Building and Environment 122 (2017) 145–160

Contents lists available at ScienceDirect

Building and Environment

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

Occupancy learning-based demand-driven cooling control for office spaces

Yuzhen Peng [a, *], Adam Rysanek [a], Zoltán Nagy [b], Arno Schlüter [a]

[a] Architecture and Building Systems, Institute of Technology in Architecture, Department of Architecture, ETH Zürich, Switzerland
[b] Intelligent Environments Laboratory, Department of Civil, Architectural, and Environmental Engineering, The University of Texas at Austin, USA

ARTICLE INFO

Article history:
Received 6 March 2017
Received in revised form
11 May 2017
Accepted 5 June 2017
Available online 7 June 2017

Keywords:
Occupancy learning
Occupancy prediction
Demand-driven control
HVAC
Energy efficiency
Intelligent systems

ABSTRACT

Occupancy in buildings is one of the key factors influencing air-conditioning energy use. Occupant presence and absence are stochastic. However, static operation schedules are widely used by facility departments for air-conditioning systems in commercial buildings. As a result, such systems cannot adapt to actual energy demand for offices that are not fully occupied during their operating time. This study analyzes a seven-month period of occupancy data based on motion signals collected from six offices with ten occupants in a commercial building, covering both private and multi-person offices. Based on an occupancy analysis, a learning-based demand-driven control strategy is proposed for sensible cooling. It predicts occupants' next presence and the presence duration of the remainder of a day by learning their behavior in the past and current days, and then the predicted occupancy information is employed indirectly to infer setback temperature setpoints according to rules we specified in this study. The strategy is applied for the controls of a cooling system using passive chilled beams for sensible cooling of office spaces. Over the period of two months both a baseline control and the proposed demand-driven control were operated on forty-two weekdays of real-world occupancy. Using the demand-driven control, an energy saving of 20.3% was achieved as compared to the benchmark. We found that energy savings potential in an individual office was inversely correlated to its occupancy rate.

1. Introduction

The latest 5th Assessment Report of the Intergovernmental Panel on Climate Change has indicated that anthropogenic greenhouse gas (GHG) emissions will continue to cause further warming of the Earth's surface and cause long-lasting changes to the world's climate system. The contribution of buildings to global energy use and energy-related GHG emissions are, in fact, significant. Globally, buildings in the residential, commercial, public and service sectors accounted for about 35% of total final energy use and were associated with 18.4% direct GHG emissions and indirect carbon dioxide (CO2) emissions (e.g. electricity) in 2010. Moreover, building-related energy demand is projected to increase by about 50% between 2010 and 2050 [1–3].

The main services consuming energy in buildings are space heating, ventilation, and air-conditioning (HVAC), domestic hot water, lighting, and electrical appliances. HVAC alone accounts for the largest share. Worldwide, HVAC services account for approximately 40% of total energy consumption in buildings [4]. In particular harsh climate, such as the tropical context of Singapore, HVAC accounts for over 50% of the building stock's electricity consumption [5].

Improving the energy efficiency and utility of existing and future HVAC systems will, therefore, be an important objective for developing future low-carbon economies. Developing a better understanding of occupants' behavior in buildings will also be an increasingly important concern in this process. The presence and absence of building occupants indicate whether indoor spaces are required to be air-conditioned or not. Building HVAC systems need to provide comfortable indoor conditions when the building spaces they serve are occupied. On the other, they do not need to ensure indoor conditions are comfortable with spaces unoccupied [6]. Whilst this may be intuitive, the poor anticipation of occupant behavior has been found to increase building energy consumption by a third [7]. Furthermore, not all occupants in buildings are sufficiently aware of this or other energy saving initiatives, especially in commercial buildings, as energy costs are not directly paid by

* Corresponding author.
E-mail address: yuzhen.peng@arch.ethz.ch (Y. Peng).

http://dx.doi.org/10.1016/j.buildenv.2017.06.010
0360-1323/

Nomenclature

Abbreviations

GHG	Greenhouse gas
CO2	Carbon dioxide
HVAC	Heating, ventilation, and air-conditioning
RFID	Radio frequency identification device
KNN	K-nearest neighbor
HMM	Hidden Markov model
M	Multi-person office
P	Private office
HMI	Human machine interface
DOAS	Dedicated outdoor air system
FCU	Fan coil unit
PCB	passive chilled beam
AHU	Air handling unit
PID	Proportional-integral-derivative
WSI	Web service interface
REST API	Representational state transfer, application programming interface
M-Bus	MeterBus
TD	Time delay
RBC	Rule-based control
BMS	Building management station
DCC	Demand-driven cooling control
MID	The measuring instruments directive
COV	Change of value
IMBPC	Inteligent Model Based Predictive Control
PIBCV	Pressure-independent balancing and control valve
CDD	Cooling degree-days
S1	Six offices that are used to evaluate the sensible cooling energy gap
S2	Six offices for the DCC study: P1, P2, P3, P4, M1, M2
CPU	Central processing unit
RAM	Random-access memory
RC	Resistance-capacitance
LCD	Liquid crystal display
VAV	Variable-air-volume

Symbols(unit)

T_{sp}	Temperature setpoint (°C)
T_{air}	Air temperature (°C)
N_x	The number of vacancy days in past x days
S_{td}	The size of the training dataset
K_{value}	The value of K
$P_{thrshld}$	The threshold of the occupancy possibility
t_{np}	Time of next presence (minute)
t_{dcc}	The time at which starting the demand-driven cooling control (minute)
t_{sd}	The time at which the facility department shuts down the air-conditioning system in the case study space (minute)
t_{arr_lmt}	The time at which the cumulative probability of the first arrivals is equal to a specified value (minute)
t_{dprtr_lmt}	The time at which the cumulative probability of the last departures is equal to a specified value (minute)
t_{drtn}	Presence duration of the remaining day (minute)
t_{drtn_lmt1}	The first threshold of presence duration (minute)
t_{drtn_lmt2}	The second threshold of presence duration (minute)
E_{nbl}	Normalized daily average cooling energy use of a room (kWh)
E_{bl}	Measured daily average cooling energy use of a room (kWh)
S_r	The area of a room (m^2)

them [8].

There are two features of conventional HVAC systems that have historically made it difficult for these systems to automatically respond to the stochastic nature of occupants' behavior in buildings [9,10]. The first regards to the behavior of physical controllers in existing HVAC systems, employing mostly two-position (i.e. on and off) control or proportional, integral and derivative (PID) control to keep indoor climates conditioned to temperature, humidity, and CO2 setpoints [11]. The second is the use of scheduled occupancy profiles to assign operating hours of HVAC control systems in commercial buildings.

Demand-driven control is an emerging HVAC control strategy that has shown promising results in coordinating real-time HVAC use to occupant presence and vacancy, reducing energy use and maintaining indoor thermal comfort in buildings [10,12–14]. Energy savings can be achieved by decreasing the temperature difference between the air-conditioned indoor climate and the outdoor weather or reducing the operating time of HVAC systems [15]. In the same manner, demand-driven HVAC control strategies decrease heating temperature setpoints or increase cooling temperature setpoints when spaces are unoccupied, and they keep the indoor spaces at comfortable levels when they are occupied. Furthermore, a demand-driven control system can automatically deactivate an HVAC system after the occupants have left a building instead of waiting for scheduled shutdown times.

Central to the effective implementation of a demand-driven HVAC control strategy is information on: 1) real-time occupancy and 2) upcoming room occupancy [10,14]. Networks of occupant-monitoring sensors are essential to measure occupants' behavior, while, at the same time, algorithms with learning capabilities are crucial for predicting future room occupancy. Prior research has shown that HVAC systems incorporating these features have yielded significant energy savings potential.

For instance, in a residential application, Scott et al. [12] developed a preheat heating system to anticipate to occupants' demand. They used radio frequency identification devices (RFID) and motion sensors to monitor real-time room occupancy status and utilized the K-nearest neighbor (KNN) algorithm to develop an occupancy forecast. Their control system then modified room temperature setpoints to preheat homes according to the expected occupancy periods. Test results showed that, on the implementation of this method, total gas consumption for heating decreased by 8%–18% over a 61-day period. Lu et al. [13] explored the energy-saving potential of a similar application in an EnergyPlus [16] simulation environment. They collected data from motion sensors and door sensors installed in each room of a house to generate room occupancy information, and they used a Hidden Markov Model (HMM) to forecast the probability of occupants' behavior (i.e. sleep, active, and not in the home) according to the generated occupancy datasets. Their simulated result produced an average energy reduction of 28% for cooling and heating over 14 days in summer and winter.

As more and more occupants in offices adopt flexible work hours [17], the total scheduled operating time of HVAC systems

may increase in duration in order to satisfy occupants of future commercial buildings [14]. This will come at the risk of consuming significant amounts of energy in space with low occupancy rates. Gunay et al. [14] employed demand-driven control to make heating and cooling systems adapt to flexible working hours in commercial buildings. They collected data from motion sensors in seven private offices and used a sequential learning process to predict occupants' arrival and departure times. Their control strategy was evaluated in the EnergyPlus simulation platform, and they predicted annual energy savings rates of 10%–15% would be possible for an academic building. Similarly, Erickson et al. [10] also evaluated their demand-driven HVAC control system in the EnergyPlus simulation environment. They gathered occupancy data from a camera sensor network in a real office setting and used a moving-window Markov Chain to infer probabilities of occupancy changes in relation to time. Their result indicated that 20% of annual energy consumption for HVAC could be saved on account of implementing their demand-driven control strategy. In another study, Ruano et al. [18] embedded an occupancy predictor in an intelligent Model Based Predictive Control (IMBPC) system. Their occupancy data was transformed from signals collected by motion sensors. They calculated an occupancy forecast from the mean of movement data ranging from the beginning of schedule to a sample point. Their approach was applied to 3 classrooms in the Gambelas Campus of the University of Algarve over 2 weeks, resulting in 56% energy savings.

Two of the above studies discussed the control accuracy achieved based on predicted occupancy information. For a household heating control system, simulation-based results in Ref. [13] illustrated that their control strategy based on the HMM approach for forecasting room occupancy status – sleep, active, and not in the home – was 88%. For the similar residential application, experimental results in Ref. [12] showed that the prediction accuracy of presence was around 80% for preheating the whole household with a 90-min look ahead time.

Whilst the above studies suggest that demand-driven HVAC control systems can yield significant energy savings, most of them, and related work in literature, undertook evaluations in the simulation environment only, with very little real-world implementation and verification. This is especially evident for the context of commercial buildings.

The contribution of this paper is, therefore, an empirical study of a demand-driven control strategy with learning capability, covering real-time occupancy detection, occupancy prediction that learns from historical datasets, and rule-based control. A holistic demand-driven HVAC control system is proposed, and results are presented and discussed upon the system's implementation and testing in a real commercial office space.

The remainder of this paper is structured as follows. Section 2 gives an introduction to the case study space and its HVAC system as well as the experimental setup. An analysis of the case study offices' occupancy profiles is presented in Section 3. In Section 4, a methodology for occupancy learning-based demand-driven HVAC control is proposed. Sections 5 illustrate test results upon implementing the proposed control system. Then, Section 6 discusses and concludes the study and indicates future work.

2. The experimental space and setup

2.1. Description of the office case study

The experiments undertaken in this paper were carried out within a 550 m^2 office space in Singapore designed and operated by researchers under the Chair of Architecture and Building Systems at ETH Zürich [19]. For matters of privacy, a floorplan of the study area cannot be provided in this paper in order to maintain the anonymity of occupancy data to be presented. In lieu of a floorplan, a thermal resistance-capacitance (RC) diagram is presented in Fig. 1, describing the thermal interactions between the studied interior rooms, and across the experimental space's boundary. The total area of six offices is 118 m^2.

Fig. 1 shows the six offices under study, depicted in gray boxes. Each office is either square or rectangular in shape, and is labeled as follows:

- M1: an office with two occupants, adjacent to the façade.
- M2: an office with four occupants.
- P1: a single-occupied office, adjacent to the façade.
- P2: a single-occupied office, adjacent to the façade.
- P3: a single-occupied office, adjacent to the façade.
- P4: a single-occupied office.

The air-conditioning system that serves the six offices is decoupled into two separate, though parallel operating systems. Decentralized outdoor air systems (DOASs) and fan coil units (FCUs) are deployed for ventilation and latent cooling, and a water-based chilled ceiling system is deployed for sensible cooling. The airflow of the DOAS and FCU is controlled such that the former unit caters for the net indoor ventilation requirement of the entire zone, and the latter caters for interior latent loads. The air supplied by the DOAS and FCU is mixed together before distribution to each room, with motorized diffusers enabling each room to be supplied with variable-air-volume (VAV) control. The sensible cooling in the each room is supplied by passive chilled beams (PCBs), with the chilled water flow rate supplied to each room's PCBs controlled by motorized pressure-independent balancing and control valves (PIBCVs). The benefit of the fully decoupled and decentralized air-conditioning systems is that it is in effect possible to independently control the temperature, humidity, and CO2 concentrations within each room. Whilst setpoints for these factors are currently assigned centrally, a human machine interface (HMI), in this case a liquid crystal display (LCD) mounted in each room, allows occupants in individual office to override the centrally-assigned room setpoint temperature if and when needed.

With regards to sensoring, each office is equipped with motion sensors, indoor climate sensors (i.e. temperature, relative humidity, and CO2 sensors), and one HMI. Motion sensors were installed on the underside of suspended ceiling panels, 2.8 m above the floor. The indoor climate sensors were placed at two height levels, 1.2 m and 2.8 m, and the HMIs were installed at the 1.2 m above the floor. Furthermore, each room's PCB array is monitored by installed energy meters, comprising a water flow meter, and temperature sensors of chilled water supply and return. The meters selected for this experiment comply with EN 1434 and the measuring instruments directive (MID) accuracy class 2 [20], with permissible relative measurement errors smaller than 7%. Overall, the quantities of installed sensors, control valves, and HMIs in each office are provided in Table 1.

2.2. Architecture of the learning-based sensible cooling control system

The learning-based sensible cooling control system in the case study offices is defined by three hierarchical layers: 1) the demand-driven control algorithm, 2) local controllers, and 3) actuators of sensible cooling together with the sensory infrastructure. The system hierarchy is illustrated in the block diagram of Fig. 2.

The major goal of the first layer is to dynamically specify setpoints of room dry-bulb temperature by learning from occupants' historical and current room occupancy. The algorithm developed

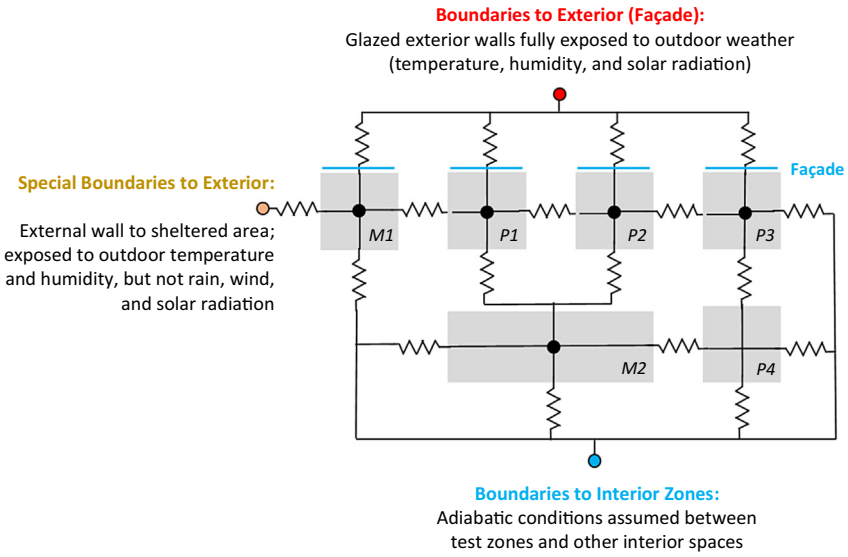

Fig. 1. Thermal resistance-capacitance (RC) diagram of the experimental space.

Table 1
The number of the actuators and sensory infrastructure for the demand-driven cooling control.

Actuators and sensors	Number of them in the offices					
	M1	M2	P1	P2	P3	P4
PIBCV	1	1	1	1	1	1
PCB	8	7	6	6	10	2
Energy meter	1	1	1	1	1	1
HMI	1	1	1	1	1	1
Motion sensor	1	2	1	1	2	1
Temperature, relative humidity, and CO2 sensor	3	3	2	2	3	2

for occupancy learning-based demand-driven cooling control was implemented in this layer, and is described further in Section 4. This algorithm has been coded in MATLAB and executed on a workstation computer that has network access to the office's building management system (i.e. the second layer). Via this network connection, the workstation, and thus the algorithms running on it, have real-time access to sensor data and actuator feedback from the office space. This communication is provided by a Web Service Interface (WSI), using the architecture of REST API (i.e. representational state transfer, application programming interface) for reading data and writing commands. The workstation itself is a desktop computer, with an Intel QuadCore i7-3770 (3.4 GHz) processor, and 12.0 GB of RAM (1600 MHz).

The second layer consists of local controllers with proportional-integral-derivative (PID) control and the office' building management station (BMS). The local controllers in this layer are primarily in charge of adjusting the actuators of the sensible cooling of the third layer to ensure adherence to indoor temperature setpoints. The local controllers send analog signals ranging from 0 V to 10 V to the PIBCVs for controlling the rate of cooling by PCBs. They use KNX and Meter-Bus (M-Bus) communication protocols to collect data from the sensory infrastructure.

The actuators and sensors used in the third layer are presented in Section 2.1, and the quantities of they are listed in Table 1.

3. Measurement and analysis of motion and occupancy datasets

Dynamic information on room occupancy is one of the key requirements for the demand-driven control. In this study, motion sensors were used to observe occupant movements (i.e. activity and inactivity) in the offices, and signals from them were processed to derive occupancy data.

Fig. 3 shows the occupants' stochastic movements during a random working day in the case study offices in 2016. The light blue background indicates the scheduled comfort period defined by the office's BMS. As shown in this figure, not all of the offices appear to have been consistently occupied during the daily comfort period. However, this is so far based on binary “active” or “inactive” signals from the motion sensors. In other words, these sensors can recognize occupants' presence promptly but cannot distinguish between an occupant's absence and sitting near-motionless in the office. A specified value of time delay (TD) is assigned to interpret motion sensor measurements in terms of occupancy. The TD is the minimum time required between motion sensor “active” signals in order to assume an occupant is not physically present [21]. It is a user-defined parameter. To define it for this case study, data from the motion sensors installed in the six rooms were gathered from seven months before the experiments in 2016. Inactive time slices (i.e. durations between “active” motion signals) were extracted from the seven-month motion dataset. Fig. 4 shows the cumulative probabilities of inactivity time slices of the offices. As shown in this figure, short inactive slices (i.e. less than 100 s) account for the majority of the collected data. To some extent, the high cumulative probabilities of short slices of inactivity are attributed to the sampling method employed by the motion sensors: change of value (COV). When a motion sensor detects a movement, a logic ‘high’ state (i.e. 1) is recorded in the air-conditioning control system. This is different from other motion sensors where a delay time setting is incorporated, which requires the logic ‘high’ state to be maintained for a minimum duration after detecting the motion. In all, the COV-based motion sensors implemented in the experiment create more short slices of inactivity. Thus, in order to obtain a high confidence level of verifying occupants' vacancy, 600 s (i.e. 10 min) was used as the TD value, which is consistent with prior research [14,22].

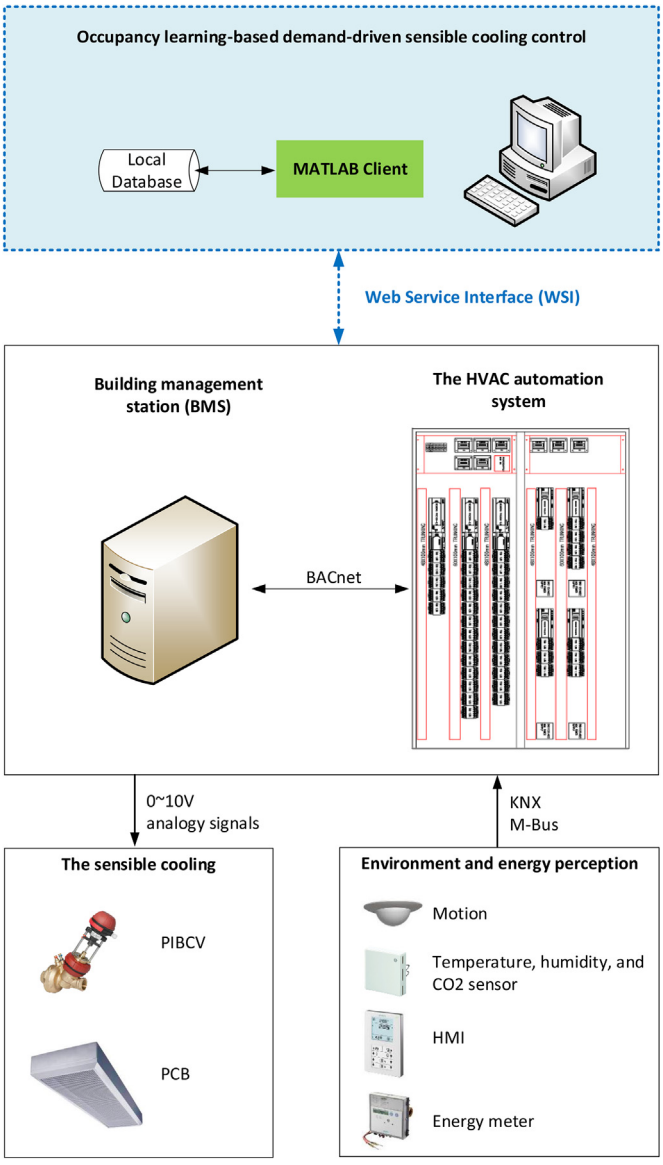

Fig. 2. Architecture of the demand-driven cooling system.

After transforming motion data to occupancy information with 1-min time intervals, occupancy probability distributions over weekdays of the monitored seven months were analyzed. As shown in Fig. 5, three key types of occupancy patterns were observed in the six offices:

- Multi-person offices M1 and M2 have high occupancy rates, and the peak occupancy probability of them is above 90%.
- Private office P4 has a medium occupancy rate, and the peak occupancy probability is above 60%.
- Private offices P1, P2, and P3 have low occupancy rates. Their occupancy rates rarely go over 50%.

As the private offices have much lower occupancy likelihood than the multi-person offices, one could expect that correct assignment of a setback temperature (i.e. higher cooling temperature) during occupants' vacancy periods would be a useful energy-saving measure for these spaces.

Two further important features of the occupancy data were extracted from the seven-month occupancy dataset: the first arrival and last departure. These two metrics can be used to determine the appropriate sensible cooling scheduled operation hours for individual offices.

Fig. 6 illustrates the cumulative probabilities of the first arrivals for the six offices. Approximately 98% of the first arrivals occurred before 8 a.m. in two multi-person offices: M1 and M2. Around 85%–90% of the first arrivals occurred before the same time in the other four private offices. Fig. 7 presents cumulative probabilities of the last departures, which were different in all the six offices. The last departure times appeared to be more stochastic and less predictable, though most of them occurred before the end of the time of the air-conditioning service (i.e. 6 p.m.). In the M1, P1, P2, P3, and P4 offices, some recorded last departures were occurred as early as the morning, likely signifying a lack of any substantial occupant presence for an entire day. In the afternoon, M1, P1, and P2 offices had steep slopes in the cumulative probability curves from 4 p.m. to 5 p.m. In all, as there are days where last departures from all offices occur before 6 p.m., shutting off the sensible cooling system entirely at these earlier times could contribute to energy savings by reducing the daily operating duration.

4. Methodology

4.1. Overview

As shown in Fig. 8, the methodology for the demand-driven control algorithm layer in Fig. 2 includes three modules: 1) occupancy data processing, 2) an algorithm for occupancy learning and prediction, and 3) rule-based control (RBC). The sensible cooling of each office in the test space is controlled by this control frame individually right up to occupants' random presence and vacancy.

As defined by the facility department's operation schedule, the test office space is normally air-conditioned to defined 'comfort' setpoint conditions between 7:00 and 18:00 on weekdays. The proposed demand-driven cooling control model discussed above takes over the overall space's existing central-control by specifying time-dependent room temperature setpoints for individual rooms between 8:00 and 18:00 on weekdays.

During the demand-driven control on weekdays, the first module collects and interprets historical and present occupancy information into datasets to be used by the second module. To do so, the module accesses the space's local controllers for reading real-time signals from the motion sensors installed in the monitored six offices, calling data reads at 1 min time intervals. This data is transformed into the occupancy information using the method described in chapter 3. The data is saved in a vector that represents the current daily occupancy information as time progresses. At the end of the day, the daily occupancy data is also updated into two historical occupancy datasets: weekday occupancy data and weekend occupancy data according to the given date.

After retrieving the occupancy information, the algorithm in the second module infers the likelihood of occupancy presence in future by learning from the profiles of occupant presence and vacancy in past days – i.e., the datasets provided by the first module. Finally, the RBC of the third module infers an appropriate room temperature setpoint for the future hours according to forecasted occupancy information, pre-defined temperature setpoints of comfort and setback modes, and constraints. The setpoints generated by this module is then sent as a command to local controllers, with any existing setpoint overridden by a new value.

The algorithm embedded in the second module and the control strategy employed in the third module are presented in Sections 4.2 and 4.3 respectively.

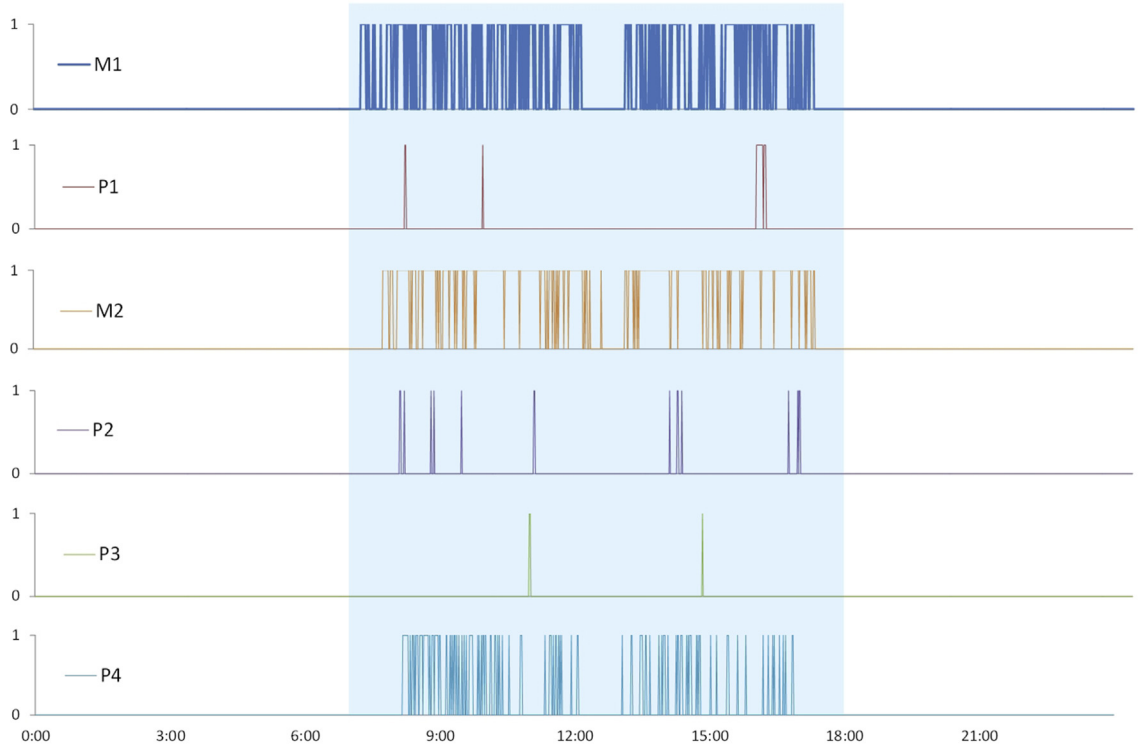

Fig. 3. An example of daily motion signals of the six offices during the experiment in 2016.

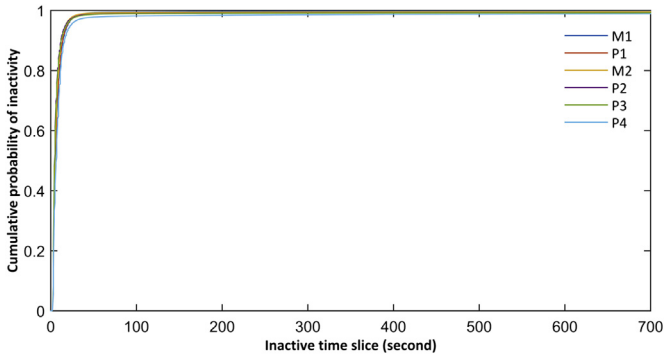

Fig. 4. Cumulative probabilities of inactivity slices in the six offices (seven-month data before the experiment in 2016).

4.2. The algorithm of occupancy learning and prediction

Inspired by the work of Scott et al. [12], K-nearest neighbor (KNN) was employed to predict future occupant behavior in this study. Occupants' behavior is represented in a daily time-frame with random times and duration lengths associated with presence and absence. To some extent, the learning process of the KNN algorithm involves computing the difference between the vector-based, time-independent occupancy data. Scott et al. [12] used KNN to undertake occupancy forecasting for the control of a residential heating system. In this work, we not only use KNN to predict occupants' next presence but also to forecast presence duration in the coming hours for demand-driven control of a cooling system in a commercial building. The formatted occupancy dataset, and the process for defining the parameters of KNN, are therefore different in this work from the strategy employed by Scott et al.

KNN is an algorithm for supervised learning that rests within the wider field of machine learning [23]. It has been recognized as one of the 'top ten' algorithms in the data mining area by the IEEE International Conference [24]. Generally, KNN is used to search which K sets of data from a prepared training set are most similar to a new piece of data. The main label of K-most-closest objects is assigned to new data for the purpose of prediction or classification [23,25].

Datasets are necessary for this data-based learning algorithm. For the prediction of occupancy, datasets are comprised of a training dataset that is selected from historical weekday occupancy data and a piece of new occupancy data that presents existing occupancy information in the current partial day. As shown in Fig. 9, the daily occupancy information of the training set and the current day is formatted to a binary vector at a resolution of 1-min intervals. In each bit of the vector, 1 indicates that the office is occupied, and 0 indicates that the office is not occupied.

KNN is an algorithm based on the geometric concept. It computes the distance between a piece of test data and a group of training data to recognize the proximity to each other [25]. In this study, the Hamming distance is used to determine the K-most-similar neighbors of the current occupancy data from the training set. The basic process is to measure the number of bits within the training vectors and the current occupancy vector at which the corresponding presence (i.e. 1) and vacancy (i.e. 0) are different [12,26]. For instance, taking a small part of elements from two daily occupancy vectors (i.e. a and b) in Fig. 9 is convenient to interpret the Hamming distance calculation process. a is a partial vector in the training set and b is a corresponding vector from the current occupancy data. The sectional vectors a and b as well as the Hamming distance $d(a,b)$ are presented in Eqs (1)–(3), respectively:

$$\boldsymbol{a} = [0,0,0,0,0,0,1,1,1,1,0,0,0] \tag{1}$$

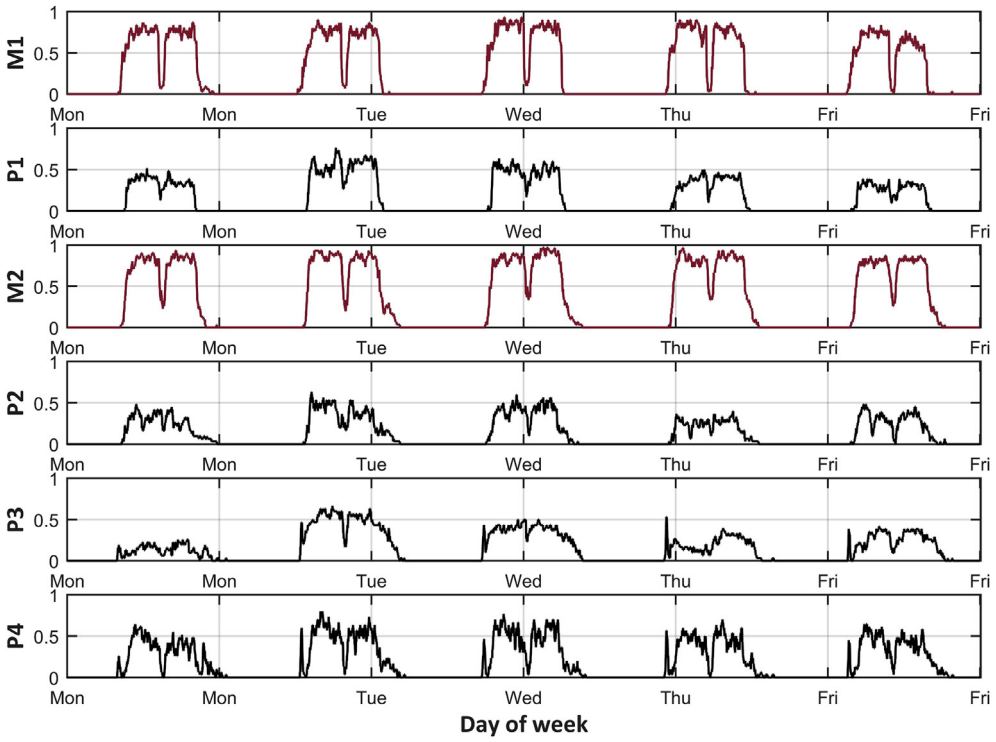

Fig. 5. Weekday occupancy probability distributions (seven-month data before the experiment in 2016).

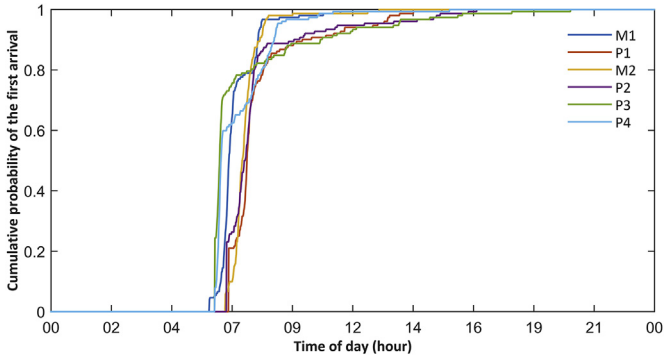

Fig. 6. The cumulative probability distributions of the daily first arrivals (seven-month data before the experiment in 2016).

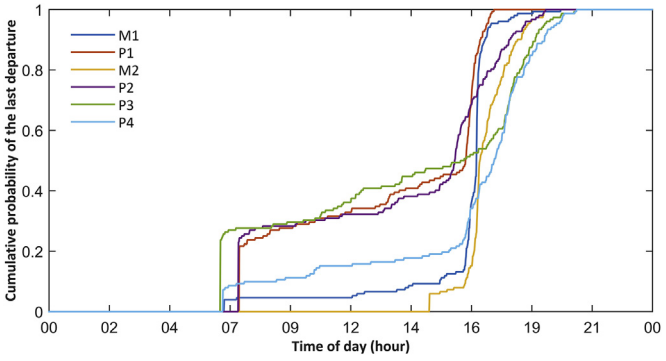

Fig. 7. The cumulative probability distributions of the daily last departures (seven-month data before the experiment in 2016).

$$\boldsymbol{b} = [0,0,0,1,1,1,1,1,1,1,1,1,1] \tag{2}$$

$$\begin{aligned}\boldsymbol{d}(\boldsymbol{a},\boldsymbol{b}) =& \boldsymbol{d}([0,0,0,0,0,0,1,1,1,1,0,0,0],\\ & [0,0,0,1,1,1,1,1,1,1,1,1,1]) = 6\end{aligned} \tag{3}$$

Additionally, another three critical parameters associated with the KNN learning process are also specified in this study, including the size of training dataset, the value of K, and the threshold of the occupancy probability. The first two parameters are integers, and the last one is a fraction.

According to observations from the case study building, occupancy is highly stochastic across the monitored offices. The room owners' daily work and life (e.g. paperwork, meetings, business trips, and annual leaves) impact on the random presence in the offices. Meanwhile, each office appears to be affected by the random presence of room owners and visitors. For instance, when the office owners are not in the offices, the presence collected by the motion sensors also include:

- The staff in other departments accesses the offices for cleaning space, or maintaining office plants and equipment, or security patrol.
- The office owners' colleagues enter into the rooms to collect files or use the space for meetings.

We account for this stochastic behavior by not using static values for the size of training dataset and K, and regard selected K-similar days as equal without assigning weights to them in the learning process. Meanwhile, both the room owners and intermittent visitors are treated as occupants in this study.

The size of the training set (S_{td}) defines the number of past daily occupancy vectors used for searching K similar neighbors. If the value of S_{td} is too large, the KNN learning can not quickly adapt to changes in the occupants' behavior as contexts vary (e.g. holiday seasons, full-load office work, and new employees). If the value of

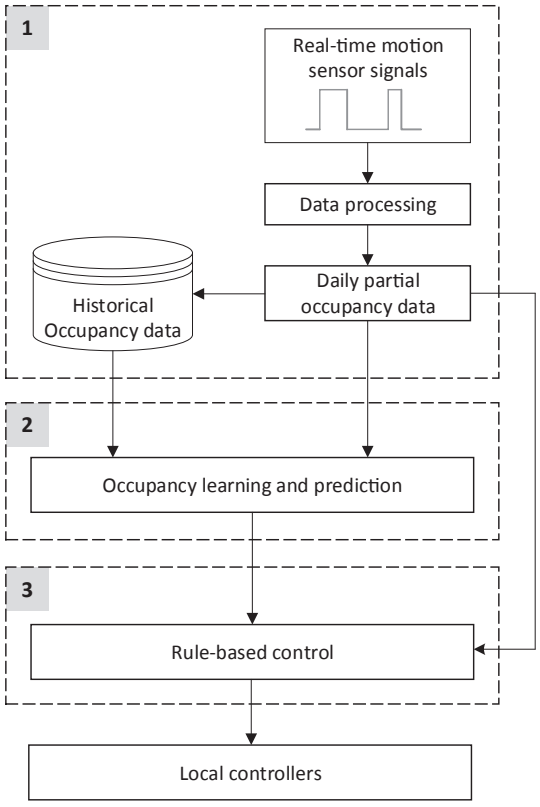

Fig. 8. The block diagram of the occupancy learning-based demand-driven cooling control.

S_{td} is too small, the training set is not enough to efficiently finding out K pieces of past daily occupancy information. To balance these two aspects, in this research we limit the number of occupancy vectors with precenses in the training set to less than 30 and more than 14. However, if a particular office tends to be empty frequently (i.e. the number of empty days is over 20 in past 40 days), the S_{td} is fixed to 45. The detailed mechanism based on pre-established rules is illustrated in Table 2. Actual values of S_{td} are determined by the number of vacant days (i.e. there is no one present) in an office in the past 20 days (N_{20}), past 30 days (N_{30}), and past 40 days (N_{40}).

Similarly, a large value for K (K_{value}) will overstate predicted occupants' presence times and durations. A small value will

Table 2
The mechanism for setting the size of the training dataset and the value of K.

$$\boldsymbol{K_{value}} = \max\left(\boldsymbol{N_{20}}, \frac{\boldsymbol{N_{S_{td}}} * 20}{\boldsymbol{S_{td}}}\right) + 1 \quad (4)$$

The number of days that the individual office is not occupied			The size of the training dataset S_{td}	The value of K K_{value}
N_{20}	N_{30}	N_{40}		
$N_{20} < 5$			20	5
$5 \le N_{20} \le 10$	$N_{30} \le 15$		30	Eq. (4)
$5 \le N_{20} \le 10$	$N_{30} > 15$		40	Eq. (4).
$N_{20} > 10$		$N_{40} \le 20$	40	Eq. (4).
$N_{20} > 10$		$N_{40} > 20$	45	Eq. (4).

underestimate these two parameters. Normally, the K_{value} is chosen by fine tuning via trial-and-error, or is assigned using cross-validation [27]. In this study, when an office is usually occupied, i.e., N_{20}<5, the K_{value} is assigned to 5 in the same manner as used and verified by Ref. [12]. However, as we recognized $K_{value} = 5$ become small as the S_{td} and the number of empty days increase, we summarized a formula for calculating K_{value} in Eq. (4). The purpose of it is to minimise underestimates of occupants' presence and durations. $N_{S_{td}}$ denotes the number of days that a room is not occupied in the selected training dataset.

Together with K_{value}, S_{td} is updated once per day. In the demand-driven control phase, the occupancy prediction algorithm is only triggered when the office is empty. In each minute of unoccupied periods, the KNN recognizes K sets of occupancy vectors from the training dataset with computing the Hamming distance between the current partial occupancy vector and all the elements of the training dataset in the same time span.

Segments of the K-closest occupancy vectors corresponding to remaining minutes in the current day are used to compute the occupancy probability of each bit. After obtaining the occupancy-probability-based vector, a probability threshold ($P_{thrshld}$), which is a constant value that can be defined between zero to one, is required to state presence or vacancy status in individual bit of the probability vector. We set $P_{thrshld}$ to zero because ensuring the thermal comfort for occupants is always the highest objective in this study. This way, presence is marked once the occupancy probability in that bit is greater than zero. For a more energy-conscious purpose, $P_{thrshld}$ can be set to a larger value to

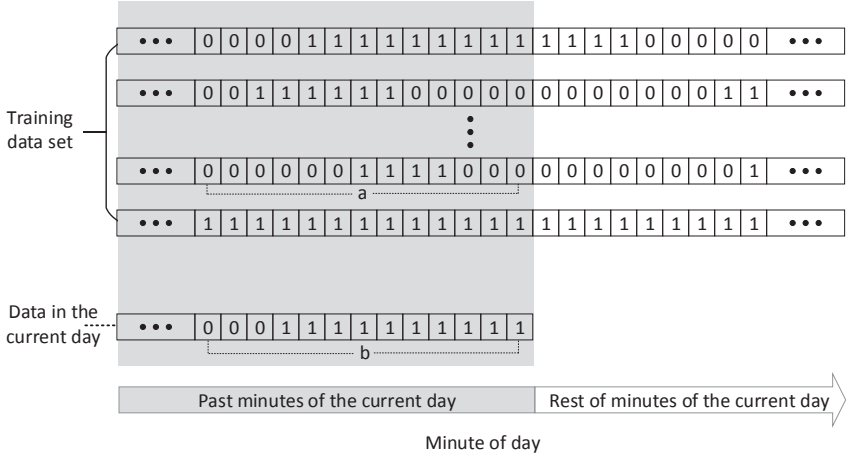

Fig. 9. The occupancy format of the training set and data in the current day.

guarantee a higher confidence of coming occupancy [12].

The remainder occupancy probability vector is transformed to a partial occupancy vector by $P_{thrshld}$. Next, this occupancy vector is regarded as forecasted occupancy information for the rest of the current day. In the last step of the occupancy prediction process, the time of the next presence and the total presence duration in the remaining day are computed from above predicted occupancy vector.

4.3. Rule-based control

The RBC module aims to determine air temperature setpoints over the demand-driven cooling control (DCC) period for the local controllers based on rules specified in this module and the occupancy information obtained from the above two modules.

Three temperature setpoints are defined for the RBC control process: comfort, idle, and economy. The last two temperature setpoints are regarded as setback temperatures, with the former representing the main temperature setpoint during occupied hours. For the six offices, the comfort temperature is fixed at 22.5 °C. The idle temperature is a higher room temperature, and represents the temperature which occupants could tolerate for short durations before becoming uncomfortable. We have assumed that an indoor temperature of 23.5 °C, 1 °C higher than the comfort setpoint, is within an appropriate tolerance to represent the idle temperature of the case study's occupants. This idle temperature is applied to save cooling energy when the offices are not occupied, but there is a modest likelihood that occupants could still arrive in the current day. An economy temperature setpoint of 35 °C is assigned to shut down the cooling system in an office when its occupants do not present for the remainder of a day.

When an office is occupied, the sensible cooling runs in the comfort mode (i.e. 22.5 °C). Conversely, when it is not occupied, the control policy of the room changes to its setback mode (i.e. 23.5 °C and 35 °C). In the DCC, the changeover between room setpoints of the setback mode does not purely depend upon the time of next presence (t_{np}) in a day. In this study, the predicted presence duration for the remaindar or the day (t_{drtn}) is also a key parameter for deducing the setback temperature setpoints. Because as the value of t_{drtn} decreases, space tends not to be occupied over the day. To effectively implement these two predicted occupancy-dependent values into the DCC, we specify prerequisite rules for the change in the current indoor temperature setpoints (T_{sp}) based on statistical analysis (i.e. t_{arr_lmt}, t_{dprtr_lmt} as below), time of the DCC operation (i.e. t_{dcc} and t_{sd} as below), and constraints of presence durations (i.e. t_{drtn_lmt1} and t_{drtn_lmt2} as below). The rules are listed in Table 3, and the unit of t-related values is minute.

The variables t_{dcc}, t_{sd}, t_{arr_lmt}, t_{dprtr_lmt}, t_{drtn_lmt1}, and t_{drtn_lmt2} are configurable and intend to make the DCC adaptable to requirements of diverse application contexts. In this study, they are defined and set as below:

- t_{dcc} is the time at which the DCC starts to specify the time-dependent temperature setpoints for the local controllers. It is set to 480 (i.e. 8:00).
- t_{sd} is the time at which the facility department shuts down the air-conditioning systems in the case study space. It is set to 1080 (i.e. 18:00).
- t_{arr_lmt} is the time at which the cumulative probability of the first arrival is 95%. Similarly, t_{dprtr_lmt} is the time at which the cumulative probability of the last dparture is 95%. Both parameters are computed according to the occupants' behavior in each office, using the same historical occupancy dataset presented in Figs. 6 and 7.
  For a more energy-conscious purpose, the thresholds of cumulative probabilities of the first arrivals and the last departures for t_{arr_lmt} and t_{dprtr_lmt} can be defined to a small value such as from 90% to 95% or more less. This is because small thresholds of these parameters make the economy temperature setpoint enter into the time-frame that is defined in Table 3 earlier. A value higer than 95% is not recommended because the first presence and the last departure in the other 5% are regards as unusual behavior patterns for the DCC.
  Additionally, for defining values of t_{arr_lmt} and t_{dprtr_lmt} according to the cumulative probabilities, it needs to be ensured that t_{arr_lmt} is smaller than t_{sd} and t_{dprtr_lmt} is greater than t_{sd}.
- t_{drtn_lmt1} and t_{drtn_lmt2} are set to 10 and 15, respectively.

According to the pre-conditions defined in Table 3, t_{dcc}, t_{arr_lmt}, t_{sd}, and t_{dprtr_lmt} divide a 24-h day into five segments: $[1,\ t_{dcc}]$, $(t_{dcc},\ t_{arr\_lmt})$, $[t_{arr\_lmt},\ t_{sd})$, $[t_{sd},\ t_{dprtr\_lmt})$, and $[t_{dprtr\_lmt},\ 1440]$. The last four are used here.

In the time between t_{dcc} and t_{arr_lmt}, occupants tend to arrive in the office. When an office is not occupied and the predicted t_{np} falls into this period, the DCC sets the room temperature to 23.5 °C for the setback mode irrespective of whether the presence duration is short or long.

The occupants' last departures are mostly distributed in $[t_{arr\_lmt},\ t_{sd})$ and $[t_{sd},\ t_{dprtr\_lmt})$. Due to the fact that the time of shutting down the air-conditioning is at t_{sd} in the experimental offices, the thresholds of presence durations in these two spans (i.e. t_{drtn_lmt1} and t_{drtn_lmt2}) are defined differently. When the t_{np} lays within these two spans, the forecasted t_{drtn} is determined to limit the office temperature to the idle level or to release it beyond the tolerance by setting the setpoint to 35 °C. For the above two time ranges, two control concepts are highlighted as follows:

1. We assume the occupants are not concerned whether the room temperature is comfortable or not when they are only present in the offices for a short time (i.e. t_{drtn_lmt1}) and leave space in a day. So if the predicted t_{np} is in the time between t_{arr_lmt} and t_{sd}, and t_{drtn} is less than or equal to t_{drtn_lmt1}, the RBC will send

Table 3
The mechanism for specifying temperature setpoints for the comfort and the setback modes.

Occupancy status in an office	Predicted occupancy information		Temperature setpoint (T_{sp})
	Time of next presence (t_{np})	Presence duration of the remaining day (t_{drtn})	
1			22.5 °C
0	$t_{dcc} < t_{np} < t_{arr\_lmt}$		23.5 °C
0	$t_{arr\_lmt} \le t_{np} < t_{sd}$	$t_{drtn} > t_{drtn\_lmt1}$	23.5 °C
0	$t_{arr\_lmt} \le t_{np} < t_{sd}$	$t_{drtn} \le t_{drtn\_lmt1}$	35 °C
0	$t_{sd} \le t_{np} < t_{dprtr\_lmt}$	$t_{drtn} > t_{drtn\_lmt2}$	23.5 °C
0	$t_{sd} \le t_{np} < t_{dprtr\_lmt}$	$t_{drtn} \le t_{drtn\_lmt2}$	35 °C
0	$t_{dprtr\_lmt} \le t_{np} \le 1440$		35 °C
0	No presence		35 °C

commands to the local controllers to change the temperature setpoint to 35 °C.

2. There is no air-conditioning service after t_{sd} even if the occupants come back office after that time. Likewise, the DCC also swich off sensible cooling with PCB valves at t_{sd}, following the policy of the facility department.
   Nonetheless, when the forecasted t_{np} falls into the period that is from t_{sd} to t_{dprtr_lmt} (i.e. the off-service period) and the t_{drtn} is greater than t_{drtn_lmt2}, the DCC changes T_{sp} to 23.5 °C instead of totally shutting down the cooling operation by setting temperature setpoint to 35 °C. Thus, if an occupant will come back during this period and prefer to stay in an office for a longer time, the cooling energy stored in the space by limiting room temperature to 23.5 °C can serve the occupant for a while. The sensible cooling is only shut down when the presence duration is relative short (i.e. less than or equal to t_{drtn_lmt2}) in this period.

As discussed previously, the occupants' presence and absence in the case study offices are stochastic. In order to avoid frequently and unnecessarily cooling down space, the DCC does not change temperature setpoints to the comfort mode from the setback mode when the presence duration is less than 3 min. To achieve this, we integrate a function to recognize current presence as a short stay or not. After 12 min, the setback mode is changed to comfort mode when presence is not a short stay.

5. Results

5.1. Contexts of the experiments

For this research, two sets of experiments were carried out over two months in the six offices (i.e. S2 includes P1 to P4 and M1 to M2): a baseline test where no DCC is employed, called the local control, and a test where the proposed DCC is employed. Each test lasted 21 weekdays. The overriding daily operational period of the offices' air-conditioning system, and the comfort setpoint temperature, were kept the same for both tests.

The local control employed in the baseline test represents the standard sensible cooling control strategy embedded in the BMS of the case study building. However, the sensible cooling energy consumed in the baseline test is not used as the energy benchmark directly to evaluate energy savings achieved by the DCC control test as the indoor and outdoor conditions during two tests are different. Assessing energy savings for different controls is a complex task [28] due to varying contexts.

On the basis that outdoor climate and room occupancy are two key factors that influence the energy use of HVAC systems [15,29]. It is intended in this study to normalize metered energy consumption above both the baseline test and the DCC test against these two conditions. Fig. 10a and b show distributions (box plot) and mean values (cross) of outdoor daily average dry-bulb temperature and horizontal solar radiation in the experimental periods, respectively. Weather data was gathered from an on-site weather station located on the roof of the case study building. The means of the outdoor daily average dry-bulb temperature and the horizontal solar radiation in the DCC test are both greater than that of the baseline test.

In Fig. 11, distributions (box plot) and means (cross) of daily occupied durations in the six offices of S2 differ both in the baseline test and the DCC test. On the whole, the occupancy duration of the six offices during the DCC test is also greater than that of the baseline test.

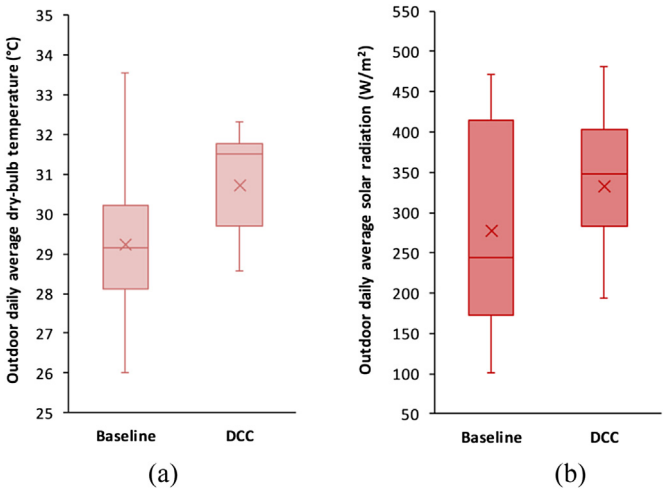

Fig. 10. (a) Outdoor daily average dry-bulb temperature distributions (box plot) and mean values (cross) (b) Outdoor daily average horizontal solar radiation distributions (box plot) and mean values (cross).

The normalization method of cooling degree-days (CDD) is mainly based on weather variation [30], and there is limited literature for the normalization of other variations such as occupancy. In this study, another six offices (S1) with the same contexts and experimental setup are employed to evaluate sensible cooling energy gaps between two experimental months. The features of S2 and S1 include:

- S1 and S2 are equipped with the same air-conditioning and sensor systems.
- S1 and S2 are on the same floor and are close to each other.
- S1 and S2 have the same number of offices adjacent to and not adjacent to façades: four and two, respectively.
- S1 and S2 had higher occupancy rates in the second month: these increased 1.44% and 21.9%, respectively.

The control strategies implemented in S1 and S2 are listed in Table 4, and the number of days of the week in the first and the second twenty-one days are same.

For the six offices of S1 with the same baseline control in two months, the daily average sensible cooling consumption in the second 21 weekdays is higher than the first 21 weekdays: increasing 0.05 kW/m^2 for the offices adjacent to the façade, and increasing 0.03 kW/m^2 for the offices not adjacent to the façade. These two values are used to calculate the sensible cooling energy benchmark for each room in S2 according to Eq (5).

$$\boldsymbol{E_{nbl} = E_{bl} + AS_r} \tag{5}$$

Where

$$A = \begin{cases} 0.05\ [kW/m^2], \text{the room adjacent to the façade (i.e. M1, P1, P2, and P3)} \\ 0.03\ [kW/m^2], \text{the room not adjacent to the façade(i.e. M2, P4)} \end{cases}$$

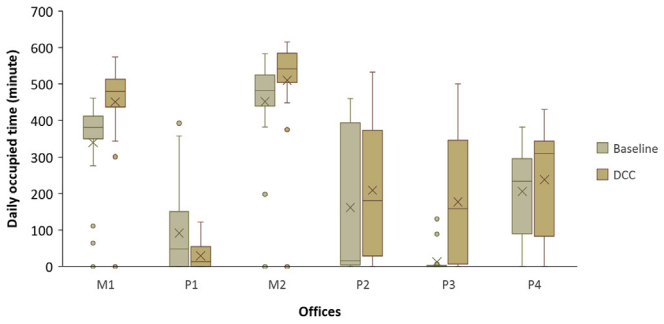

Fig. 11. Daily occupancy durations (box plot) and mean values (cross).

Table 4
The control strategies are deployed in S1 and S2 during experimental months.

	Control strategies	
	S1	S2
1st twenty-one weekdays	The baseline test	The baseline test
2nd twenty-one weekdays	The baseline test	The DCC test

In Eq (5), E_{nbl} and E_{bl} are respectively the normalized and measured daily average cooling energy use of a room in S2 during the baseline test, the unit of which is kWh. S_r is the area of a room in S2, the unit is m^2.

The following subsections will present test results associated with the proposed DCC in details.

5.2. Energy savings

Each office has one thermal energy meter installed to record the real-time sensible cooling energy load, and the energy use of the sensible cooling is calculated from them. For the evaluation of energy savings achieved by the DCC in S2, the benchmark of sensible cooling use in each office is calculated according to Eq (5), and energy use in the DCC test is extracted from measured cooling energy data. Fig. 12 shows the daily average sensible cooling energy consumption of the benchmark and the DCC in relation to the six offices of S2, and Table 5 presents energy savings of them and the whole space of S2.

The test results showed that energy savings achieved by the DCC were closely inversely correlated to occupancy rates within the individual offices. The energy savings in the offices with low or

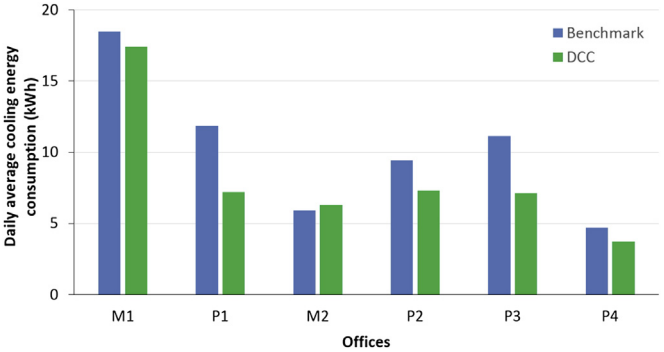

Fig. 12. Normalized daily average sensible cooling energy use in each office during the experiments.

Table 5
Sensible cooling energy savings of the offices and the whole space of S2 in the DCC test.

	Offices in S2						The whole space of S2
	M1	P1	M2	P2	P3	P4	
Energy saving (%)	5.7	39.4	−6.1	22.3	36.0	21.2	20.3

medium occupancy rates (i.e. P1 to P4) reached 21.2%–39.4% by the DCC. As the occupancy rates increase, the energy savings were less. For M1, with long daily average occupancy durations (i.e. from 5.7 h to 7.5 h), the energy savings achieved by the DCC was 5.7%. It is anticipated that if the daily average occupancy time increase to a duration close to or greater than the scheduled air-conditioning operating hours, there would be no energy-savings achievable by the DCC, such as the room M2.

Overall, for the six offices in S2, the proposed DCC strategy saved 20.3% sensible cooling energy.

5.3. Control performance

Two examples of the DCC test in the private offices P2 and P3 are shown in Fig. 13 and 14 respectively. Each figure depicts four curves: office occupancy, time-dependent room T_{sp}, measured room average temperature (T_{air}), and PCB cooling energy. In the occupancy curves, short presence of less than 3 min during the setback mode are drawn in gray. Likewise, for the PCB cooling power curves, the periods of the setback mode that are set T_{sp} to 23.5 °C or 35 °C are in the light and dark green backgrounds, respectively. For the DCC during 8:00 to 18:00, Figs. 13 and 14 cover the main scenarios defined earlier in Table 3. They also present the successful operation of the comfort mode and the setback mode according to the occupants' presence and vacancy in the offices. The DCC also effectively ignores occupants' quick presence in the gray curve during the setback mode. If the occupants are not in their offices, the DCC sets the indoor temperature to 23.5 °C or 35 °C in the setback mode according to the RBC rules associated with prediction occupancy data. The main observation from the DCC test, as shown in Figs. 13 and 14, is that chilled water supply to the PCBs in the individual offices is reduced, or shut off, during unoccupied periods. This illustrates the main mechanism by which the DCC achieves energy savings.

The predicted next-presence time and presence duration for the remainder of a day are used indirectly to infer setback temperature setpoints according to the pre-conditions described in Table 3: predicted presence falls into which time segment and forecasted occupancy duration is greater or less than which threshold. Therefore, to quantify the control accuracy of the DCC based on the predicted occupancy, three key criteria that potentially affects occupants' thermal comfort or tolerance are extracted from Table 3:

1 The DCC should switch the setback mode to the comfort mode when occupants are back to the office with a stay that is greater than 2 min before t_{sd}.
2 The DCC should set the setback mode with 23.5 °C when an office is empty and will be occupied before t_{sd} with a duration that is greater than t_{drtn_lmt1}.
3 The DCC should set the setback mode with 23.5 °C when an office is empty and will be occupied between t_{sd} and t_{dprtr_lmt} with a duration that is greater than t_{drtn_lmt2}.

To calculate the overall accuracy of the DCC, the daily control accuracy is set to zero when temperature setpoints reasoned by the

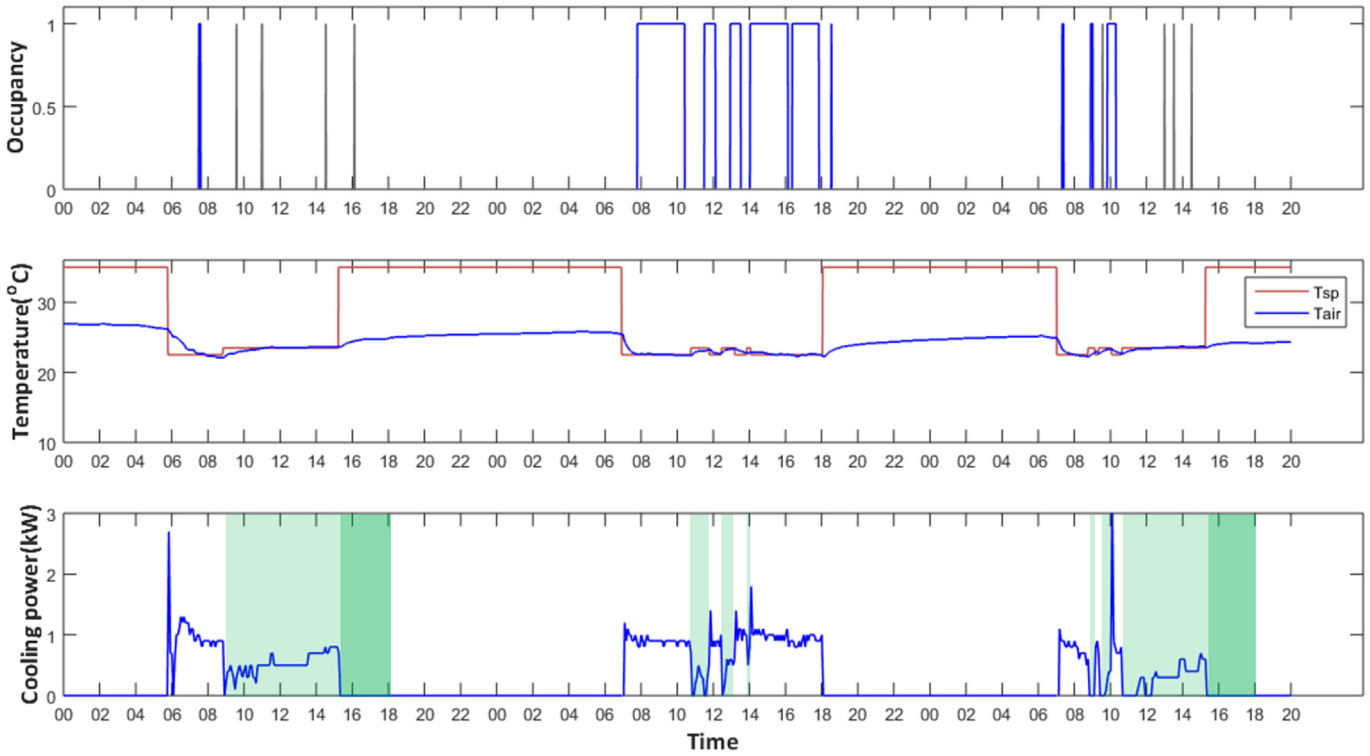

Fig. 13. Three-day sensible cooling control in the private office P2 during the DCC test (explained in the above text).

DCC do not meet the above criteria, even though it only happens in a control cycle (i.e. 1 min) in a day and the air temperature of space is still in the occupants' thermal tolerance (i.e. between 22.5 °C and 23.5 °C). On the other hand, the daily control accuracy is set to one when the DCC achieves the three criteria during all minutes of the DCC control period in a day.

Table 6 presents the control accuracy of the DCC in correctly anticipating occupant behavior and expectations. It shows the

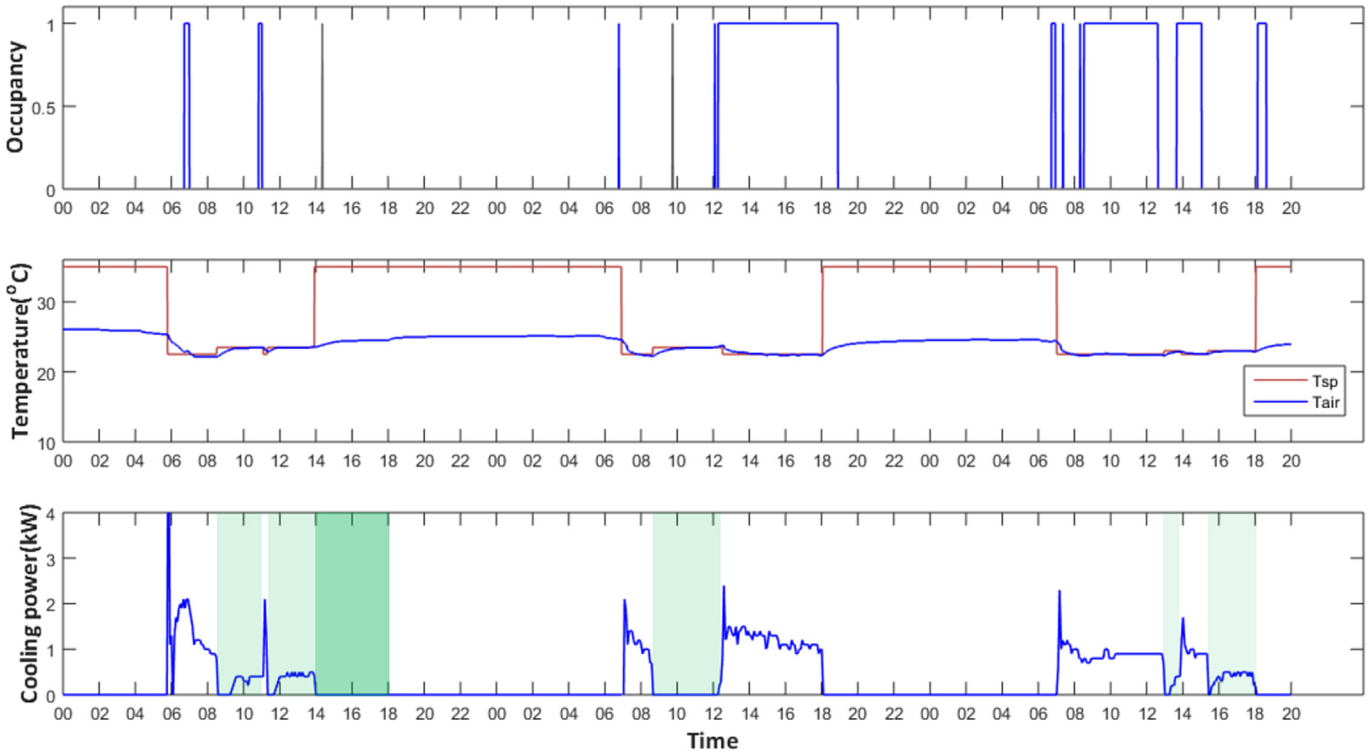

Fig. 14. Three-day sensible cooling control in the private office P3 during the DCC test (explained in the above text).

Table 6
The accuracy of the DCC algorithm as tested.

Offices	M1	P1	M2	P2	P3	P4	Total
Control accuracy (%)	95.2	81.0	100.0	81.0	90.5	81.0	88.1

Table 7
The number of days that the DCC did not meet the criteria.

	The number of days					
	M1	P1	M2	P2	P3	P4
The 1st criterion	0	0	0	0	0	0
The 2nd criterion	1	4	0	4	1	4
The 3rd criterion	0	0	0	0	1	0
Overall	1	4	0	4	2	4

percent number of days, out of 21 tested weekdays, that the DCC satisfies the three criteria at all times. The results are shown on a per-room basis. Overall, the DCC shows good control performance on the multi-person offices M1 and M2 as well as the private office P3, and accuracy rates of them are over 90%. The control performance of another three private offices P1, P2, and P4 reached an intermediate level: 81.1%. The average control accuracy of the entire tested space was 88.1%. In order to analyze how daily occupancy patterns affect decision makings of the DCC, the above control accuracy is calculated in days. This will increase to 90.5% if we use the percent number of hours, versus days.

As the stochastic presence of occupants in a multi-person office overlaps with one another, intermittent periods of occupancy are relatively rare for the M1 and M2 spaces versus others. This also suggests why the DCC accuracy in M1 and M2 is higher than that of the single person offices.

To look deeper into the accuracy of DCC, Table 7 specifies the number of days that the DCC did not fulfill the three criteria. The largest cause of error in this period were times which the DCC incorrectly set the setback temperature to 35 °C where the idle setback temperature 23.5 °C would have been required. This occurred when:

- An office would become occupied randomly towards the end of the working day.
- An office would become irregularly occupied for short periods.

Taking the office P4 as an example, the space in all four days that DCC did not meet the second criterion is only occupied for short periods. According to our records, these short durations were triggered by visitors and the room owner of office P3, because both needed to pass the office P4 to get access to the office P3. As this type of short presence is easier to deviate from occupancy patterns in the training set, the KNN learning tends to regard it as no presence in the space.

This reasoning can be used to apply to the days that did not meet the second and the third criteria in all test offices. When the occupancy pattern of the current day does not belong to the predominant patterns or deviates from them in the training dataset, the DCC in this paper set the room temperature according to patterns of the main scenarios that are most close to the pattern of the current day.

5.4. Room temperatures of occupied periods

As well as improving the energy efficiency of the office space, assuring that the occupants of the test spaces would be thermally comfort is a primary objective of this research. Each room in the test is equipped with several temperature sensors to monitor real-time air temperatures, as well as an HMI to let the occupants view their room's temperature and adjust their room's temperature setpoint according to their own thermal preferences.

For analyzing temperatures during occupied periods of the six offices of S2, the average air temperature of each occupied slice during the baseline and DCC test is calculated from the raw dataset. The room temperature cannot be maintained at a fixed setpoint constantly, and is usually kept within a small band compared to the setpoint (e.g. $\pm 0.5\ ^{\circ}C + T_{sp}$). Three temperature levels are employed here to compare room temperatures in the occupied slices during the baseline and DCC test: 22.5 °C, 23 °C, and 23.5 °C. In this research, the comfort setpoint is defined as 22.5 °C, the 23 °C is the assumed upper boundary of the comfort setpoint, and the assumed temperature tolerance for the setback mode is 23.5 °C.

Fig. 15 illustrates percentages of total occupied durations, in terms of minutes, for room temperatures less than or equal to the three levels. The occupied durations of room temperatures less than or equal to the comfort boundary (i.e. 23 °C) achieved 97.8% and 93.6% in the baseline and DCC test respectively, with a small difference between them (i.e. 4.1%). In the DCC test, room temperatures above 23.5 °C during the occupied periods account for less than 0.5% of the total occupied duration, and all of them are less than 17 min.

For investigating whether occupants changed temperature setpoints to adjust the controlled room temperature during the DCC test, temperature setpoint shifts modified by the occupants were collected from the six HMIs. It was found that the temperature setpoint shift was done only once throughout the entire DCC test period. One occupant in M2 changed it to the 22.5 °C from the setback mode, though the air temperature in that time was 22.7 °C with only a 0.2 °C deviation. As presented in Section 4, to avoid frequently cooling down space relative to the quick presence, the DCC delayed 12 min to make a decision whether or not to change room temperature back to the comfort mode. This temperature adjustment in M2 occurred in the 12-min period. It is possible to overcome this situation by disabling the function of the quick presence detection and making the DCC switch the comfort mode back as soon as a motion sensor is triggered.

5.5. Computation time

The time step of the DCC is set to 1 min, and the computation process during a cycle is broken down into four steps: 1) reading sensor data, 2) processing data, 3) occupancy prediction and RBC, and 4) writing commands.

In each control cycle, the computation time of the above four steps was recorded into a log to evaluate the computational resources for such control. Fig. 16 shows the mean computing time of

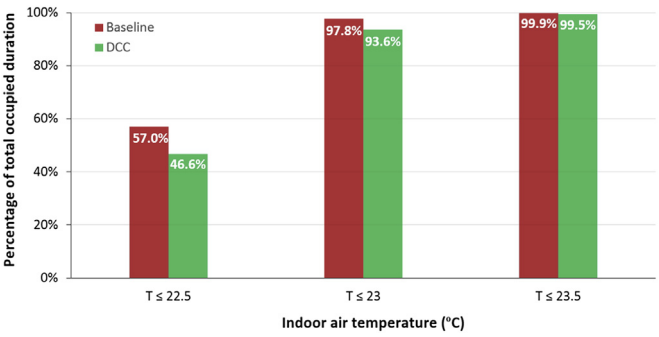

Fig. 15. Office temperatures during the occupied time in the baseline and DCC test.

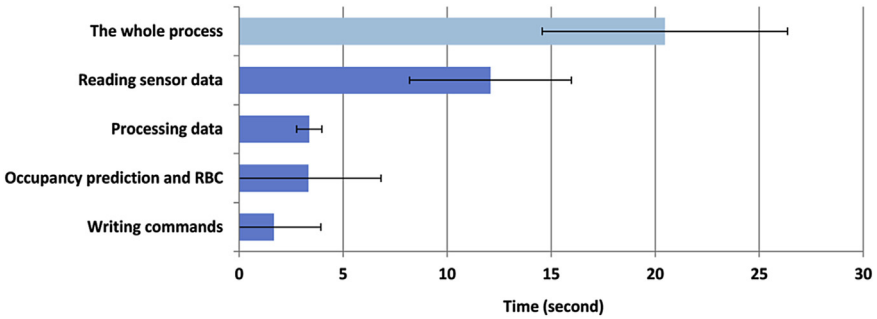

Fig. 16. The mean computation time of each computational step and the overall process in a control cycle.

the four phases together with the overall process. It can be observed that the DCC computation only needed 20.5 s in one cycle, and the remaining 39.5 s would be a waiting period for entering into the next control loop. Among the computational stages, reading sensor data accounted for the largest percentage of the computation time — 59%. The occupancy prediction and RBC that are concerned most is just used 3.3 s with a small share — 16.5%, similar to the time consumed by the data processing.

Normally, the KNN algorithm requires a comparatively expensive computational resource due to calculating the distance between new data and all elements in a training dataset, especially when the training dataset has a large size [23,24]. For predicting occupancy, the Hamming distance for vectors can be computed quickly. Meanwhile, the size of the occupancy training dataset is also limited to a small amount, considering the adaptation to changes in occupants' behavior. Overall, the computing time of the proposed DCC was adequate for this study.

6. Discussion and conclusion

In this paper, we presented a control strategy with learning capacity for demand-driven sensible cooling and implemented it in six offices of a case study building in Singapore. Test results of the control strategy have been provided and analyzed in this paper.

The proposed demand-driven control can not only be used for sensible cooling systems but also can be implemented into heating systems and ventilation systems. For demand-driven heating, the control process is exactly the same as the proposed cooling control, only using different temperature setpoints for the comfort mode and the setback mode. For the ventilation control, time-dependent setpoints can be employed to regulate the room humidity or CO2 level.

The following subsections draw out the discussion and conclusions of this work and lay out areas of future work.

6.1. Energy savings potential

The DCC employed in this work can automatically adapt to current and future occupancy states of different offices for saving energy efficiently and maintaining room temperature during the occupied periods with similar performance as the baseline system. The results showed that the DCC contributed to reducing overall sensible cooling energy requirements by 20.3% over the test period.

Our experimental result indicates that energy use in the office spaces of the case study building could be reduced by decreasing unnecessary sensible cooling requirements during non-occupied hours. The results also showed that achieved energy savings were inversely correlated to occupancy rates within the individual case study offices.

For example, occupancy rates during both baseline and DCC tests are low in the single person office P3, as shown in Fig. 11. During the baseline test, this office had a lower occupancy rate, though a comfortable temperature level was maintained throughout the entire daily air-conditioning system operating period. In contrast to the baseline test, whilst this office was occupied with a higher occupancy rate in the DCC test, the DCC made the sensible cooling adapt to its occupancy behavior. In doing so, the DCC test attained an energy reduction of 36% in this office. Similar energy savings were attained in the other single person offices with low occupancy rates (i.e. P1, P2, and P4). Generally, these findings are supported by related studies in the building sciences. For instance, Masoso et al. [31] reported that energy waste in commercial buildings was mainly caused by equipment kept in operation during unoccupied periods. The DCC has addressed this by ensuring that the operation of sensible cooling in the case study building follows the stochastic presence of occupants.

Meanwhile, even though the DCC achieved high control accuracy for two multi-person offices, the energy savings potential, in this case, is limited by when experiencing high occupancy rates. During the DCC test, office M2 displayed even higher energy consumption than the baseline as the occupancy rate increased further.

6.2. The control performance

The average control accuracy of the DCC based on the predicted occupancy data attained in the six offices was 88.1% according to specified criteria. Meanwhile, the controlled room temperatures above assumed thermal tolerance (i.e. 23.5 °C) during occupied time only represented 0.5% of the total occupied durations, and all were only short occupied slices (i.e. less than 17 min). Under the tested conditions, we did not observe that the occupants used the HMI to change the setpoints except one situation described in Section 5.3.

Two of reviewed systems [12,13] illustrated the control accuracy based on the forecasted occupancy, as presented in Section 1. In contrast to their research on residential buildings, our control strategy attained a similar control performance for demand-driven cooling of office spaces. To further improve the control performance, more learning algorithms such as [32,33] can be explored and embedded into the current strategy to analyze occupancy patterns for the occupancy learning and prediction module, and more statistical analysis can be integrated into the RBC module. These will be explored in our future development of the DCC strategy.

6.3. Computation time

The assessment of the computation time illustrated in this paper gives a rough view of the computing distribution across the DCC structure based on our experimental system and raises possibilities

to reduce the computational time of the DCC process further in the development and engineering phase.

The communication between the workstation that the DCC algorithm has been coded and the BMS computer is only evoked by the client side (i.e. the DCC workstation) for reading sensor data and sending control commands with the REST API communication. Therefore, the time step of the DCC was defined to 1 min to make the sensible cooling system respond quickly to changes in occupants' behavior in the research phase. The average computation time in each cycle is calculated for evaluating the possibility of future applications. For the proposed DCC, the average computation time in one control cycle was 20.5 s for the six offices. That is, the sensible cooling control in each office only required about 3.4 s. Adding demand-driven ventilation into the current DCC would not require more computation time.

Due to limits of computing and memory, complicated algorithms may be difficult to integrate into local controllers [34]. For controlling a much larger number of offices in commercial buildings, it is possible to embed such DCC into a computer-based BMS itself working as an application function. Collecting sensor data from the BMS, and transmitting it to the DCC via REST API, took 59% of the total computational time in the undertaken experiment. Once the DCC is integrated into the BMS workstation, the 1-min inter-computer data sampling process would no longer be needed. Further computational savings could be found by increasing the time resolution of the occupancy learning and prediction cycle and optimizing the local data reading process. The best values of them need to be evaluated in the development and engineering phase.

6.4. Occupancy detection

In this study, movement signals are only transformed to occupancy data with a binary format, and the proposed DCC makes the cooling operation adapt to such binary behavior. In a room with a larger maximum capacity (e.g. big open office) compared to our case study rooms, there are usually more than two motion sensors. In such a space, each motion sensor only monitors movements occurring in a designated floor area, and multi-motion signals collected from this space can be processed to discrete values to represent what proportion of the room is occupied or what proportion of people are in the room. This data can be used to infer thermal loads generated by occupants and their personal equipment for the control of PIBCVs, VAV boxes or other control components of an air-conditioning system.

6.5. Limitations of the experiments

The primary purpose of the commercial building used in this study is not academic research, so the baseline control and DCC were implemented in the 6 offices with 10 occupants over 2 months with 42 weekdays. Ideally, such an experiment would be beneficial to be conducted over a longer time horizon of several months. The small number of the case study offices and the short experimental time may restrain the universality of the results presented in this paper. However, even in this small experimental scale, the offices assessed by our experiment captured large variations in occupants' behaviors across various private offices and multi-person offices. This was sufficient to illustrate a correlation between energy savings potential and occupancy rates as well as demonstrate how a demand-driven control strategy could reduce building energy consumption in practice.

Acknowledgements

The authors would like to thank Siemens Building Technologies and the United World College South East Asia for their support and contributions.

References

[1] IPCC, Contribution of Working Groups I, II and III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change [Core Writing Team, in: R.K. Pachauri, L.A. Meyer (Eds.), Climate Change 2014: Synthesis Report, IPCC, Geneva, Switzerland, 2014, p. 151.
[2] IPCC, Contribution of Working Group III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change, in: O. Edenhofer, R. Pichs-Madruga, Y. Sokona, E. Farahani, S. Kadner, K. Seyboth, A. Adler, I. Baum, S. Brunner, P. Eickemeier, B. Kriemann, J. Savolainen, S. Schlömer, C. von Stechow, T. Zwickel, J.C. Minx (Eds.), Climate Change 2014: Mitigation of Climate Change, Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA, 2014.
[3] Transition to Sustainable Buildings, Strategies and Opportunities to 2050, International energy agency, 2013.
[4] H. Mirinejad, K.C. Welch, L. Spicer, A review of intelligent control techniques in HVAC systems, IEEE Energytech 2012 (2012) 1–5, http://dx.doi.org/10.1109/EnergyTech.2012.6304679.
[5] K.J. Chua, S.K. Chou, W.M. Yang, J. Yan, Achieving better energy-efficient air conditioning – a review of technologies and strategies, Appl. Energy 104 (2013) 87–104, http://dx.doi.org/10.1016/j.apenergy.2012.10.037.
[6] 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, http://dx.doi.org/10.1016/j.enbuild.2014.09.046.
[7] T.A. Nguyen, M. Aiello, Energy intelligent buildings based on user activity: a survey, Energy Build. 56 (2013) 244–257, http://dx.doi.org/10.1016/j.enbuild.2012.09.005.
[8] T. Labeodan, W. Zeiler, G. Boxem, Y. Zhao, Occupancy measurement in commercial office buildings for demand-driven control applications—a survey and detection system evaluation, Energy Build. 93 (2015) 303–314, http://dx.doi.org/10.1016/j.enbuild.2015.02.028.
[9] Y. Agarwal, B. Balaji, R. Gupta, J. Lyles, M. Wei, T. Weng, Occupancy-driven energy management for smart building automation, in: Proc. 2Nd ACM Workshop Embed. Sens. Syst. Energy-effic. Build, ACM, New York, NY, USA, 2010, pp. 1–6, http://dx.doi.org/10.1145/1878431.1878433.
[10] V.L. Erickson, A.E. Cerpa, Occupancy based demand response HVAC control strategy, in: Proc. 2Nd ACM Workshop Embed. Sens. Syst. Energy-effic. Build, ACM, New York, NY, USA, 2010, pp. 7–12, http://dx.doi.org/10.1145/1878431.1878434.
[11] S. Wang, Z. Ma, Supervisory and optimal control of building HVAC systems: a review, HVACR Res. 14 (2008) 3–32, http://dx.doi.org/10.1080/10789669.2008.10390991.
[12] J. Scott, A.J. Bernheim Brush, J. Krumm, B. Meyers, M. Hazas, S. Hodges, et al., PreHeat: controlling home heating using occupancy prediction, in: Proc. 13th Int. Conf. Ubiquitous Comput, ACM, New York, NY, USA, 2011, pp. 281–290, http://dx.doi.org/10.1145/2030112.2030151.
[13] J. Lu, T. Sookoor, V. Srinivasan, G. Gao, B. Holben, J. Stankovic, et al., The smart thermostat: using occupancy sensors to save energy in homes, in: Proc. 8th ACM Conf. Embed. Networked Sens. Syst, ACM, New York, NY, USA, 2010, pp. 211–224, http://dx.doi.org/10.1145/1869983.1870005.
[14] Burak Gunay H, O'Brien W, Beausoleil-Morrison I. Development of an occupancy learning algorithm for terminal heating and cooling units. Build Environ n.d. http://dx.doi.org/10.1016/j.buildenv.2015.06.009.
[15] F. Nicol, M. Humphreys, S. Roaf, Adaptive Thermal Comfort: Principles and Practice, first ed., Routledge, London, New York, 2012.
[16] EnergyPlus n.d. https://energyplus.net/.
[17] B. Bordass, R. Cohen, M. Standeven, A. Leaman, Assessing building performance in use 2: technical performance of the Probe buildings, Build. Res. Inf. 29 (2001) 103–113, http://dx.doi.org/10.1080/09613210010008027.
[18] A.E. Ruano, S. Pesteh, S. Silva, H. Duarte, G. Mestre, P.M. Ferreira, et al., The IMBPC HVAC system: a complete MBPC solution for existing HVAC systems, Energy Build. 120 (2016) 145–158, http://dx.doi.org/10.1016/j.enbuild.2016.03.043.
[19] A. Schlueter, A. Rysanek, C. Miller, J. Pantelic, F. Meggers, M. Mast, et al., 3for2: realizing spatial, material, and energy savings through integrated design, CTBUH J. (2) (2016) 40–45.
[20] EN Standard, EN-1434:Thermal energy (Btu meter or heat meter) n.d.
[21] Z. Nagy, F.Y. Yong, M. Frei, A. Schlueter, Occupant centered lighting control for comfort and energy efficient building operation, Energy Build. 94 (2015) 100–108, http://dx.doi.org/10.1016/j.enbuild.2015.02.053.
[22] Dong B. Sensor-Based Occupancy Behavioral Pattern Recognition for Energy and Comfort Management in In[TEL]ligent Buildings. n.d.
[23] P. Harrington, Machine Learning in Action, first ed., Manning Publications, Shelter Island, N.Y, 2012.
[24] X. Wu, V. Kumar, J.R. Quinlan, J. Ghosh, Q. Yang, H. Motoda, et al., Top 10 algorithms in data mining, Knowl. Inf. Syst. 14 (2008) 1–37, http://dx.doi.org/10.1007/s10115-007-0114-2.
[25] A Course in Machine Learning n.d. http://ciml.info/(Accessed 23 December 2016).

[26] D.J.C. MacKay, Information Theory, Inference and Learning Algorithms, first ed., Cambridge University Press, Cambridge, UK, New York, 2003.
[27] D. Barber, Bayesian Reasoning and Machine Learning, Cambridge University Press, New York, NY, USA, 2012.
[28] J. Široký, F. Oldewurtel, J. Cigler, S. Prívara, Experimental analysis of model predictive control for an energy efficient building heating system, Appl. Energy 88 (2011) 3079–3087, http://dx.doi.org/10.1016/j.apenergy.2011.03.009.
[29] J.E. Braun, C. Johnson, Reducing energy costs and peak electrical demand through optimal control of building thermal storage, ASHRAE Trans. Am. Soc. Heat. Refrig Air-Cond Eng. U. S. 96 (1990) 2.
[30] J.A. Eto, Comparison of Weather Normalization Techniques for Commercial Building Energy Use, 1987.
[31] O.T. Masoso, L.J. Grobler, The dark side of occupants' behaviour on building energy use, Energy Build. 42 (2010) 173–177, http://dx.doi.org/10.1016/j.enbuild.2009.08.009.
[32] X. Liang, T. Hong, G.Q. Shen, Occupancy data analytics and prediction: a case study, Build. Environ. 102 (2016) 179–192, http://dx.doi.org/10.1016/j.buildenv.2016.03.027.
[33] S. D'Oca, T. Hong, Occupancy schedules learning process through a data mining framework, Energy Build. 88 (2015) 395–408, http://dx.doi.org/10.1016/j.enbuild.2014.11.065.
[34] ASHRAE, 2011 ASHRAE handbook - HVAC applications - IP. Atlanta, Ga.: refrigerating and air-conditioning engineers, Am. Soc. Heat. (2011).