of the building

model so as to minimize the gap between the simulated indoor temperature and the measured one. To this end, we use an inverse calibration approach, whereby the model's outputs are used to calibrate the parameters of the model. This calibration process is carried out using the co-simulation framework illustrated in Fig. 9. In particular, the framework couples the EnergyPlus building model with our calibration algorithm which we implemented using Python; this coupling is done using the BCVTB (Building Controls Virtual Test Bed) software tool [47], which provides a data-exchange interface between EnergyPlus and Python. Moreover, BCVTB allows EnergyPlus to take into consideration the real-world data measured from our testbed, such as observed occupancy and temperature.

Out of the numerous parameters that can be calibrated in EnergyPlus, we focus on a particular set of uncertain parameters as our candidates for calibration, because they are either unknown or likely to deviate from the datasheet. Specifically, these parameters are the cooling capacity and air flow rate of the HVAC system, as well as the thickness and conductivity of the roof, the walls, and the windows.

Now, we will explain how the calibration process is done. To this end, we use a dedicated algorithm, the flowchart of which is depicted in Fig. 10. The algorithm is based on the notion of gradient descent, where values of those parameters are updated iteratively until the error between the simulated indoor temperature and the measured temperature is within an acceptable threshold. Next, we explain each step of this algorithm.

Let $X = \{x_1, x_2, \dots, x_K\}$ denote the set of parameters to be calibrated. For every $x_k \in X$, the algorithm runs a series of N co-simulations, updating the model in every iteration using a different value of x_k while keeping the values of the remaining parameters unchanged. Specifically, in the i th iteration of this process (where $i \in \{1, \dots, N\}$) the error is calculated as the difference between the measured temperature and the simulated model temperature; this error is denoted by:

$$\varepsilon(x_1, \dots, x_{k-1}, x_k^i, x_{k+1}, \dots, x_K)$$

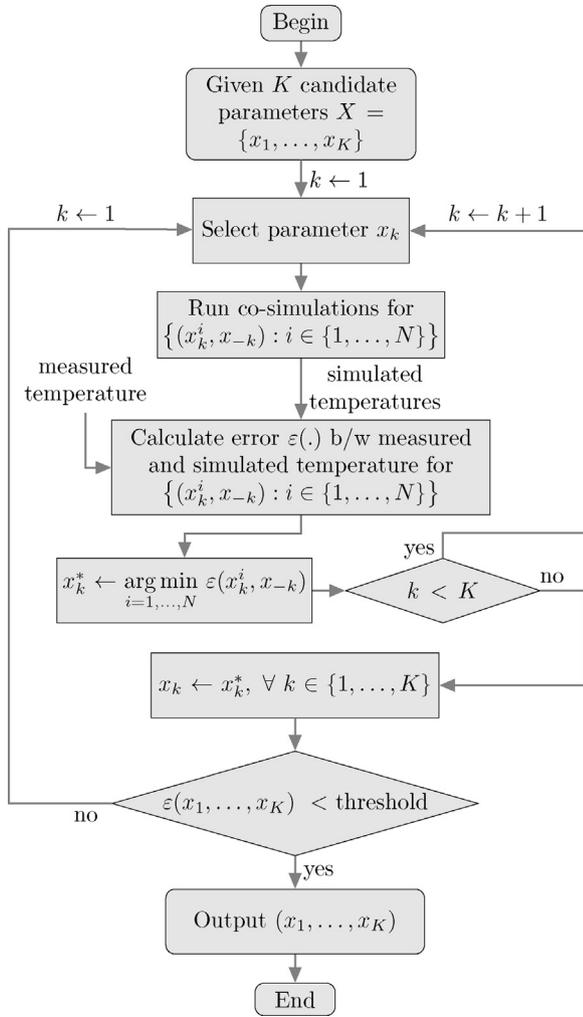


Fig. 10. Flowchart of the model-calibration algorithm.

where x_k^i takes the value of x_k in the i th iteration. For notational convenience, let us write (x_k^i, x_{-k}) instead of writing $(x_1, \dots, x_{k-1}, x_k^i, x_{k+1}, \dots, x_K)$. Then the optimal value for x_k , denoted by x_k^* , can be computed as follows:

$$x_k^* = \arg \min_{i=1, \dots, N} \varepsilon(x_k^i, x_{-k}) \quad (4)$$

After computing x_k^* for every $k \in \{1, \dots, K\}$, the algorithm checks whether $\varepsilon(x_1^*, x_2^*, \dots, x_K^*)$ is within the acceptable threshold. If so, then it outputs $(x_1^*, x_2^*, \dots, x_K^*)$ and terminates; otherwise it repeats the entire process but after updating the parameters, i.e., after setting $x_k \leftarrow x_k^*$ for every $k \in \{1, \dots, K\}$. We note that the outcome of the calibration process is not affected by the order in which the algorithm iterates over the parameters.

5.3. Evaluating the model-calibration algorithm

To evaluate the algorithm from Fig. 10, we use two standard metrics, namely: the Coefficient of Variation of Root Mean Squared Error (CVRMSE) and the Mean Bias Error (MBE); these are computed as follows:

$$\text{CVRMSE} = \frac{\sqrt{\sum_{t=1}^M \left[\frac{(T_t - \hat{T}_t)^2}{M} \right]}}{\frac{1}{M} \sum_{t=1}^M T_t} \quad (5)$$

$$\text{MBE} = \frac{\sum_{t=1}^M (T_t - \hat{T}_t)}{\sum_{t=1}^M T_t} \quad (6)$$

where T_t and \hat{T}_t are the measured and simulated temperature values at time t , respectively, and M is the total number of simulation time steps. Both MBE and CVRMSE provide different insights to the calibration process. However, the CVRMSE is arguably superior because unlike MBE it does not suffer from the cancellation effect [48].

The default and calibrated values of our model parameters are listed in Table 6. Note that we do not manually specify the allowed parameter ranges. EnergyPlus performs automatic parameter range checking using Input Data Dictionary [49], which specifies the maximum and/or minimum values of parameters.

Using a simulation period of 24 h, we tested the temperature response of our model before and after the calibration process (we used 24 h since it is adequate for our purpose of HVAC control). Specifically, the calibration process takes around 100 co-simulations on average, each of which uses sub-hourly data with a sampling time period of 10 minutes. A CVRMSE of less than 2.0% was treated as an acceptable calibration threshold. The results are depicted in Fig. 11. In particular, Fig. 11(a) plots the actual temperature as well as the simulated temperature, given the default parameter values from SketchUp. In Fig. 11(b), the default parameter values are replaced by the ones calibrated using our algorithm from Section 5.2. As can be seen, the simulated temperature of the calibrated EnergyPlus model closely matches the measured temperature. Quantifying the performance gains using the aforementioned metrics, we find that the algorithm achieves $\text{MBE} < 1\%$ and $\text{CVRMSE} < 1\%$. This seems satisfactory, bearing in mind that the ASHRAE guidelines, which require that $\text{MBE} < 10\%$ and $\text{CVRMSE} < 30\%$ for hourly data, and $\text{MBE} < 5\%$ and $\text{CVRMSE} < 15\%$ for monthly data [50]. We tested our calibrated model with 24-h periods taken from different times of the year to cover hot and cold days. The results are shown in Fig. 12. In particular, Fig. 12(a) shows the temperature response of the model on a typical winter day when the HVAC system is turned off. It can be seen that CVRMSE is still below the 2.0% threshold and the model is considered calibrated. Fig. 12(b) shows the model response on a hot summer day. In this case, CVRMSE has become above the threshold and the algorithm will need to re-calibrate the model to once again make CVRMSE below the threshold. These figures also show the detected occupancy to demonstrate how the model response varies with changing occupancy patterns. We note that in our testbed building there are no days without occupancy because there are always multiple prayers and multiple worshippers every day.

We observe that the calibrated values of the uncertain HVAC parameters are very close to the manufacturer's specifications. In other words, even if those specifications were not available, our approach can still be applied without any noticeable change in performance. This makes our approach more practical in situations where such information is not available. We experimented with parameters other than those outlined in Table 6, and found that their calibration did not yield any noticeable improvements. An example is illustrated in Fig. 13. In particular, Fig. 13(a) depicts the impact of calibrating two parameters from Table 6, whereas Fig. 13(b) depicts the impact of calibrating two parameters that are not outlined in Table 6, which are: (i) the roofs thickness and conductivity, and (ii) the windows thickness and conductivity. Here, the x -axis and y -axis represent the relative change (compared to the default value) in the first parameter and in the second parameter, respectively, while the z -axis represents the corresponding CVRMSE value. As can be seen, the performance improvements from the calibration process vary significantly depending on the parameters being calibrated.

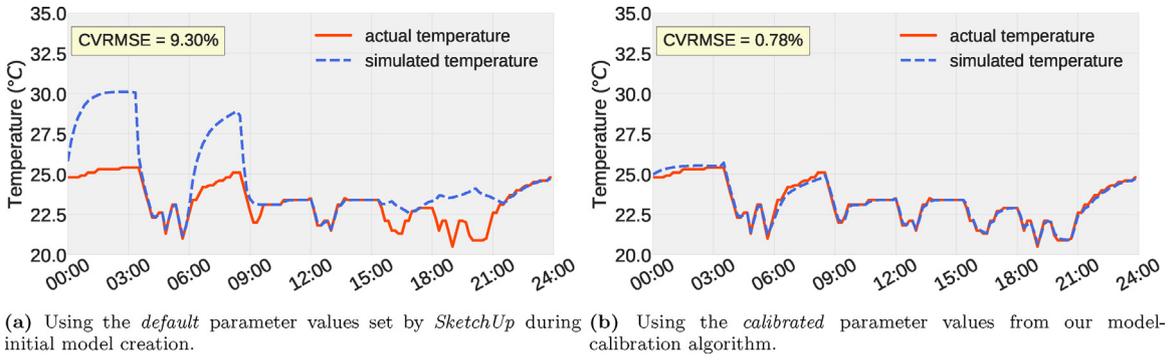


Fig. 11. Results of model calibration.

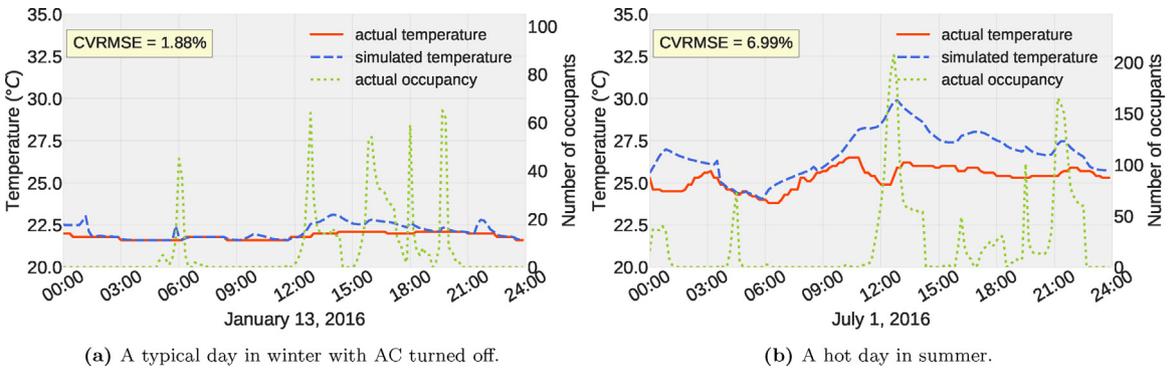


Fig. 12. Results of testing calibrated model with validation sets.

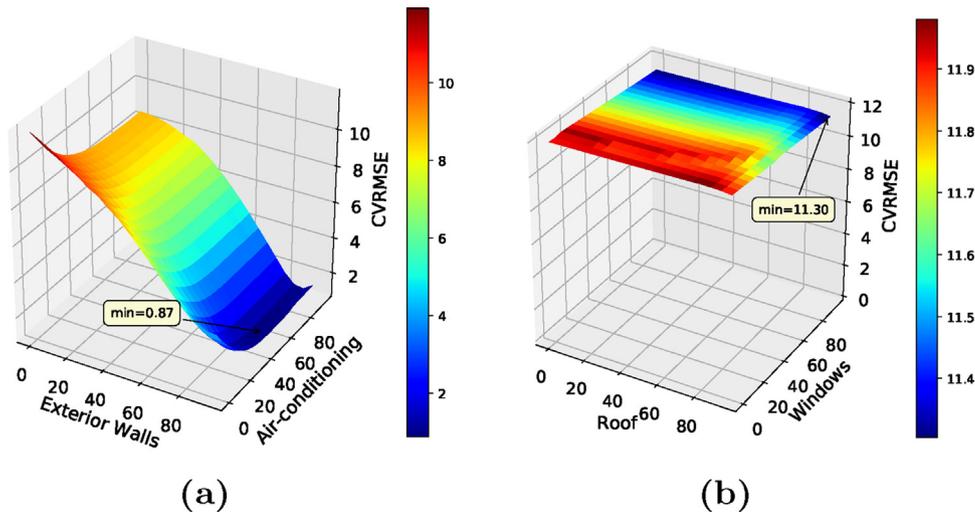


Fig. 13. Examples of how the performance improvements from the calibration process can vary significantly depending on the parameters being calibrated.

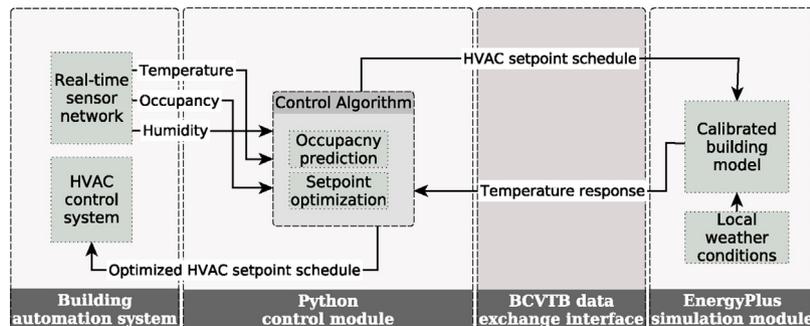


Fig. 14. Framework of model predictive control algorithm.

Table 6
Parameters used for model calibration.

Parameter	Field	Initial value	Calibrated value
Exterior	Outer layer	Thickness	2.53 cm
	(Stucco)	Conductivity	0.69 W/(m K)
	Layer 2	Thickness	20.32 cm
	(Concrete)	Conductivity	1.31 W/(m K)
	Layer 3	Thickness	1.27 cm
HVAC	(Gypsum)	Conductivity	0.16 W/(m K)
	Cooling capacity		87,920 W
	Air flow rate		3.78 m ³ /s
			164,850 W
			7.08 m ³ /s

Table 7
Setpoints used by HVAC-MPC. Here, setpoint_{oc} is used when the building is occupied; setpoint_{uo} is used when unoccupied.

Occupied (setpoint _{oc})	Unoccupied (setpoint _{uo})
24 °C	28 °C

Finally, commenting on the runtime of the calibration algorithm, our implementation on Raspberry Pi 3 took an average of 43 s per co-simulation, which allows for about 2000 co-simulations per day (recall that our satisfactory results from Fig. 11 required only 100 co-simulations). These results demonstrate that our calibration methodology can be used in practice, to automatically and continuously correct the building model for effective real-time HVAC control using low-cost embedded computers such as Raspberry Pi.

6. Simulation-guided model predictive control

Building upon our occupancy-prediction from Section 4 and our temperature-response simulation from Section 5, we now propose a model predictive control algorithm in Section 6.1, and evaluate it in a real-world HVAC system in Section 6.2.

6.1. HVAC control framework

The framework of our model predictive control consists of several parts that interact with one another as illustrated in Fig. 14. Here, a control algorithm provides an HVAC setpoint schedule to the EnergyPlus simulator through the BCVTB interface. Based on this schedule and the local weather information, the EnergyPlus simulator performs simulations and provides temperature response to the control algorithm. Note that the local weather information is provided by the *EnergyPlus Weather File* [46] which includes temperature, humidity, pressure, wind speed, solar radiation, luminance, precipitation, etc. Now, based on the temperature response as well as the occupancy prediction, the control algorithm checks whether thermal-comfort criteria will be satisfied; if so, then there is no need for any modifications to the HVAC setpoint schedule; if not, then the HVAC setpoint schedule is modified subject to the thermal-comfort threshold, before being sent again to the EnergyPlus simulator, and so on. Once the HVAC setpoint schedule is finalized, the HVAC system starts following this schedule, and the control algorithm starts monitoring the real-time temperature, humidity, and occupancy data, to determine the actual thermal comfort that resulted from its control decisions, and adjust its thermal-comfort threshold accordingly.

Algorithm 1. HVAC-MPC

```

Input:  $t_{\text{now}}$ , Initialization:  $\tau \leftarrow 0$ , timer  $\leftarrow 0$ 
1 if occupnow = 0 and timer = 0 then
2    $T_{\text{target}} \leftarrow \text{setpoint}_{\text{uo}}$ 
3   Set HVAC setpoint to  $T_{\text{target}}$   $\triangleright$  increase setpoint
4    $\tau \leftarrow \text{Pre-cooling}(t_{\text{now}}, t_{\text{pd.oc}})$   $\triangleright$  pre-cooling time
5   timer  $\leftarrow t_{\text{now}} + \tau$ 
6 end
7 if  $t_{\text{now}} = \text{timer}$  then
8    $T_{\text{target}} \leftarrow \text{setpoint}_{\text{oc}}$ 
9   Set HVAC setpoint to  $T_{\text{target}}$   $\triangleright$  start pre-cooling
10  timer  $\leftarrow 0$ 
11 else
12  Wait for timer expiry  $\triangleright$  no pre-cooling yet
13 end

```

Having provided an overview of our framework for model predictive control, we will now explain the control algorithm therein. In particular, we call this algorithm HVAC-MPC, the pseudocode of which is outlined in Algorithm 1. First, in line 1 the algorithm checks whether the building is unoccupied at the current time. If so, then in lines 2 and 3 it instructs the HVAC system to increase the temperature setpoint (see Table 7 for the exact setpoints used by the algorithm).

After that, in lines 4 and 5, the occupancy prediction is used to determine when the building is expected to become occupied in the future, and determine the optimal pre-cooling time accordingly (here the algorithm uses a procedure called *Pre-cooling*, the workings of which will be explained later on in this section). Then, in line 7, the algorithm checks whether the time has come for the pre-cooling to start. If so, then it instructs the HVAC system to decrease the temperature setpoint, before resetting the corresponding timer (see lines 8–10). On the other hand, if the time has not yet come to start pre-cooling, the algorithm simply waits until it does (see lines 11 and 12).

Having explained the pseudocode of HVAC-MPC, we now explain the workings of the *Pre-cooling* procedure used therein. Fig. 15 illustrates the flowchart of this procedure. Basically, the goal here is to determine the time t^* at which the pre-cooling should start; this time falls within a certain range of feasible times, denoted by $[t_{\text{min}}, t_{\text{max}}]$. Initially, the algorithm sets t_{min} to be equal to the current time (i.e., t_{now}), sets t_{max} to be equal to the predicted time of occupancy (i.e., $t_{\text{pd.oc}}$), and sets t^* to be right in the middle between the two (i.e., $t^* = t_{\text{min}} + (t_{\text{max}} - t_{\text{min}})/2$). After that, t^* is adjusted iteratively as follows. In each iteration, the algorithm runs an EnergyPlus simulation in which the HVAC is switched off from t_{min} to $t^* - 1$, and then the pre-cooling starts at t^* , leaving the HVAC on from t^* to t_{max} . Thus, we obtain the simulated temperature, $T_{\text{sim}}[t]$, for every $t \in [t_{\text{min}}, t_{\text{max}}]$. These simulated temperatures should ideally satisfy the following conditions:

$$T_{\text{sim}}[t_{\text{pd.oc}}] = \text{setpoint}_{\text{oc}} \quad (7)$$

$$T_{\text{sim}}[t] > \text{setpoint}_{\text{oc}}, \quad \forall t : t^* \leq t < t_{\text{pd.oc}} \quad (8)$$

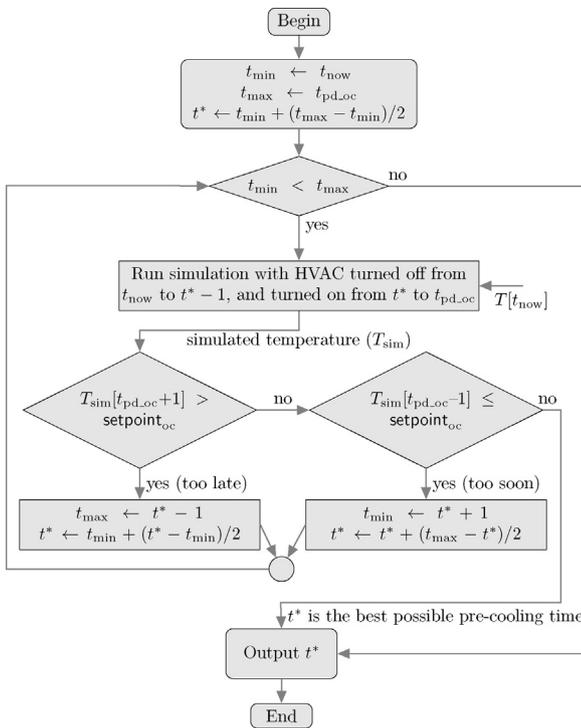


Fig. 15. Flowchart of our *Pre-cooling* procedure, which determines the pre-cooling time based on the predicted occupancy.

which basically means that the temperature becomes satisfactory exactly when needed, and not before. The algorithm checks whether the above two conditions hold. Now:

- if Condition (7) does not hold, it implies that the pre-cooling process started too late, and so t^* should be set to an earlier time. To this end, the algorithm sets t_{\max} to be equal to $t^* - 1$, and then updates t^* as follows:

$$t^* \leftarrow t_{\min} + (t^* - t_{\min})/2 \quad (9)$$

- if Condition (8) does not hold, it implies that the pre-cooling process started too soon, and so t^* should be delayed. To this end, the algorithm sets t_{\min} to be equal to $t^* + 1$, and then updates t^* as follows:

$$t^* \leftarrow t^* + (t_{\max} - t^*)/2 \quad (10)$$

After that, the algorithm proceeds to the next iteration to check whether the new t^* needs any further adjustments. This process is repeated until either $t_{\min} \geq t_{\max}$, or conditions (7) and (8) are both met. Either way, the best possible t^* is found.

6.2. Evaluation results

In this subsection, we evaluate the performance of our HVAC-MPC algorithm. To this end, we consider two standard measures of thermal comfort, namely *Predicted Mean Vote* (PMV) and *Predicted Percentage Dissatisfied* (PPD) [51]. Specifically, PMV quantifies the human perception of thermal sensation on a scale that runs from -3 to $+3$, where -3 is very cold, 0 is neutral, and $+3$ is very hot. On the other hand, PPD is built upon PMV to quantify the percentage of occupants that are dissatisfied given the current thermal conditions. The recommended PMV range for thermal comfort is between -0.5 and $+0.5$ for indoor spaces, while the acceptable PPD range is between 5% and 10% [52].

Table 8
Summary of energy savings attained by HVAC-MPC.

Metric	Value
Average daily savings over experiment period	456 kWh
Standard deviation of daily savings	118 kWh
Total savings over the seven-day experiment	3197 kWh

We compared our HVAC-MPC algorithm against a baseline alternative, where the HVAC is turned on throughout the day, regardless of the varying occupancy. The evaluation results are shown in Fig. 16. Specifically:

- Fig. 16(a) depicts the observed occupancy, the indoor and outdoor temperatures, and the HVAC status given the baseline HVAC control;
- Fig. 16(b) depicts the thermal comfort according to PMV and PPD given the baseline HVAC control;
- Fig. 16(c) depicts the predicted occupancy, the indoor temperature, and the HVAC status given HVAC-MPC algorithm;
- Fig. 16(d) depicts the thermal comfort according to PMV and PPD given HVAC-MPC algorithm.

Let us first comment on the results of the baseline HVAC control. As can be seen in Fig. 16(a), the HVAC is turned on throughout the day, despite the significant changes in both the occupancy and the external temperature. In terms of thermal comfort, the baseline HVAC control performs rather poorly, as shown in Fig. 16(b). Here, the recommended range for PMV is highlighted by the dotted region. As can be seen, the PMV index is outside the recommended range throughout the day. Furthermore, according to the PDD index, a considerable percentage of occupants are dissatisfied most of the day.

Moving on to the results of our HVAC-MPC algorithm, Fig. 16(c) shows how the HVAC is switched off for a considerable number of hours. This is because, whenever the building becomes unoccupied, HVAC-MPC increases the temperature setpoint and predicts the future temperature and occupancy to find the best possible time for pre-cooling. Fig. 16(d), on the other hand, depicts the PMV and PPD indices throughout the day. As can be seen, when the building is occupied, the PMV index is almost always within the recommended range. Likewise, when the building is occupied, the PPD is close to 5% (which is the best possible PPD score that can be achieved).

Finally, we present the energy savings that are attained by our HVAC-MPC algorithm. The algorithm was activated in the testbed building for a duration of one week in total and the results are provided in Fig. 17. In particular, Fig. 17(a) shows the results for three different days, where each day is in a different month. For each day shown, the algorithm performance is compared with the immediately preceding day where HVAC-MPC is not activated. Fig. 17(b) shows the results of another experiment where we activated HVAC-MPC for four consecutive days in the same week. For each day, the results are compared with the corresponding day of the previous week in which HVAC-MPC was not activated. The purpose of this experiment was to find out how much energy is saved by the MPC algorithm relative to the same day of the previous week. The energy savings attained by HVAC-MPC range from 23% to 39% . Table 8 lists the average daily savings, the standard deviation, and the total energy savings over the experiment period.

7. Testbed design and implementation

Our automatic HVAC control system with real-time video-based occupancy recognition and EnergyPlus-guided model predictive control is implemented on the Raspberry Pi 3 platform. Our choice is motivated by the fact that Raspberry Pi 3 is a low-cost embedded system platform that costs as little as $\$35$ per unit, as of 2016.

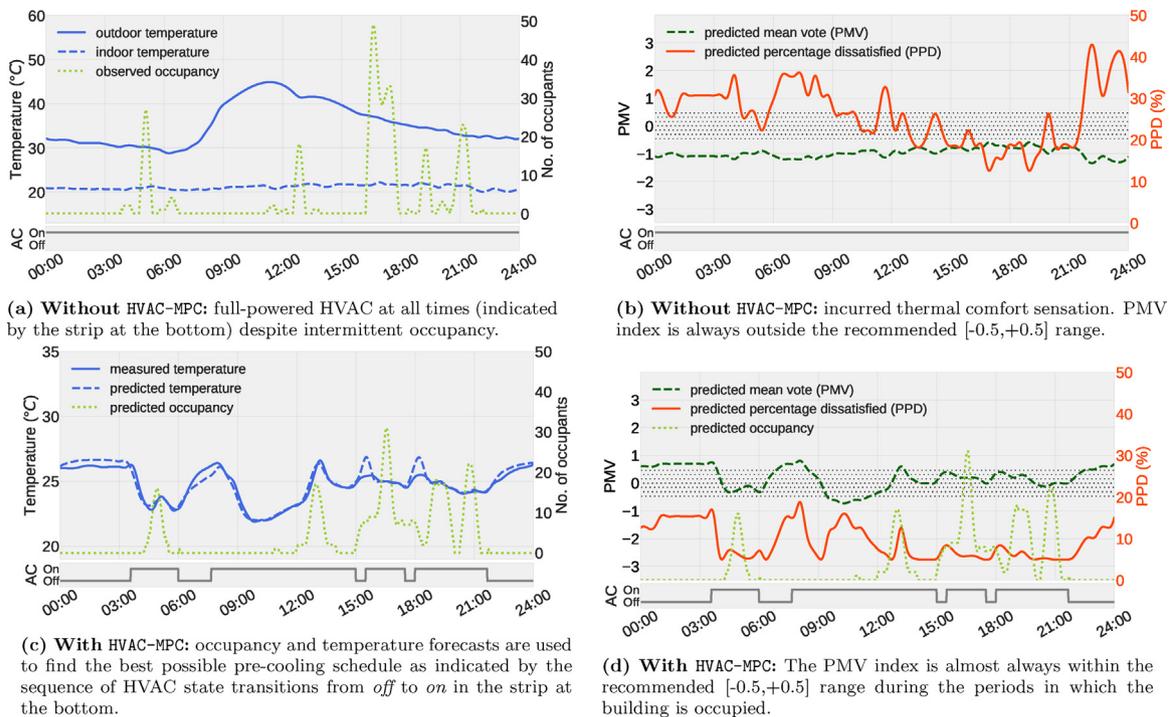


Fig. 16. Results of HVAC model predictive control (HVAC-MPC) in our testbed.

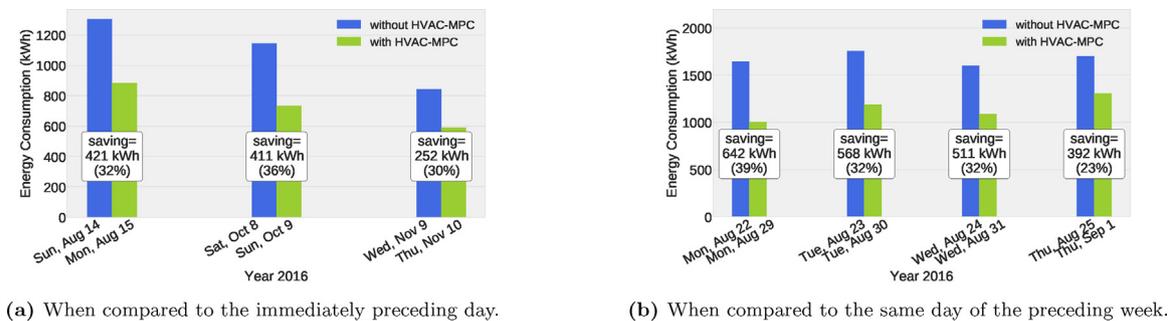


Fig. 17. Reduction in HVAC energy consumption achieved by HVAC-MPC (Algorithm 1).

Although there are other embedded systems (e.g., Intel Galileo), to the best of our knowledge, Raspberry Pi is the cheapest in the market for our purpose of HVAC control. A comparison of different embedded systems in terms of cost and performance can be found in [53]. Another advantage of Raspberry Pi 3 is its rich set of specifications, which include: 1.2 GHz 64-bit quad-core CPU, 1 GB SDRAM memory, SD card based expandable external data storage, WiFi & Bluetooth connectivity, and GPU. The remainder of this section provides more details about our testbed design and implementation.

7.1. System design

Fig. 18 illustrates the overall design of our system. This system is comprised of two subsystems: (1) the building automation system, and (2) the cloud-based management system. Next, we explain each of these systems.

Starting with the building automation system, it consists of the hardware components installed in the building, including temperature and humidity sensors, occupancy monitors, energy meters, and customized wireless HVAC controllers. The temperature and humidity sensors are based on *EnOcean* technology, and have the advantage of being self-powered and maintenance-free, which

makes them highly flexible for deployment anywhere in the target environment. The occupancy monitor consists of a dedicated Raspberry Pi and a fish-eye camera for real-time tracking of occupancy. An Arduino-based energy meter is used to measure the energy consumption of the HVAC units. A data daemon running on the base station Raspberry Pi stores the data received from these components. The data daemon also provides the received data to a control daemon (also running on the base station Raspberry Pi), which uses it to make HVAC control decisions. To this end, the control daemon interfaces with the HVAC system through wireless controllers implemented in a Raspberry Pi. The base station module is shown in Fig. 19(d).

Having described the building automation system, we now move on to the cloud-based management system. In particular, this system is comprised of the remote server and the web graphical user interface (GUI). The remote server receives the collected data and stores it in a server database for later analysis and permanent storage. The web GUI retrieves data from the server database for visualization and analysis. The web GUI also provides remote access to the building control system. Secure Shell (SSH) protocol is used as the underlying protocol to securely provide remote access.

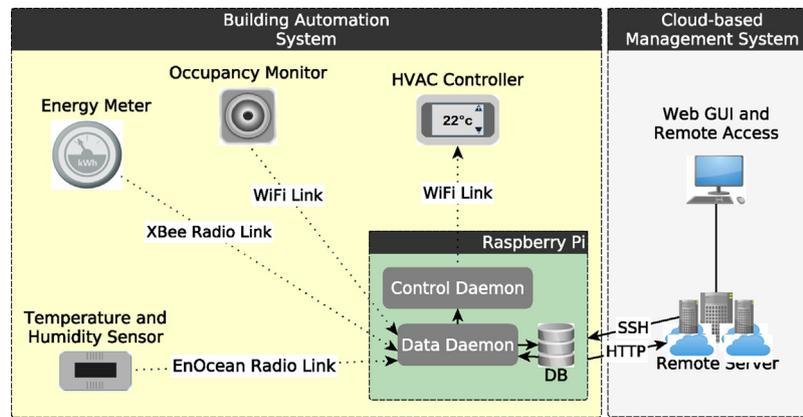


Fig. 18. Outline of our system, which is comprised of two subsystems: (1) the building automation system, and (2) the cloud-based management system.

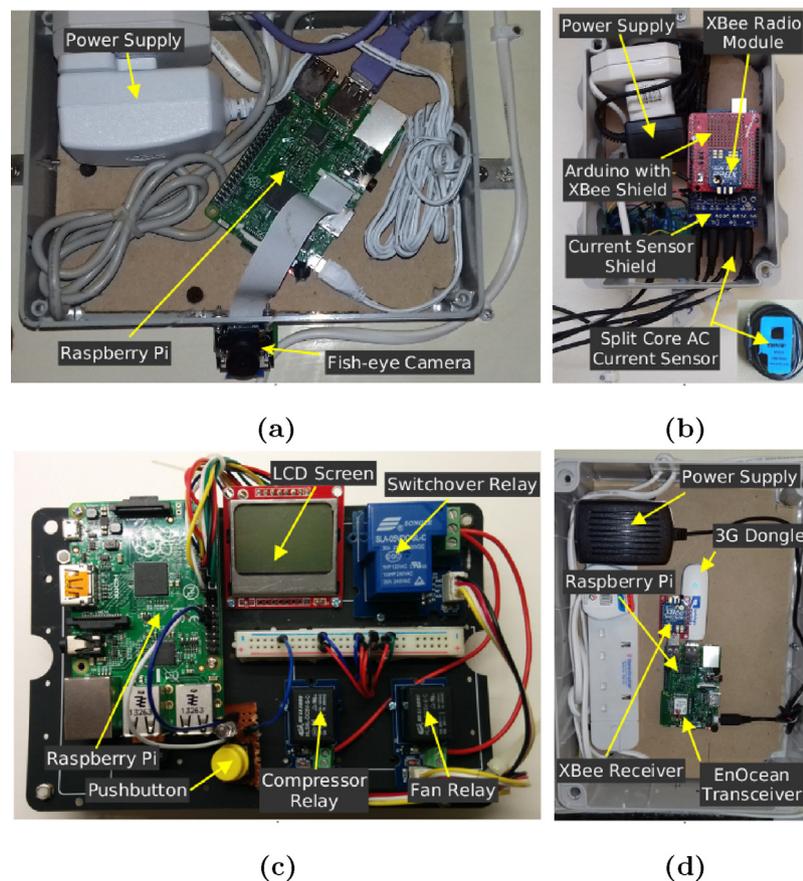


Fig. 19. Testbed system components: (a) occupancy detection module, (b) HVAC energy measurement module, (c) HVAC controller, and (d) base station module.

7.2. Occupancy detection module

Our system monitors the building occupancy using a Raspberry Pi with a fish-eye camera. In order to keep track of the number of occupants in the building, a video processing algorithm runs on the Raspberry Pi to analyze the real-time video stream captured by the fish-eye camera. Fig. 19(a) shows the occupancy module developed for this research. The module should be installed above the building entrance door so that it can monitor the occupancy inside the building. Multiple modules may be needed if the building has multiple entry/exit points, such that each module covers a single point.

7.3. HVAC energy measurement module

To collect the HVAC energy consumption data, we developed an Arduino-based module shown in Fig. 19(b). The module design is adapted from the *OpenEnergyMonitor* framework, which is an open source project for developing energy monitoring and analysis tools [54]. The module uses non-invasive alternating current (AC) sensors to measure the HVAC energy consumption. These sensors can be clipped onto either the live or ground wire coming into the HVAC unit without needing to strip the wire, thus avoiding any high-voltage work. We used a dedicated sensor for each HVAC unit installed in the building. The collected data is sent to the base sta-

tion through an XBee radio link. It was decided to use XBee radio technology because the HVAC unit power supply can be in a separate room or even on the roof. In such cases, the XBee modules with their extended transmission range can penetrate concrete walls and roofs and are able to transmit the sensor readings to the base station.

7.4. HVAC controller

Fig. 19(c) shows the wireless HVAC controller implemented in Raspberry Pi. The module operates the HVAC unit according to the setpoint schedule, which is determined by the MPC algorithm. The module has three relays, two of which are used to open and close the compressor and fan circuits of the HVAC unit to keep the indoor temperature within fixed bound of the MPC setpoint. The third relay (labeled as *Switchover Relay* in Fig. 19(c)) is used to switch control between our wireless controller and the default HVAC thermostat controller, allowing the building occupants to disable the automatic HVAC control if they are dissatisfied with the current thermal comfort level or if the system is malfunctioning. The module is also equipped with an LCD display to provide status information. By adjusting the setpoints, we actually control the HVAC system directly. This generic approach can be applied to different HVAC systems, with varying specifications.

8. Conclusion

This paper presents an automatic HVAC control system, featuring real-time occupancy recognition, dynamic occupancy prediction, and simulation-guided model predictive control, implemented in a low-cost embedded system (Raspberry Pi). We deployed and evaluated our system for providing automatic HVAC control in the large public indoor space of a mosque. Our experiments showed that our real-time occupancy recognition system can reach 90% accuracy, whereas our occupancy prediction system can reach 85% accuracy. We employ real-time HVAC control guided by an on-board EnergyPlus simulator, which is able to achieve more than 30% energy saving while maintaining the comfort level within acceptable range. Importantly, our system is sufficiently general and can be deployed in other types of buildings with large public indoor spaces.

Notably, we ported the EnergyPlus simulator to the Raspberry Pi embedded system platform. We release our Raspberry Pi version of EnergyPlus publicly [15] to enable other researchers to take advantage of our work for future building automation projects.

Our testbed has been evaluated in buildings with somewhat predictable occupancy patterns. It is challenging to apply our system to settings with irregular occupancy patterns. In future work, we will explore robust online HVAC control using minimal occupancy prediction. Online algorithms have been applied to wireless sensor based building control [55,56]. Robust online control can ensure good performance in the presence of dynamic irregular environmental factors.

Acknowledgments

This research was conducted with permission of Abu Dhabi's General Authorities of Islamic Affairs and Endowments for the mosque testbed. We would like to thank Afshin Afshari and Prashant Shenoy for helpful discussions.

References

- [1] L. Pérez-Lombard, J. Ortiz, C. Pout, A review on buildings energy consumption information, *Energy Build.* 40 (3) (2008) 394–398.
- [2] F. Oldewurtel, A. Parisio, C.N. Jones, D. Gyalistras, M. Gwerder, V. Stauch, B. Lehmann, M. Morari, Use of model predictive control and weather forecasts for energy efficient building climate control, *Energy Build.* 45 (2012) 15–27.
- [3] J. Álvarez, J. Redondo, E. Camponogara, J. Normey-Rico, M. Berenguel, P. Ortigosa, Optimizing building comfort temperature regulation via model predictive control, *Energy Build.* 57 (2013) 361–372.
- [4] S. Salakij, N. Yu, S. Paolucci, P. Antsaklis, Model-based predictive control for building energy management. I: energy modeling and optimal control, *Energy Build.* 133 (2016) 345–358.
- [5] Y. Kwak, J.-H. Huh, C. Jang, Development of a model predictive control framework through real-time building energy management system data, *Appl. Energy* 155 (2015) 1–13.
- [6] P.-D. Moroşan, R. Bourdais, D. Dumur, J. Buisson, Building temperature regulation using a distributed model predictive control, *Energy Build.* 42 (9) (2010) 1445–1452.
- [7] F. Ascione, N. Bianco, C.D. Stasio, G.M. Mauro, G.P. Vanoli, Simulation-based model predictive control by the multi-objective optimization of building energy performance and thermal comfort, *Energy Build.* 111 (2016) 131–144.
- [8] H. Park, M. Ruellan, A. Bouvet, E. Monmasson, R. Bennacer, Thermal parameter identification of simplified building model with electric appliance, in: 11th International Conference on Electrical Power Quality and Utilisation, 2011, pp. 1–6.
- [9] R. Bălan, J. Cooper, K.-M. Chao, S. Stan, R. Donca, Parameter identification and model based predictive control of temperature inside a house, *Energy Build.* 43 (2) (2011) 748–758.
- [10] J. Hu, P. Karava, A state-space modeling approach and multi-level optimization algorithm for predictive control of multi-zone buildings with mixed-mode cooling, *Build. Environ.* 80 (2014) 259–273.
- [11] C.-H. Jan, 10 years of transistor innovations in system-on-chip (soc) era, in: 12th IEEE International Conference on Solid-State and Integrated Circuit Technology (ICSICT), IEEE, 2014, pp. 1–4.
- [12] Raspberry pi foundation. Available from: <https://www.raspberrypi.org>, retrieved on 31.01.17.
- [13] B. Pavlin, G. Pernigotto, F. Cappelletti, P. Bison, R. Vidoni, A. Gasparella, Real-time monitoring of occupants' thermal comfort through infrared imaging: a preliminary study, *Buildings* 7 (1) (2017) 10.
- [14] D.B. Crawley, L.K. Lawrie, F.C. Winkelmann, W. Buhl, Y. Huang, C.O. Pedersen, R.K. Strand, R.J. Liesen, D.E. Fisher, M.J. Witte, J. Glazer, Energyplus: creating a new-generation building energy simulation program, *Energy Build.* 33 (4) (2001) 319–331. Special issue: (BUILDING) SIMULATION'99.
- [15] Energyplus for raspberry pi. Available from: https://github.com/muhaftab/energyplus_rpi, retrieved on 05.05.17.
- [16] Number of mosques in U.A.E. Available from: <https://www.awqaf.gov.ae/Affair.aspx?SectionID=3&RefID=18>, retrieved on 20.01.17.
- [17] Number of mosques in k.s.a. Available from: <http://www.moia.gov.sa/menu/pages/statistics.aspx>, retrieved on 20.01.17.
- [18] D.T. Delaney, G.M. O'Hare, A.G. Ruzzelli, Evaluation of energy-efficiency in lighting systems using sensor networks, in: Proceedings of ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Buildings, ACM, 2009, pp. 61–66.
- [19] A. Barbato, L. Borsani, A. Capone, S. Melzi, Home energy saving through a user profiling system based on wireless sensors, in: Proceedings of ACM workshop on embedded sensing systems for energy-efficiency in buildings, ACM, 2009, pp. 49–54.
- [20] R.S. Hsiao, D.B. Lin, H.P. Lin, S.C. Cheng, C.H. Chung, A robust occupancy-based building lighting framework using wireless sensor networks, in: Applied Mechanics and Materials, vol. 284, Trans Tech Publ., 2013, pp. 2015–2020.
- [21] J. Li, L. Huang, C. Liu, Robust people counting in video surveillance: dataset and system, in: 8th IEEE International Conference on Advanced Video and Signal-Based Surveillance (AVSS), IEEE, 2011, pp. 54–59.
- [22] L. Chen, F. Chen, X. Guan, A video-based indoor occupant detection and localization algorithm for smart buildings, in: Emerging Intelligent Computing Technology and Applications, Springer, 2009, pp. 565–573.
- [23] M. Jia, R.S. Srinivasan, Occupant behavior modeling for smart buildings: a critical review of data acquisition technologies and modeling methodologies, in: Winter Simulation Conference (WSC), IEEE, 2015, pp. 3345–3355.
- [24] W. Kleiminger, F. Mattern, S. Santini, Predicting household occupancy for smart heating control: a comparative performance analysis of state-of-the-art approaches, *Energy Build.* 85 (2014) 493–505.
- [25] B. Dong, K.P. Lam, C. Neuman, Integrated building control based on occupant behavior pattern detection and local weather forecasting, in: Twelfth International IBPSA Conference, IBPSA Australia, Citeseer, Sydney, 2011, pp. 14–17.
- [26] B. Dong, K.P. 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, Springer, 2014, pp. 89–106.
- [27] P. Ferreira, A. Ruano, S. Silva, E. Conceição, Neural networks based predictive control for thermal comfort and energy savings in public buildings, *Energy Build.* 55 (2012) 238–251, cool Roofs, Cool Pavements, Cool Cities, and Cool World.
- [28] Identification of nonlinear systems using narmx model, *Nonlinear Anal. Theory Methods Appl.* 71 (12) (2009) e1198–e1202.
- [29] D.B. Crawley, J.W. Hand, M. Kummert, B.T. Griffith, Contrasting the capabilities of building energy performance simulation programs, *Build. Environ.* 43 (4) (2008) 661–673 (part Special: Building Performance Simulation).

- [30] D. Coakley, P. Raftery, M. Keane, A review of methods to match building energy simulation models to measured data, *Renew. Sustain. Energy Rev.* 37 (2014) 123–141.
- [31] J. Cigler, D. Gyalistras, J. Široky, V. Tiet, L. Ferkl, Beyond theory: the challenge of implementing model predictive control in buildings, in: *Proceedings of 11th Rehva World Congress, Clima*, vol. 250, 2013.
- [32] J. Zhao, K.P. Lam, B.E. Ydstie, V. Loftness, Occupant-oriented mixed-mode energyplus predictive control simulation, *Energy Build.* 117 (2016) 362–371.
- [33] D. Sturzenegger, D. Gyalistras, M. Gwerder, C. Sagerschnig, M. Morari, R.S. Smith, Model predictive control of a Swiss office building, in: *Clima-Rheva World Congress*, 2013, pp. 3227–3236.
- [34] B. Dong, B. Andrews, Sensor-based occupancy behavioral pattern recognition for energy and comfort management in intelligent buildings, in: *Proceedings of Building Simulation*, 2009, pp. 1444–1451.
- [35] B. Dong, K.P. Lam, Building energy and comfort management through occupant behaviour pattern detection based on a large-scale environmental sensor network, *J. Build. Perform. Simul.* 4 (4) (2011) 359–369.
- [36] S. Goyal, H.A. Ingle, P. Barooah, Occupancy-based zone-climate control for energy-efficient buildings: complexity vs. performance, *Appl. Energy* 106 (2013) 209–221.
- [37] G. Bradski, *Dr. Dobb's Journal of Software Tools*.
- [38] openFrameworks Community, *openframeworks*.
- [39] Z. Zivkovic, F. van der Heijden, Efficient adaptive density estimation per image pixel for the task of background subtraction, *Pattern Recogn. Lett.* 27 (7) (2006) 773–780.
- [40] Z. Zivkovic, Improved adaptive Gaussian mixture model for background subtraction, in: *Proceedings of IEEE International Conference on Pattern Recognition*, vol. 2, IEEE, 2004, pp. 28–31.
- [41] S. Suzuki, K. Abe, Topological structural analysis of digitized binary images by border following, *CVGIP* 30 (1) (1985) 32–46.
- [42] N. Halko, P.-G. Martinsson, J.A. Tropp, Finding structure with randomness: probabilistic algorithms for constructing approximate matrix decompositions, *SIAM Rev.* 53 (2) (2011) 217–288.
- [43] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, E. Duchesnay, Scikit-learn: machine learning in Python, *J. Mach. Learn. Res.* 12 (2011) 2825–2830.
- [44] Y. Yang, X. Liu, A re-examination of text categorization methods, in: *Proceedings of the 22nd Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, ACM, 1999, pp. 42–49.
- [45] W. Cohen, P. Ravikumar, S. Fienberg, A comparison of string metrics for matching names and records, in: *Kdd Workshop on Data Cleaning and Object Consolidation*, vol. 3, 2003, pp. 73–78.
- [46] Energyplus standard weather data. Available from: <https://energyplus.net/weather>, retrieved on 08.02.17.
- [47] M. Wetter, Co-simulation of building energy and control systems with the building controls virtual test bed, *J. Build. Perform. Simul.* 4 (3) (2011) 185–203.
- [48] M. Royapoor, T. Roskilly, Building model calibration using energy and environmental data, *Energy Build.* 94 (2015) 109–120.
- [49] Energyplus input data dictionary (idd). Available from: <https://energyplus.net/sites/default/files/pdfs/pdfs.v8.3.0/InputOutputReference.pdf>, retrieved on 05.05.17.
- [50] A. Guideline, *Ansi/ashrae standard 14-2014, Measurement of energy, demand, and water savings*.
- [51] P.O. Fanger, et al., *Thermal Comfort. Analysis and Applications in Environmental Engineering*, 1970.
- [52] A. Standard, *Ansi/ashrae standard 55-2004, Thermal Environmental Conditions for Human Occupancy*.
- [53] Embedded Linux Board Comparison. Available from: <https://learn.adafruit.com/embedded-linux-board-comparison?view=all>, retrieved on 05.05.17.
- [54] Open energy monitor – a project to develop and build open source energy monitoring and analysis tools. Available from: <https://openenergymonitor.org/>, retrieved on 07.02.17.
- [55] C.-K. Chau, M. Khonji, M. Aftab, Online algorithms for information aggregation from distributed and correlated sources, *IEEE/ACM Trans. Netw.* 24 (6) (2016) 3714–3725.
- [56] M. Aftab, C.-K. Chau, P. Armstrong, Smart air-conditioning control by wireless sensors: an online optimization approach, in: *Proceedings of ACM International Conference on Future Energy Systems (e-Energy)*, 2013, pp. 225–236.