Building and Environment 93 (2015) 71–85

Contents lists available at ScienceDirect

Building and Environment

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

Development of an occupancy learning algorithm for terminal heating and cooling units

H. Burak Gunay [a], William O'Brien [a, *], Ian Beausoleil-Morrison [b]

[a] Department of Civil and Environmental Engineering, Carleton University, Canada
[b] Department of Mechanical and Aerospace Engineering, Carleton University, Canada

ARTICLE INFO

Article history:
Received 21 March 2015
Received in revised form
5 June 2015
Accepted 7 June 2015
Available online 11 June 2015

Keywords:
Occupancy
Learning controls
Automation in buildings
Energy use

ABSTRACT

A significant portion of the North American workforce reports having the ability to alter their daily arrival and departure times for work. As a result, personal preferences translate into individual occupancy profiles. To accommodate these diverse personal schedules, building operators tend to use conservatively short temperature setback periods. In this paper, a year's worth of data gathered by motion sensors placed in private offices in an academic building were analyzed. The predictability of the recurring occupancy patterns was assessed. Drawing upon this, an adaptive occupancy-learning control algorithm which learns the arrival and departure times recursively and adapts the temperature setback schedules accordingly, was developed. Later, the algorithm was implemented in the Energy Management System (EMS) application of the building performance simulation (BPS) tool EnergyPlus. Simulations conducted with this tool indicate that a 10–15% reduction in the space heating and cooling loads can be achieved by applying individual and dynamically evolving temperature setback periods.

1. Introduction

As reported by Bordass, et al. [1], many office occupants today tend to have flexible work hours. For example, 27% of the workforce in the United States were reported having flexible work hours in 2004 — up from 15% in 1995 [2–4]. Only two in five of the workforce in Canada has a standard work schedule [5]. As a result, personal preferences and job restrictions translate into office occupancy patterns. Some occupants can prefer to arrive as early as 6h00 and others as late as 12h00 [6] and occupancy may even extend to weekends and holidays [7]. Consequently, peak occupancy rarely exceeds 50% [8–12]. For example, Mahdavi, et al. [9] monitored 48 offices in different types of buildings: a university building, a large office complex, and a government building. In all office types, it was found that the workstations were unoccupied at least half of the time, and the occupancy differed significantly from one office space to another.

To accommodate this diversity, building operators tend to choose conservative operating schedules: specifically they shorten temperature setback periods. The temperature setback in commercial buildings refers to reducing heating setpoint temperature and/or increasing the cooling setpoint temperature to reduce the space heating and/or cooling loads during unoccupied periods. It is a simple and effective way of reducing heating and cooling demand [13]. Conservatively short setback schedules are intended to keep the majority of the occupants satisfied with a fixed operating schedule (for heating and cooling systems). Although many office occupants work less than 40 h/week [15], operators tend to conservatively operate office buildings as if they are working much longer than that [6,16,14,17]. If the occupancy patterns in each office can be predicted reliably, customized operating schedules can be adapted for individual offices — eliminating the need for these wastefully conservative fixed operating schedules.

One of the major challenges in predicting the recurring occupancy patterns in an office is the fidelity of occupancy detection technologies. The passive-infrared (PIR) motion sensors are the most commonly used type of sensors to detect occupants' presence in offices [18]. As the presence is inferred through the movement of the occupants, it is nearly impossible to distinguish a short intermediate absence (e.g., a coffee break) from mere occupant immobility [19,20]. Even with a fifteen minute delay, the motion sensors (for office lighting systems) can lead to frequent false absence detections [21,22]. Recently, Nagy, et al. [23] demonstrated that the optimal time delay value for a motion sensor network is unique in

* Corresponding author. Carleton University, Department of Civil and Environmental Engineering, 1125 Colonel by Drive, Ottawa, Ontario K1S 5B6, Canada.
E-mail address: Liam_OBrien@carleton.ca (W. O'Brien).

http://dx.doi.org/10.1016/j.buildenv.2015.06.009

each office space — depending on the PIR sensor's position and office layout. In other words, the PIR sensors can reliably identify the first arrival and the last departure of the day; and they can identify long intermediate vacancy periods. Carbon dioxide (CO_2) sensors have been employed to complement the deficiency of PIR sensors during immobility and to acquire an estimate of the number of occupants in the space [24,25]. As the user presence is inferred through the CO_2 concentration, the detections become susceptible to a number of uncertainties such as window or door positions, sensor calibration, outdoor air supply rate, and the proximity of the occupants to the sensor. Occasionally, other detection methodologies such as acoustic sensors, wearable sensors, and cameras have been employed [26–29]. These occupant detection methodologies, despite being promising to those who study occupant behaviour in buildings, could be a violation of privacy for building operations and add considerable cost [27]. Furthermore, somewhat surprisingly, Hailemariam, et al. [24] reported that inclusion of other sensor types beyond motion sensors did not improve overall accuracy, indicating that the error propagating from different sensor types can undermine the accuracy of occupancy detections.

In recognition of the transient nature of the heat transfer within building fabric, real-time occupancy detections are not adequate to control heating and cooling systems. In order to adapt the operation of heating and cooling systems, the likelihood of occupancy needs to be predicted well before the space is occupied. The challenge in predicting the recurring occupancy patterns in offices is the randomness inherent in occupants' schedules. Several researchers have carried out long-term observational studies and developed models that predict occupants' presence in offices [8,10,19,30–35]. Reinhart [30]'s occupancy model suggests that occupants' arrival and departure times are normally distributed, and the intermediate vacancy periods are represented as constants throughout the year. Wang, et al. [31]'s iteration to this modelling approach suggests that the vacancy periods can be approximated with an exponential decay function. These models assume that the number of intermediate vacancy periods remain constant between individuals and throughout the year. In reality, Rubinstein, et al. [36] reported that the average number of intermediate vacancy periods can be as small as three or as large as nine for different individuals and on different days. Page, et al. [19] elaborated the earlier occupancy models by incorporating the inherent randomness in occupants' mobility — i.e., frequency of intermediate vacancy periods. After analyzing the occupancy patterns on different weekdays, Yang and Becerik-Gerber [8] did not observe a notable difference in occupancy on different weekdays. On the contrary, D'Oca and Hong [37] — after clustering occupants in four distinct occupancy types — identified that for some occupants the day of the week can be an influential factor over the occupancy patterns. In some cases, the occupancy models derived from a specific building's dataset were used to employ temperature setback in vacant or intermittently used thermal zones [8,38–40]. These occupancy-driven temperature setback strategies have proven to yield substantial reductions in the space heating and cooling loads [10].

Because the occupancy models have been derived from a building's dataset retrieved at a certain point during operation, they may become unrepresentative if the occupancy characteristics change during the operation. In fact, the evidence suggests that occupancy is a non-stationary process meaning that its characteristics tend to change slowly in time [41] and space [11]. The development of a different offline occupancy model for each thermal zone in a large office building is a cumbersome and intrinsically labour-intensive process.

Therefore, this paper puts forward a self-adaptive occupancy-learning control algorithm. It recursively identifies the patterns of occupancy in an office space based upon motion sensor detections; and it employs individual and dynamically evolving nighttime temperature setback schedules in each thermal zone. The algorithm was developed after analyzing a year's worth of motion sensor data gathered in seven private offices in an academic building. Later, it was implemented in four local controllers that actuate eight radiant ceiling panel heaters and two variable air volume (VAV) terminal units — serving eight private offices. The algorithm's ability in learning recurring occupancy patterns and maintaining the daytime temperature setpoints during occupied hours were discussed. In order to analyze the energy-savings potential of the control algorithm subject to the environmental conditions specified in Canadian Weather Year for Energy Calculations (CWEC) [42], it was implemented in the Energy Management System (EMS) application of the BPS tool EnergyPlus to determine the temperature setback schedules of a BPS model representing one of the monitored offices.

2. Analysis of the observational dataset

An observational dataset was analyzed to better understand the characteristics of the occupancy patterns detected by motion sensors. This analysis was intended to help us shape the recursive occupancy-learning algorithm.

The observational data includes a year's worth of motion sensor measurements from seven private offices in an academic building. In all seven of the offices, the office furniture did not obstruct the motion sensors. The offices belong to a full-time faculty member. The occupants were typically engaged in sedentary office activities — typing, reading, and writing. The occupants had flexible work schedules interrupted with irregular absences as they were also engaged in work outside their offices (e.g., meetings, conferences, teaching, and work from home) [43].

Motion sensors can detect movements. A movement detection guarantees occupancy; but the absence of a movement detection does not necessarily indicate absence. However, as time elapses in absence of a movement detection, the likelihood of presence decreases. The time between two consecutive movement detections were grouped in one minute bins, they form the empirical likelihood distribution shown in Fig. 1. Simply stating that if a movement detection is not followed by another detection for more than ten minutes, it is very unlikely to observe a new movement — meaning that the office is most likely vacated. Thus, in this data analysis, the occupants' presence was inferred through their movements within a ten minute time-frame. In other words, if there is at least one movement detection within the ten minute time-horizon, the room is assumed occupied; otherwise it is assumed unoccupied. This is consistent with the time delay values in the literature, such as two minutes [19] and ten minutes [20].

Fig. 2 presents the individual mean weekday occupancy profiles for the seven private offices over all weekdays during the

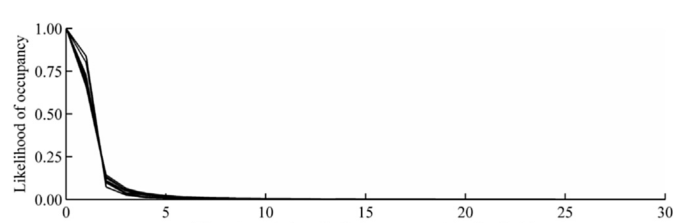

Fig. 1. Empirical likelihood distribution of the frequency of the movement detections. Each line represents one of the seven rooms.