Occupancy-based predictive control of an outdoor air intake damper: A case study

Brodie W. Hobson[1], H. Burak Gunay[1], Araz Ashouri[1,2], Guy R. Newsham[2]
[1]Department of Civil and Environmental Engineering, Carleton University, Ottawa, Canada
[2]Construction Research Centre, National Research Council Canada, Ottawa, Canada

Abstract

Most commercial HVAC systems provide ventilation assuming near-full occupancy during working hours. As conditioning outdoor air requires significant energy, reducing the ventilation requirements based on actual occupancy can improve energy efficiency without significantly impacting occupant comfort. An occupancy-based predictive control program was implemented in a case study building's BAS to control the outdoor air dampers. Changepoint models show the program saved 38.4% and 10.3% of heating and cooling energy during a 24-week implementation period, respectively, with a negligible impact on indoor CO_2 concentrations or the buildings ability to meet temperature setpoints. This study highlights how occupancy data can improve building operational efficiency without the need for additional infrastructure.

Introduction

As many modern workplaces begin to transition towards flexible work hours (Zeytinoglu *et al.*, 2019), commercial and institutional buildings are hosting a fraction of their maximum occupancy on a typical workday (e.g., Barbour *et al.*, 2019). Despite this, most heating, ventilation, and air conditioning (HVAC) systems still operate on static pre-set schedules that assume the building is almost fully occupied during working hours (e.g., ASHRAE 90.1-2019), resulting in chronic overventilation. Considering that indoor climate control systems of these buildings account for up to 40% of the buildings' energy use in Canada (NRCan, 2012), there is a growing oppourtunity to reduce energy use and greenhouse gas emissions by providing HVAC services at an appropriate level for the actual occupancy of the building. This is traditionally attempted with demand control ventilation (DCV), which depends on CO_2 sensors to control a variable air volume (VAV) system by increasing the airflow rate to zones with higher CO_2 levels and therefore higher occupancy. However, the accuracy of commercial-grade CO_2 sensors can often be poor (Shrestha and Maxwell, 2009). A study by Fisk *et al*. (2008) found commercial CO_2 sensors frequently had measurement errors in excess of 20%. Additionally, the lag time between occupant arrival and the build-up of CO_2 has been identified as a weakness to this approach, especially in large spaces with highly variant occupancy profiles (Lu *et al.*, 2018). If improvements are to be made to the operation of system level equipment, such as air handling units (AHUs), the actual occupancy of the space must be predicted beforehand to ensure adequate indoor air quality (IAQ) is maintained. Occupancy-based control of the outdoor air damper in AHUs has been shown in simulation to realize significant energy savings (Cui *et al.*, 2017). Reducing the outdoor air fraction in climates with lengthy heating and/or cooling seasons could reduce the temperature differential between the mixed air and the desired supply air setpoint, lowering the heating and cooling loads.

In addition to occupancy forecasting and occupant-count estimation, an emerging research challenge is to implement occupancy-centric control algorithms in commercially available building automation systems (BASs) at the system level (Park *et al.*, 2019) due to the implementation challenges associated with the limited computational power of these systems. Additionally, while data sources like Wi-Fi have been identified as having strong correlations with occupant count (e.g., Zou *et al.*, 2017; Çiftler *et al.*, 2018; Wang *et al.*, 2018, 2019), the integration of these data streams into BASs is not prevalent due to technological limitations, and privacy and security concerns. However, a ubiquitous data stream within BASs is electrical loads (ISO, 2005) which – while having a weaker correlation with occupant count than Wi-Fi (Hobson *et al.*, 2019) – follow high-level trends with occupancy levels throughout the day (Rafsanjani *et al.*, 2018). Therefore, combining these two disparate data sources may allow for occupancy forecasting to be implemented within a commercial BAS without additional sensing or control infrastructure.

This paper explores the implementation of an occupancy forecasting technique to control the outdoor air damper within a case study building in Ottawa, Canada. Occupancy count estimates from Wi-Fi enabled device count data were used to identify daily occupancy patterns (Hobson *et al.*, 2019). These patterns were compared to patterns in building-level plug-in equipment and lighting loads to determine representative electrical load profiles for high, medium, low, and very low occupancy days. A proposed rules-based occupancy forecasting approach based on these representative profiles (Hobson *et al.*, 2020) was implemented as a program in the BAS. Forecast accuracy, energy savings, IAQ impacts, challenges and relevant anecdotes from the implementation are examined and discussed.

Methods

Representative occupancy profiles from Wi-Fi data (isolated from the BAS) were used to calculate damper profiles. Representative electrical load profiles (available in the BAS) were used as a proxy for the occupancy level, and a classification tree was used to predict day-ahead occupancy levels. To overcome the limited computational power of the building controllers, these models were trained offline using the programming language *R* (R Core Team, 2017). The resulting values and Boolean logic from this process were translated into the building's BACnet controls programming language, GCL+ (Delta Controls Inc., 2006), and implemented in the case study building. It is worth noting that this program only effects the damper position when the dampers are in their minimum position mode. If another operating or safety mode is enabled, the program rightfully loses authority over the outdoor air dampers.

Damper profiles

The case study building is a 6650 m^2 academic office building located in Ottawa, Canada (i.e., ASHRAE climate zone 6A). It is served by two AHUs (AHU1 and AHU2) with a combined outdoor airflow rate of 10,000 L/s at minimum damper positions of 40% and 30%, respectively. The outdoor air dampers are both opposed blade dampers. Assuming the common rule-of-thumb of 10 L/s·person for combined outdoor airflow rates from ASHRAE 62.1 (ASHRAE, 2019b), the building is providing enough outdoor air for 1000 occupants during occupied hours.

Previous work (Hobson *et al.*, 2019) explored different data streams for occupancy-count estimation to determine the actual occupancy of the case study building. Wi-Fi enabled device counts were collected over a seven-month study and the trends in occupancy were analysed. Using k-means clustering, it was found that four distinct daily occupancy profiles were commonly observed in the case study building, denoted 'h', 'm', 'l', and 'L' for high, medium, low, and very low peak daily occupancy, respectively (Hobson *et al.*, 2020), see Figure 1. These profiles represent the maximum occupancy expected at each timestep for a day that falls within each respective cluster. The occupancy profiles used in this study were condensed to this handful of representative profiles to reduce the program size and laboriousness of implementation, as each corresponding damper position profile must be hardcoded into a matrix within the BAS. It was found that the overall peak occupancy rate of the building rarely exceeded 50%, while the average peak occupancy rate was between 17% and 25%. Recall the study by Barbour *et al.* (2019), which found a similar peak occupancy rate of 20% in their case study buildings.

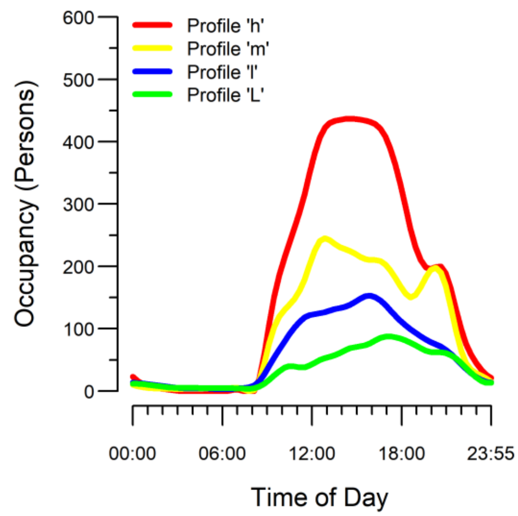

Figure 1: Representative occupancy profiles previously developed for the building.

Area (R_a) and person (R_p) outdoor air rate requirements set forth in ASHRAE 62-1 (ASHRAE, 2019b) were used to calculate minimum breathing zone air rates (V_{bz}) and thus the outdoor air fraction for each of the occupancy profiles. As-built drawings of the case study building were used to determine the area fraction of the outdoor air requirements which are summarized in Table 1, which was approximately 3000 L/s of the 10,000 L/s provided at the current minimum outdoor air dampers' positions. The remaining 7000 L/s accounts for the person outdoor airflow rate requirements for 1000 persons, which can be increased or reduced depending on the occupancy of the building (e.g., if the peak occupancy of the building is actually 500 persons, the breathing zone air rate can be reduced from 10,000 L/s to the 3000 L/s area outdoor air rate plus half of the 7000 L/s person outdoor air rate, or 6500 L/s total).

Table 1: Area outdoor air rate of the case study building.

Occupancy Category	**Area (m^2)**	**Area Outdoor Air Rate, R_a (L/s)**
Lecture classroom	472	142
University/college laboratory	700	630
Computer lab	670	744
Cafeteria/fast-food dining	74	67
Corridor	1437	431
Lobby/prefunction	150	45
Office space	1486	446
Storage/utility	1092	327
Total	**6650**	**2831**

When in economizer mode (i.e., when the outdoor air fraction is at or near 100%), a one-to-one linear relationship between the outdoor air damper position and outdoor air fraction has always been used in the case study building.

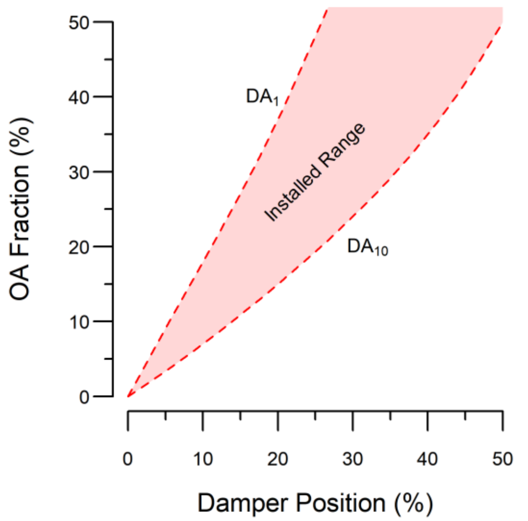

Figure 2: Damper position and air fraction for opposed blade dampers.

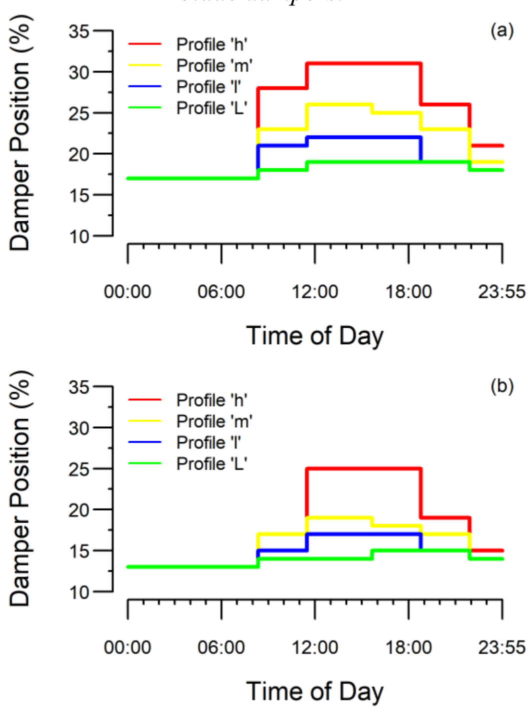

Figure 3: Damper position profiles for (a) AHU1 and (b) AHU2

However, the relationship between the outdoor air fraction and the outdoor air damper position is not necessarily linear, especially at the lower airflow rates that are the subject of this study. The ASHRAE Handbook of Fundamentals (ASHRAE, 2017) specifies a family of curves for installed flow characteristics of different dampers with different control authority based on experimental data from the Air Movement and Controls Association (AMCA, 2006). Felker and Felker (2009) note the actual installed characteristics of dampers typically fall in between the 1% and 10% damper authority (DA$_1$ and DA$_{10}$) curves, see Figure 2. The DA$_{10}$ curve was used in this study as it is a conservative estimation of the amount of airflow that will be achieved at each damper position (i.e., the outdoor air fraction should always be higher than the damper position when using this curve in lower airflow regimes). The outdoor air fraction for each occupancy profile was then translated into its respective damper position profile using this relationship, see Figure 3. The new minimum damper position is between 17% and 31% from zero to maximum observed occupancy for AHU1 and 13% and 25% for AHU2. Recall the previous minimum damper positions were 40% and 30%, respectively, regardless of occupancy levels. To reduce wear on the actuators and network traffic on the BAS, the profiles were discretized into multi-hour increments, and the maximum damper position for each increment was used for the entire increment period as shown in Figure 3. Higher occupancy increments were extended by an hour before and after as a conservative form of pre-conditioning.

Occupancy forecast

A forecast was developed to determine when each damper position profile should be used. Plug-in equipment and lighting load data available within the case study building's BAS were examined at the same time as the previous occupancy study (Hobson *et al.*, 2020). Mean daily electrical load profiles were developed from the days clustered under each occupancy profile, see Figure 4.

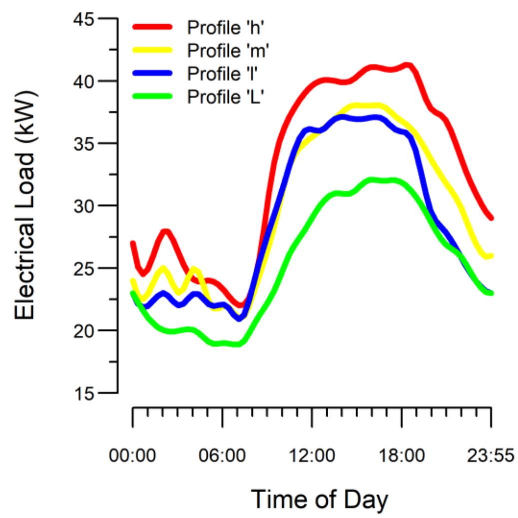

Figure 4: Representative plug-in equipment and lighting load profiles previously developed for the building.

Each day of the study was assigned to an appropriate letter based on which daily profile had the lowest sum of square errors (SSE) compared to the BAS metered load. Motif and discord identification were used on the string of letters to determine commonly repeated weekly patterns in these electrical load data as a proxy to occupancy. The letters for these weekly patterns were analysed to develop a classification tree following the same methodology of the previous study, see Figure 5. Using this forecasting method

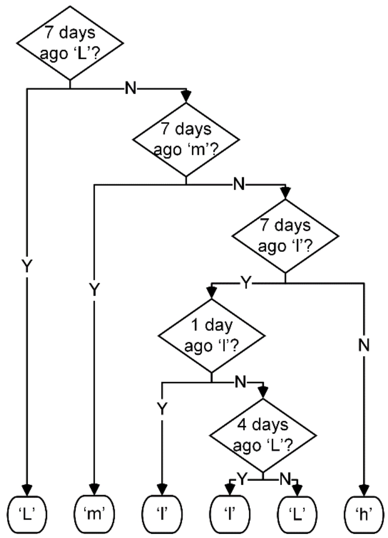

Figure 5: Classification tree for day-ahead occupancy prediction using previous week of electrical load profiles.

on the previously recorded data showed that days were successfully classified 70.4% of the time with a mean occupant-count estimation error of 47 ± 69 persons at 95% confidence. It should be noted that Notably, the forecast was heavily biased towards over-classification of occupancy due to conservative assumptions made in the development of the occupancy and electrical load profiles.

Implementation

The representative electrical load profiles from Figure 4 were hardcoded into a matrix in the GCL+ program. Each day, the daily electrical load profile was recorded at hourly intervals. The sum of the square errors (SSE) between the recorded daily electrical load profile and each of the representative profiles from Figure 4 were computed. Due to limited program size (i.e., 10,240 bytes) the number of timesteps that could be used to calculate the SSE was limited. To determine the number of timesteps and which timesteps to be used, the previously recorded data were analysed using a forward stepwise approach. This revealed that the electrical loads at 1 pm, 5 pm, and 7 pm contributed significantly to forecast accuracy, with a successful classification rate of 67.7% using only the three timesteps compared to the 70.4% achievable when all 24 hourly timesteps were used, see Figure 6. This reduced the dimensions of the electrical load profile matrix to four by three (i.e., four profiles with three values) while simultaneously decreasing the polling rate for recording the electrical loads to only three times per day.

The profile with the lowest SEE (i.e., the closest match to the daily profile) was taken as the representative profile for that day. This process was repeated each day, with the profile membership data being recorded each day of the week in a seven-character long vector. A series of 'IF' statements developed from the classification tree then examine the seven-character long vector once at 10 pm (i.e.,

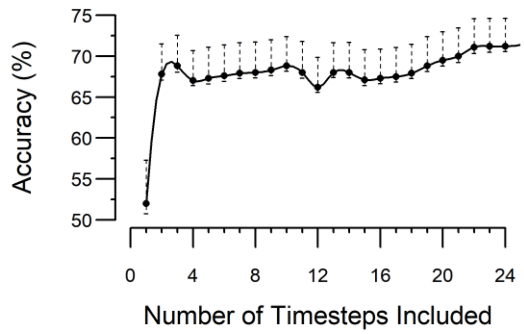

Figure 6: Forward stepwise selection of timesteps for GCL+ forecast.

network traffic is low) to predict the occupancy level for the next day (i.e., high, medium, low, or very low occupancy). The damper positions that correspond to this predicted occupancy are then used throughout the day to control the outdoor air damper. These damper positions are stored in a four by six matrix (i.e., the four profiles and damper positions for six discrete time increments illustrated in Figure 3). One limitation of this method is that the forecast cannot make predictions for the first week of implementation as the program has not yet been able to record the electrical load data off which the forecast is based. To overcome this, when the program is initialized, the seven-character long profile membership vector is initially filled with dummy data that corresponds to five high 'h' occupancy weekdays and two low 'l' occupancy weekend days (i.e., a conservative weekly profile). Therefore, the first week after initialization is based on this dummy data (i.e., until it is overwritten with recorded data).

Results and discussion

Data for the case study building were collected from the third week of August until the beginning of February (i.e., 24-weeks) from the previous year (baseline year) and during the same period for the implementation year for comparison. These data include chiller electrical use data, AHU fan power data, as well as building level plug-in equipment and lighting loads and bulk metered electrical load data. Note that the outdoor air damper program was turned off in error from November 17th to December 20th during the implementation year. Operational data from this period are neglected from both the baseline and implementation year in the analysis. The building is served by a district heating system; steam flow and temperature data were available for the building's heat exchanger. CO_2 sensors in 125 zones and temperature sensors in 138 zones, as well as AHU operational data, were also available. Wi-Fi enabled device counts were also collected for this study via IT services; no MAC addresses were monitored or stored over the course of this study as per the institution's research ethics board.

Accounting for the changes in climate between the baseline and implementation year is critical to ensure the appropriate levels of savings are attributed to the new outdoor air control program. Note that the implementation year had comparatively milder weather than the baseline year; the baseline monitoring period had 2454 heating degree days (HDD) and 157 cooling degree days (CDD), whereas the implementation year had 2131 HDD and 96 CDD during the 24-week study period. Switchover from heating to cooling occurred October 2nd and October 9th during the baseline and implementation year, respectively.

Comparing the outdoor air damper positions from the baseline year and implementation year reflects the lower minimum damper positions, see Figure 7. While it may be expected that the building's AHUs spent more time in the minimum position mode during the colder baseline year, they actually spent 12.0% and 16.6% of hours collectively in minimum position mode during the baseline and implementation period, respectively. This was due to a soft fault in the economizer program of the AHUs that was discovered and altered during the implementation process. This fault was causing the units to go into the economizer mode when the return air temperature exceeded the outdoor air temperature, regardless if it was heating or cooling season. As such, the unit was often fully economizing at low outdoor air temperatures; this caused the supply air temperature to remain at the lower end of the temperature threshold (13°C), resulting in increased loads on the building's perimeter heaters to reach zone setpoints. On a particularly cold day during the implementation period when the program was erroneously turned off, the extreme volume of cold outdoor air caused a heating coil in one of the units to freeze. This fault was identified by facilities personnel who were able to manually reset the unit without incident, at which point the outdoor air damper program was reimplemented.

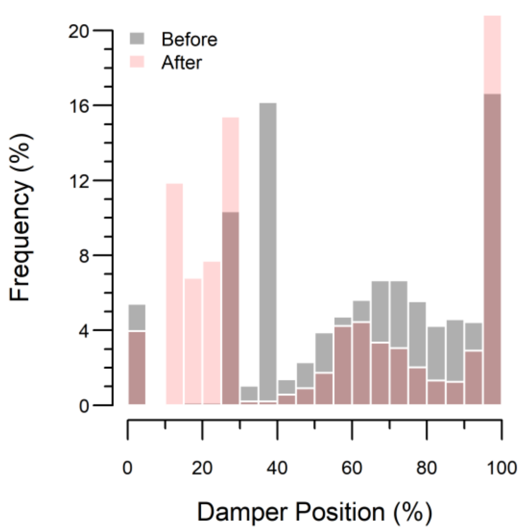

Figure 7: Outdoor air damper positions from the baseline (before) and implementation (after) period.

Forecast accuracy

After the implementation period, Wi-Fi enabled device count data were retroactively examined to determine a baseline for the occupancy forecast accuracy. The analysis showed that a handful of days exceeded the highest expected occupancy profile, with 12 total occupied hours spent above the high 'h' occupancy threshold over the entire 24-week study period. It was found 50, 22, 15, and 80 days during the study period could be classified as high 'h', medium 'm', low 'l', and very low 'L' occupancy, respectively, see Figure 8a. This was then compared to the electrical load forecast-based results from the implementation period. The program classified 101, 3, 13, and 50 days during the study period as high 'h', medium 'm', low 'l', and very low 'L' occupancy, respectively, see Figure 8b. The overall prediction accuracy for the forecast implemented was 58%. However, it should be noted that the forecasts were intended to be conservative; 97% of forecasted days had an observed occupancy that was at or below forecasted levels, see Table 2.

Overall, the accuracy of the electrical load-based occupancy level forecast was lower than expected. This is in part due to the fact that the algorithm had to be initialized at the start

Table 2: Confusion matrix for occupancy forecast.

Electrical forecast \ Wi-Fi measured	h	m	l	L
h	48	20	11	22
m	1	1	1	0
l	0	0	1	12
L	1	1	2	46

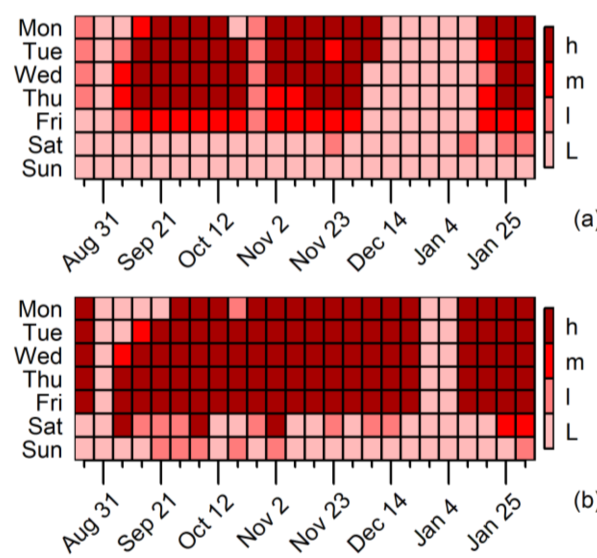

Figure 8: Occupancy profiles (a) measured retroactively with Wi-Fi and (b) forecasted during the implementation period with electrical load.

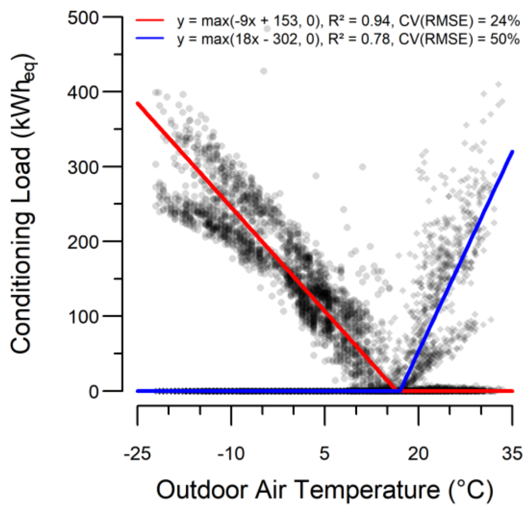

Figure 9: Changepoint model for baseline year.

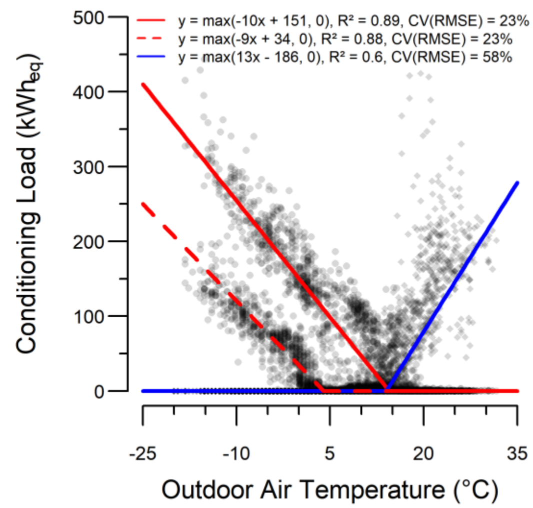

Figure 10: Changepoint model for implementation year.

of the study period and reinitialized during the week of December 20th as the outdoor air damper program was erroneously turned off; this led to the over-classification of ten days of low 'l' and very low 'L' occupancy days to high 'h' during the initialization process, reducing the apparent forecast accuracy. However, despite the highly conservative nature of the program, significant energy savings were still achieved, which may largely be attributed to the outright reduction in outdoor air fraction associated with even the highest new outdoor air damper profile.

Energy savings

Energy use data for the implementation period were compared to energy use data from the previous year during the same period. Three-parameter univariate changepoint models (Gunay *et al.*, 2019) were used to account for the temperature differences between the two years. Note that the heating energy use from the steam loop data include heating coils and perimeter heaters. Fitting the implementation year temperature data to the changepoint models for the baseline year (Figure 9) indicates that the building would have used an estimated 85 MWh_{eq} for cooling and 297 MWh_{eq} for heating if there was no implementation. Actual cooling and heating energy use during the implementation year were 76 MWh_{eq} and 183 MWh_{eq}, representing 10.3% and 38.4% savings, respectively. This change is reflected in the lower slopes of the changepoint models shown in Figure 10. Two distinct groupings of points emerge; the upper family of heating days is characteristic of the energy use on days following the high 'h' occupancy profile, whereas the lower family represents the heating days where low 'l' and very low 'L' occupancy was predicted. Note that the fan energy profiles show a negligible decrease in fan energy consumption of 2.8% (from 71 MWh to 70 MWh using the baseline model and the implementation year itself, respectively).

Comfort impact

In addition to changes in energy use, changes in indoor air quality were explored. CO_2 concentrations from 125 zones were compared between the baseline year and implementation year, see Figure 11. A mean increase of 33 ppm was observed for occupied hours during the implementation year. The increase in hours where indoor CO_2 concentration exceeds atmospheric CO_2 concentrations by 700 ppm – the threshold set out in ASHRAE 62-1 (2019b) – was from 0.06% to 0.23% of occupied hours from before and after implementation, respectively. Further work needs to be done to determine how much of this increase is attributable to the outdoor air program versus drift in the CO_2 sensor accuracy over time.

Changes in indoor air temperature for 138 zones were also examined, see Figure 12. While the mean zone air temperature did not change significantly before and after implementation, the time spent within ± 1°C of the 22°C typical setpoint decreased from 73.8% to 70.0% (3.8% reduction) after implementation. However, this change in ability to meet setpoints may be partially attributable to the milder climate conditions during the implementation year, which led to a prolonged shoulder season (recall that switchover from heating to cooling occurred October 2nd and October 9th during the baseline and implementation year, respectively) and several seasonally warm days during the heating season. To explore this phenomenon, temperature data from 32 zones in another academic office building located on the same campus as the case study building were examined. It was similarly found that the mean zone air temperature did not change significantly, however the number of occupied hours spent within ± 1°C of the setpoint decreased from 65.9% to 63.5% (2.4% reduction) during the same study period.

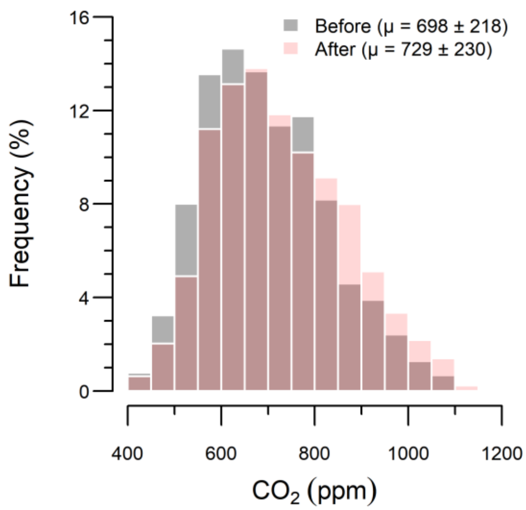

Figure 11: Changes in CO_2 concentration.

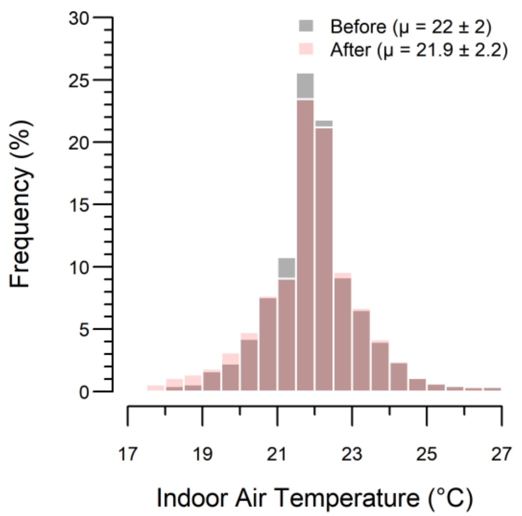

Figure 12: Changes in indoor air temperature.

Unresolved issues

Reducing the outdoor air fraction and increasing the return air fraction increased the mixed air temperature during the study period, reducing the heating coil output to meet temperatures within the supply air temperature setpoint range (13°C to 20°C). As a result, supply air temperatures during the implementation year were higher than the baseline year. The baseline year had a mean temperature of 16.1 ± 4.1°C with 39.1% of hours spent in the upper half of the setpoint range (>16˚C), compared to 16.5 ± 4.7°C with 50.9% of hours spent above 16˚C during the implementation year. This significantly reduced the loads on the perimeter heaters and resulted in energy savings. A disadvantage of the implemented program was that higher return air (and hence higher supply air) temperatures would occur as the building's average room temperature increased, reducing the availability of cooling in problem zones and in some instances causing localized overheating, especially during very low 'L' occupancy days where the outdoor air temperature rose above 0°C. This problem could potentially be addressed by increasing the airflow rate at the zone level in rooms where this issue was identified, or by introducing an override during days that match these criteria to ensure adequate outdoor airflow for cooling purposes.

While overly-excessive ventilation rates have little benefit to occupant comfort or satisfaction, well-ventilated buildings (e.g., above 10 L/s·person) have been associated with better human factor outcomes (Charles and Veitch, 2002). If ventilation can be reduced to levels that still provide IAQ benefits, significant savings may still be possible. For example, the minimum design ventilation rate in this case study was 13 L/s·person (6500 L/s for an assumed 500 maximum occupants), which was sufficiently low to generate HVAC savings. Marginally increasing this ventilation rate may mitigate changes to the building's IAQ with a small cost to energy savings. The optimal balance between these two factors warrants further investigation.

Additionally, while the energy benefit of reducing the outdoor air fraction outright based on occupancy is apparent, the savings directly attributable to the forecasting method used should be further explored. As the variability of occupancy in a space increases, so does the energy savings potential (Ouf, O'Brien and Gunay, 2019). While the case study building had highly variable occupancy levels, the occupancy patterns of an academic office building are relatively predictable and may not require such complex forecasting mechanisms where a schedule may be reliably used. Therefore, while the case study building facilitated a useful methodological exercise in creating these rules-based forecasts, the benefits of such an approach should be thoroughly explored for other occupancy types and patterns.

Conclusion

This study found that the case study building was providing enough ventilation for 1000 occupants constantly during occupied hours, whereas the peak one-time occupancy during the entire study period was less than 600 occupants. A GCL+ script was deployed on the BACnet controllers within the BAS to create a simple program for predicting occupancy, based on clustering and motif identification of Wi-Fi and electrical load data. This program was conservative, classifying 97% of days at or above their actual occupancy. However, the program was still able to reduce heating and cooling energy use by 38.4% and 10.3%, respectively, with a 33 ppm increase in CO_2 concentrations and a 3.8% reduction in hours spent within ± 1°C of temperature setpoints (likely attributable to the milder climate during the implementation year). This case study highlights an example of the benefits to be gained by including occupancy information in building operations.

Acknowledgement

This research is supported by research funding provided by Natural Resources Canada, CopperTree Analytics, Sensible Building Science, Delta Controls, Bentall GreenOak, and National Research Council Canada. The authors acknowledge Gavin Symonds, Scott Macdonald, and their team at Carleton University's Facilities Management and Planning, as well the IEA EBC Annex 79 researchers, as we have greatly benefitted from the associated discussions.

References

AMCA (2006) *AMCA Publication 502-06 (R2009) - Damper Application Manual for Heating, Ventilating, and Air Conditioning*. Arlington Heights, IL, IL: AMCA.

ASHRAE (2017) *2017 ASHRAE Handbook: Fundamentals (SI Edition)*. Atlanta, GA: ASHRAE.

ASHRAE (2019a) *ANSI/ASHRAE/IES 90.1-2019: Energy Standard for Buildings Except Low-Rise Residential Buildings*. Atlanta, GA.

ASHRAE (2019b) *ANSI/ASHRAE Standard 62.1-2019: Ventilation for Acceptable Indoor Air Quality*.

Barbour, E. *et al.* (2019) 'Planning for sustainable cities by estimating building occupancy with mobile phones', *Nature Communications*. Springer US, 10(1).

Charles, K. E. and Veitch, J. A. (2002) *Outdoor Ventilation Rates in Offices and Occupant Satisfaction (IRC-RR-160)*. Ottawa, Canada.

Çiftler, B. S. *et al.* (2018) 'Occupancy Counting With Burst and Intermittent Signals in Smart Buildings', *IEEE Internet of Things Journal*, 5(2), pp. 724–735.

Cui, X. *et al.* (2017) 'Modelling and Performance Evaluation of an Air Handling Unit for an Air Treatment System with Regulated Outdoor-air Fraction', *Energy Procedia*. The Author(s), 105, pp. 4718–4723.

Delta Controls Inc. (2006) *ORCAview Version 3.33 Technical Reference Manual*. 1.6. Vancouver, B.C.

Felker, L. G. and Felker, T. L. (2009) *Dampers and Airflow Control*. Atlanta, GA: ASHRAE.

Fisk, W. J., Faulkner, D. and Sullivan, D. P. (2008) 'A pilot study of the accuracy of CO2 sensors in commercial buildings', *IAQ Conference*, (October), pp. 1–11.

Gunay, H. B. *et al.* (2019) 'Detection and interpretation of anomalies in building energy use through inverse modeling', *Science and Technology for the Built Environment*. Taylor & Francis, 25(4), pp. 488–503.

Hobson, B. W. *et al.* (2019) 'Opportunistic occupancy-count estimation using sensor fusion: A case study', *Building and Environment*. Elsevier, 159(January), p. 106154.

Hobson, B. W. *et al.* (2020) 'Wi-Fi Based Occupancy Clustering and Motif Identification : A Case Study', *ASHRAE Transactions*, 126(1), pp. 256–264.

ISO (2005) *ISO 16484-3:2005 Building automation and control systems (BACS) — Part 3: Functions*. Geneva.

Lu, S., Karaguzel, T. O. and Hameen, E. C. (2018) 'Simulations of dynamic ventilation operations with real-time occupancy estimation', in *ACEEE Summer Study Proceedings*. Pacific Grove, CA, pp. 1–10.

Natural Resources Canada (2019) *Survey of Commercial and Institutional Energy Use - Buildings 2014 - Detailed Statistical Report 2019*.

Ouf, M. M., O'Brien, W. and Gunay, H. B. (2019) 'On quantifying building performance adaptability to variable occupancy', *Building and Environment*. Elsevier, 155(February), pp. 257–267.

Park, J. Y. *et al.* (2019) 'A critical review of field implementations of occupant-centric building controls', *Building and Environment*. Elsevier, 165(May), p. 106351.

R Core Team (2017) 'R: A language and environment for statistical computing'. Vienna, Austria: R Foundation for Statistical Computing.

Rafsanjani, H. N., Ahn, C. R. and Eskridge, K. M. (2018) 'Understanding the recurring patterns of occupants' energy-use behaviors at entry and departure events in office buildings', *Building and Environment*. Elsevier, 136(March), pp. 77–87.

Shrestha, S. S. and Maxwell, G. M. (2009) 'An experimental evaluation of HVAC-grade carbon dioxide sensors -Part I: Test and evaluation procedure', *ASHRAE Transactions*, 115 PART 2, pp. 471–483.

Wang, W. *et al.* (2018) 'Occupancy prediction through Markov based feedback recurrent neural network (M-FRNN) algorithm with WiFi probe technology', *Building and Environment*. Elsevier, 138(April), pp. 160–170.

Wang, W. *et al.* (2019) 'Linking energy-cyber-physical systems with occupancy prediction and interpretation through WiFi probe-based ensemble classification', *Applied Energy*. Elsevier, 236(August 2018), pp. 55–69.

Zeytinoglu, I. U., Cooke, G. B. and Mann, S. L. (2019) 'Flexibility : Whose Choice Is It Anyway ?', *Industrial Relations*, 64(4), pp. 555–574.

Zou, H. *et al.* (2017) 'Non-intrusive occupancy sensing in commercial buildings', *Energy and Buildings*. Elsevier B.V., 154, pp. 633–643.