Implementation of next-generation occupant-centric sequences of operation in an office building using supervisory control

Hobson, Brodie W.; Markus, Andre A.; Bursill, Jayson; Burak Gunay, H.; Darwazeh, Darwish; O'Neill, Zheng

Publisher's version / Version de l'éditeur:

https://doi.org/10.1016/j.enbuild.2024.115087

Energy and Buildings, 327, C, 2024-11-24

Energy & Buildings 327 (2025) 115087

Contents lists available at ScienceDirect

Energy & Buildings

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

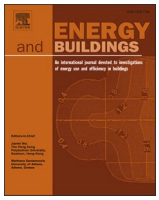

Implementation of next-generation occupant-centric sequences of operation in an office building using supervisory control

Brodie W. Hobson [a,*], Andre A. Markus [a], Jayson Bursill [a,b], H. Burak Gunay [a], Darwish Darwazeh [c], Zheng O'Neill [d]

[a] Department of Civil and Environmental Engineering, Carleton University, 1125 Colonel By Drive, Ottawa, Ontario K1S 5B6, Canada
[b] Delta Controls Inc., 17850 56 Ave, Surrey, British Columbia V3S 1C7, Canada
[c] Construction Research Centre, National Research Council Canada, 1200 Montreal Road, Ottawa, Ontario K1K 2E1, Canada
[d] J. Mike Walker '66 Department of Mechanical Engineering, Texas A&M University, 202 Spence Street, College Station, TX 77840, USA

ARTICLE INFO

Keywords:
Demand-controlled ventilation
DCV
HVAC
VAV
AHU
RP-1747
ASHRAE
Guideline 36
Sequences of operation
Occupant-centric control
Occupancy-centric control
OCC
Supervisory control
Supervisory controller
Implementation
Case study

ABSTRACT

ASHRAE RP-1747 is a CO_2-based demand-controlled ventilation (DCV) approach which uses trim and respond logic to dynamically adjust variable air volume (VAV) terminal units' minimum airflow setpoints based on zones' ventilation requirements. While simulation results and laboratory testing have estimated the impact on heating, ventilation, and air conditioning (HVAC) systems compared to traditional ventilation, there have been limited applications of RP-1747 in occupied, real-world buildings to date. This paper introduces a real-world implementation of RP-1747 DCV in an institutional office building over an eight-month period, using a supervisory control approach. During this implementation, a complementary temperature setback approach was also developed and employed in the case study building. These changes to the sequences of operation, as well as corrections to other sub-optimal sequences of operation discovered during implementation, resulted in a 36 ± 2 % and 2 ± 6 % reduction in heating and cooling energy use, respectively, while improving per person ventilation rates in the case study. The results aim to contribute to the body of literature on this emerging DCV approach and provide anecdotal evidence of its benefits and interactions with other control logic in a real-world application, while also demonstrating the benefits of supervisory control when implementing complex process control functions.

1. Introduction

The field of occupant-centric controls (OCC) in commercial buildings has evolved rapidly through concerted research efforts such as the recently completed International Energy Agency's Energy in Buildings and Communities (IEA-EBC) Annex 79 *Occupant-Centric Building Design and Operation* [1]. This method of control strays from conventional building operations – which are typically based on static and rigid schedules and sequencing of equipment – by employing control strategies that are occupant behaviour- or occupancy-centric. The control strategies which fall under the umbrella of OCC range in scale and complexity, from the detection of occupant presence at the building level (i.e., the lowest resolution of occupant data, or grade 1 data as defined in Gunay et al. [2]) to improve scheduling of plant-level equipment (e.g., [3,4]), to the monitoring of individual occupants and their activity at the zone level (i.e., the highest resolution of occupant data, or grade 6 data as defined in Gunay et al. [2]) to adjust the delivery of heating, ventilation, and air conditioning (HVAC) services and lighting automatically based on occupant location and activity level (e.g, [5–7]), for example.

Park et al. [8] provide the first literature review dedicated explicitly to OCCs that have been implemented in the field. They found that the most common OCC implementations were targeted at zone-level HVAC equipment, such as variable air volume (VAV) terminal units. Similarly, the most commonly employed form of OCC in the commercial building industry is CO_2-based demand-controlled ventilation (DCV). These strategies rely on the differential between zone and supply air CO_2 concentrations to estimate ventilation demand in individual zones and adjust the delivery of outdoor air accordingly. Ventilation rates are

* Corresponding author at: Carleton University, Department of Civil and Environmental Engineering, 1125 Colonel By Drive, Room 3432, C.J. Mackenzie Building, Ottawa, Ontario K1S 5B6, Canada.
E-mail address: brodie.hobson@carleton.ca (B.W. Hobson).

https://doi.org/10.1016/j.enbuild.2024.115087
Received 10 September 2024; Received in revised form 1 November 2024; Accepted 18 November 2024
Available online 22 November 2024

Nomenclature

A_z	zone floor area (m^2)
C_{bz}	breathing zone CO_2 concentration (ppm)
CR_z	zone criticality ratio (%)
C_s	supply air CO_2 concentration at the air handling unit (ppm)
E_z	air distribution effectiveness (%)
I	number of ignored requests (−)
k	CO_2 generation rate (L/s/met/person)
M	importance multiplier
m	metabolic rate (met)
P_z	zone population (persons)
R	number of requests from zone(s) (−)
R_a	area outdoor air rate (L/s)
R_p	people outdoor air rate (L/s)
S	binary occupancy status (1 = occupied, 0 = unoccupied)
SP_0	initial setpoint (L/s)
SP_{max}	maximum setpoint (L/s)
SP_{min}	minimum setpoint (L/s)
SP_{res}	respond amount (L/s)
$SP_{res-max}$	maximum respond amount per time interval (L/s)
SP_t	setpoint at time t (L/s)
SP_{trim}	trim amount (L/s)
S_{rad}	radiator valve position (1 = fully open, 0 = fully closed)
T_{sp}	zone air temperature setpoint (°C)
$T_{sp,afterhours}$	zone air temperature setpoint during afterhours (°C)
$T_{sp,default}$	zone air temperature default setpoint during operating hours (°C)
$T_{sp,occupant}$	zone air temperature setpoint specified by occupant (°C)
V_{bz}	breathing zone outdoor air flow rate (L/s)
V_{oz}	zone outdoor airflow (L/s)
V_{pz}	zone primary air flow to the ventilation zone (L/s)
V_{ps}	system primary air flow (L/s)
Z_p	maximum zone primary outdoor air fraction (%)
Z_{pz}	zone primary outdoor air fraction (%)

increased to zones with higher CO_2 concentrations and presumably denser occupancies, whereas ventilation rates are reduced in zones where CO_2 concentrations are low and occupancy is sparser. This approach improves indoor air quality (IAQ) by diluting indoor contaminants (e.g., CO_2, bio-effluents and other volatile organic compounds (VOCs) [9,10], etc.) and reducing HVAC energy use by only providing ventilation services when and where they are needed, and in the amount that they are needed [11]. This approach differs from the traditional approach to ventilation in ASHRAE Standard 62.1-2022 [12] which requires a static level of ventilation air to be delivered regardless of the actual number of occupants present in a given space, potentially wasting HVAC energy and even causing thermal discomfort [13].

The benefits of DCV are well known in industry and practice: since 1999, ASHRAE Standard 90.1 has mandated the use of CO_2-based DCV in zones with specific occupancy categories, areas, and/or occupant densities. Since its inception, both iterations of ASHRAE Guideline 36 [14,15] have referred to DCV as best-practice and refer to ASHRAE Standard 62.1–2019 [16] for details on its implementation. However, despite its requirement in some cases and recommendation as best-practice in others, the specific operational sequences which constitute DCV have historically been left to individual controls integrators and/or operators to define. This lack of standardization has led to a wide range of CO_2-based DCV approaches being deployed across the commercial building stock, which undoubtedly vary in both complexity and quality. To address this gap, ASHRAE Research Project- (RP-) 1547 *CO_2-based Demand Controlled Ventilation for Multiple Zone HVAC Systems* [17] was undertaken to create a standardized approach to CO_2-based DCV. Though this effort was completed in 2013, the results from the project still lacked "practical, programmable control sequences" [18] that could be used by industry practitioners in commercial buildings. A follow-up project, ASHRAE RP-1747 *Implementation of RP-1547 CO_2-based Demand Controlled Ventilation for Multiple Zone HVAC Systems in Direct Digital Control Systems* [19] was commissioned to provide these practical, programmable sequences for CO_2-based DCV logic. Although this project was completed in 2017, its adoption by ASHRAE Guideline 36 is still pending as its efficacy in real-world applications must be verified [15]. Once this is completed, RP-1747 DCV is set to become the industry-standard DCV strategy.

Therefore, the aim of this paper is to provide novel analyses which speak to the efficacy of RP-1747 DCV in a real-world occupied building to further support its adoption by forthcoming versions of ASHRAE Guideline 36. The RP-1747 DCV control logic was deployed to a densely instrumented case study building over an eight-month implementation period; this represents the first documented real-world implementation of RP-1747 in a cold climate to date. To achieve this, a supervisory control approach was employed which allowed for RP-1747 to be deployed with minimal invasiveness in the case study building. Several deficiencies in the case study buildings operation – namely an inappropriate minimum outdoor air damper (OAD) position and a suboptimal supply air temperature (SAT) reset – were corrected over the course of the implementation. Another OCC, referred to as a 'probabilistic temperature setback', was deployed during the implementation to enhance the efficacy of RP-1747 DCV. The reasoning behind and the impact of these upgraded sequences of operation on the operation and energy use of the case study building are presented and discussed throughout this paper.

The remainder of this paper is structured as follows. In the background section, a brief overview of the established control logic behind ASHRAE RP-1747 DCV is provided, and considerations for the use of a supervisory control approach in its implementation are discussed. Then, the methodology section presents the case study building and data, the supervisory control approach and sequences of operation implemented over the duration of the study period, and the methods used for quantifying changes in building energy use and occupant comfort. The results from this analysis are presented and the limitations are discussed. Lastly, the conclusions and recommendations from this study are developed and presented.

2. Background and previous work

The section introduces the established control logic behind ASRHAE RP-1747 DCV, as well as the use of supervisory control for the purpose of process control in BASs.

2.1. ASHRAE RP-1747 demand-controlled ventilation

ASHRAE RP-1747 [18,19] is a CO_2-based DCV approach which uses trim and respond (T&R) logic [20,21] to dynamically adjust VAV terminal units' minimum primary airflow setpoints (V_{pz-min}) based on breathing zone outdoor airflow (V_{bz}) requirements and zone air distribution effectiveness (E_z). Briefly, this logic uses VAV airflow, CO_2 sensor, and occupancy sensor (e.g., motion detector) data at the zone level and the air handling unit (AHU) outdoor air ratio (OAR) to determine criticality ratios (CR_z), which essentially represents how much outdoor air is entering a zone, relative to ventilation requirements and the maximum amount that the VAV can provide at the current system-level OAR. Further information on the logic of this DCV approach is provided in Section 3.2.1.

To date, studies focused on RP-1747 DCV have been simulation- (e.g., [18,22]) or laboratory-based [18], with a single preliminary study on its implementation in a real-world building [23]. In their study, which utilized the co-simulation of CONTAM and EnergyPlus, O'Neill et al. [18] found between 9 % and 33 % annual HVAC energy savings using RP-1747 DCV compared to traditional ventilation controls in a 200 m^2 institutional office building. These savings varied depending on the climate zone considered in the study (i.e., ASHRAE Climate Zones 1A, 3A, 3B, and 5A), with savings generally increasing in colder climates. Similarly, Hobson et al. [22] found between 4 % and 41 % annual HVAC energy savings in their EnergyPlus simulations of a 600 m^2 institutional office building in ASHRAE Climate Zone 6A; these savings varied depending on the density of the CO_2 and motion sensor grid in the building. The real-world implementation of RP-1747 in the same building during the cooling season [23] found up to 25 % HVAC energy savings during a brief two-week period. However, there is a notable lack of longitudinal studies on the implementation of RP-1747 DCV in real-world buildings to date.

2.2. Supervisory control for process control functions

One obstacle to the adoption of RP-1747 DCV – and other control logic outlined in ASHRAE Guideline 36 – is that building automation systems (BASs) and the sequences of operation contained therein are growing increasingly complex. Consider the fact that ASHRAE Guideline 36 has increased from 102 pages in its initial 2018 release [14] to 292 pages in the latest 2021 version [15]. While the ASHRAE Guideline 36 has shown potential to generate up to 31 % energy savings in HVAC systems [24], the number of lines of code to implement these high-performance sequences of operation in multi-zone VAV systems has increased sevenfold compared to the status quo in the shorter 2018 edition alone [25]. Considering that just over 30 % of control-related problems in commercial and institutional buildings were attributable to programming errors before the introduction of ASHRAE Guideline 36 [26] – and the same amount of excess energy is consumed by commercial buildings as a result [27] – it is not unreasonable to assume that the number of programming errors will continue or increase as a result of lengthier and more complex control sequences. Despite these challenges, the benefits of ASHRAE Guideline 36 if properly implemented are evident. It is therefore important to consider how buildings can be controlled to allow for rapid alteration to either rectify programming errors where they arise or incorporate forthcoming innovative control sequences offered by each new addition of ASHRAE Guideline 36, such as RP-1747 DCV.

One of the potential solutions to this problem is supervisory control. Supervisory control is defined in the ASHRAE Handbook as the upper control level which "specifies the set points and other modes of operation that are time dependent" [28]. Supervisory control of HVAC systems exists in a layer above local controllers and is used to supervise the control loops typically handled by local controllers, such as sequencing and process control [29]. Simply put, the optimal operation of a local proportional-integral-derivative (PID) loop typically used for process control may not be globally optimal when considering its impact on the entire system; supervisory control strategies can monitor collections of local-loop controllers to optimize the operation of plant-, system-, and zone-level HVAC equipment. However, because supervisory control inherently requires the monitoring of the control variables used by local process controllers, there exists the potential to perform more complex process control functions directly on the supervisory controller itself [30]. This is especially true given the "significant increases in computational power and communication capabilities" [28] that supervisory control systems can now take advantage of, which allow them to "incorporate many new data steams" [28]. For example, supervisory control strategies are often employed to enable model predictive control (MPC) (e.g., [31–33]) which requires increased computational power to perform optimizations using a combination of control variables from the building and external data sources, such as weather data. In this case, a parallel supervisory controller with enhanced computational capabilities may be used in tandem with the buildings existing supervisory controls in the event that communications are lost with external data streams. Such an approach offers several benefits. For example, the basic controls needed for the local controller to perform process functions can be left on the existing controllers, ensuring system services continue to operate properly and adequately in the event of lost communication with the enhanced supervisory controller. The enhanced supervisory controller can then shoulder the burden of more complex process functions by taking advantage of additional computational resources outside of the BAS, such as server- or cloud-based computing infrastructure. Additionally, the variables and control logic needed to perform these process functions can be stored on the supervisory controller rather than cluttering the local controller with the additional variables needed to perform the process functions.

This approach creates a paradigm whereby the existing controllers provide the building with its baseline level of operation needed to ensure adequate service delivery and life-safety operations, and the enhanced supervisory controller acts as a layer above the existing controllers to provide enhanced sequences of operation that, while beneficial to optimal operation and energy use, are non-critical to the functionality of the BAS. As such, this approach could also be extended to high performance sequences of operation such as those outlined in ASHRAE Guideline 36 [34]. This lessens the burden on operations personnel as the critical control logic and variables for regular system operation will be readily available on the local controllers, with extraneous control logic and variables handled (with appropriate precautions for control logic intervention) by the enhanced supervisory controller, thus reducing the sevenfold increase in lines of code from ASHRAE Guideline 36–2018 [14] on the BAS itself. While the length of the enhanced process control logic within the supervisory controller itself would be non-negligible, the centralized nature and potential use of server- or cloud-based computing infrastructure for supervisory controls means that the overall length and relative complexity of the code can be comparatively reduced. This will be explored in Section 3.3.

3. Methodology

In this section, the case study building and dataset are introduced. Implemented changes to the sequences of operation, including RP-1747 DCV, probabilistic zone temperature setbacks, and other sub-optimal sequences are presented. Then, the supervisory control framework is outlined. The measurement and verification approach used to quantify changes in energy use and occupant comfort before and after implementation of the high-performance sequences of operation are also described.

3.1. Building and dataset

The case study building is a 28-zone floor of an institutional office building located in Ottawa, Canada (i.e., ASHRAE Climate Zone 6A), as illustrated in Fig. 1. The building, which received its occupancy permit less than a year before the initiation of this study, has a total conditioned area of 626 m^2. The predominately south- and east-facing portions of the façade are adjacent to another building, with only the north- and west-facing perimeter zones facing the exterior. The window-to-wall ratio (WWR) of the exterior facing façades is approximately 42 %. The opaque portion of the exterior wall assembly has an effective R-SI value of 6.21 m^2K/W, while the triple-glazed argon-filled glazing units and their frames have an effective overall U-SI value of 0.9 W/m^2K and average solar heat gain coefficient (SHGC) of 0.27. The roof R-SI value is approximately 6.45 m^2K/W. This represents a high-performance building envelope which exceeded energy code requirements at the time of design [35].

The floor is served by a single dedicated AHU with a total supply

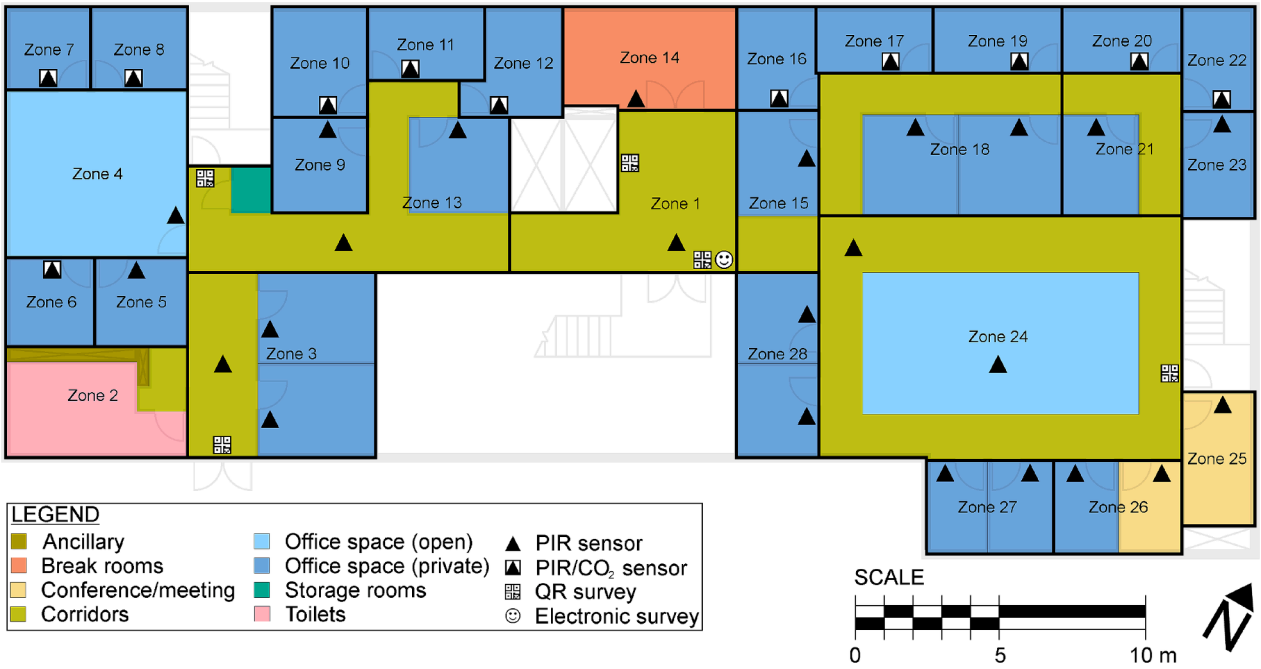

Fig. 1. Case study building floor plan, sensor locations, and VAV zoning.

airflow capacity of 5147 L/s provided by a 15 kW variable frequency drive (VFD) supply fan. The AHU is equipped with a hot water heating coil, a chilled water cooling coil, and a resistive electric steam humidifier with nominal rated capacities of 43 kW, 90 kW (60 kW sensible), and 6 kW, respectively, and does not have any heat recovery capabilities. The AHU utilizes a dry-bulb temperature differential-based economizer control scheme whereby the OAD position is minimized when the building is in the heating mode (OS#1) or mechanically cooling mode (OS#4), see Fig. 2, during which the hot water heating coil and chilled water cooling coil, respectively, actuate to meet the SAT. The SAT is dynamically adjusted between 13 °C and 27 °C depending on the heating or cooling demand in the building. However, the SAT range specified during design was 13 °C to 20 °C. The specifics of the SAT reset scheme utilized and the implications of the elevated SAT will be discussed in Section 3.2.3. The minimum OAD position utilized by the AHU was 20 %, however, it should be noted that the minimum OAD position specified during design was 40 %. The implications of this programming error during commissioning will also be discussed in Section 3.2.3. At intermediate temperatures, the OAD position increases above the minimum OAD position as needed (i.e., free cooling mode (OS#2)), until reaching 100 % outdoor air during economizer and mechanical cooling mode (OS#3). The AHU operates between 6 AM and 10 PM Monday to Friday and from 8 AM to 6 PM on weekends and holidays (referred to as

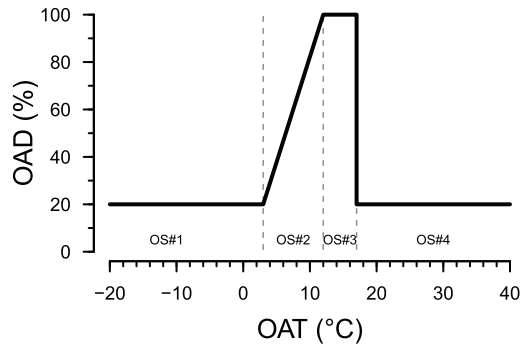

Fig. 2. Outdoor air damper (OAD) position for different air handling unit (AHU) operating states (OSs).

operating hours). Outside of operating hours (referred to as afterhours), the AHU can cycle on if a single zone's temperature drops below 17 °C in the heating season or increases above 27 °C during the cooling season.

The AHU provides supply air to a single-ducted VAV terminal unit without reheat in each of the 28 zones, see Fig. 1. Displacement ventilation is used in all zones except zone 2 (i.e., the toilets), with floor-level wall-mounted diffusers for supply air and ceiling-level wall-mounted grilles for return air. Ventilation is provided to zone 2 via ceiling-mounted diffusers for supply air; the air from zone 2 is exhausted directly to the outdoors. Hydronic baseboard units with BAS-integrated controllers provide additional heating in each zone as needed; heating is generally available from October through April inclusively, at which point seasonal switchover to mechanical cooling occurs. Water for heating and cooling is conditioned by a steam and chilled water loop, respectively, which is fed by a central physical plant that serves multiple buildings. See Table 1 for further information on zone and zone-level equipment sizing.

Each of the 28 zones contains a wall-mounted BAS-integrated thermostat with an integrated passive infrared (PIR) sensor which provides zone air temperature (ZAT) and setpoint and binary motion detection (S) data. The occupants can control their ZAT setpoint via these thermostats within a limited range; the default ZAT setpoint was 22 °C year-round, with occupants able to increase or decrease their setpoint by ±2 °C during operating hours. Afterhours, the ZAT setpoint defaults to 18 °C during the heating season and 26 °C during the cooling season. The zone primary airflow rate (V_{pz}) and setpoint and valve position (S_{rad}) are collected from the single-ducted VAV terminal units and hydronic baseboards, respectively. At the system level, relevant data collected from the AHU includes the supply air CO_2 concentration (C_s), outdoor air temperature (OAT), supply fan energy use, heating and cooling coil load, system primary airflow rate (V_{ps}) and setpoint, OAD position and setpoint, and SAT.

Additional sensing and metering are available in the 11 exterior perimeter private offices (i.e., zones 6–8, 10–12, 16, 17, 19, 20, and 22), see Fig. 1. The thermostats in these offices contain integrated non-dispersive infrared (NDIR) CO_2 sensing to measure zone-level CO_2 concentrations (C_{bz}). A ceiling-mounted sensor hub provides an additional PIR sensor for binary motion detection. The hydronic baseboards in these zones are also equipped with BAS-integrated energy valves

Table 1
Zone and equipment sizing.

Zone	A_z(m^2)	P_z(persons)	V_{pz-min}(L/s)	V_{pz-max}(L/s)	Baseboard heating capacity (W)
1	29	0	21	106	0
2	15	2	62	224	450
3	34	2	31	232	230
4	44	3	59	416	2205
5	12	1	21	63	130
6	12	1	21	170	630
7	13	1	31	212	1260
8	13	1	21	170	755
9	13	1	21	63	130
10	16	1	21	106	755
11	14	1	21	170	932
12	15	1	21	106	466
13	42	1	21	169	130
14	29	8	66	424	1398
15	15	1	21	169	100
16	13	1	21	107	630
17	13	1	21	170	932
18	22	2	21	232	260
19	14	1	21	169	932
20	13	1	21	170	1500
21	33	1	21	169	130
22	12	1	21	107	630
23	11	1	21	63	207
24	103	4	59	318	518
25	14	4	31	212	207
26	17	2	31	212	260
27	17	4	21	94	260
28	30	2	21	126	230
Total	626	50	790	4949	16,267

which measured the baseboard output based on the hot water flow rate and temperature differential for each unit. Data from the AHU and each zone were collected at concurrent 10-minute timesteps for a consecutive 14-month period from February 1st, 2023, up to and including April 30th, 2024.

Prior to the start of this study, the authors undertook a measurement and verification campaign to test the functionality of the sensing and metering infrastructure in the 11 exterior perimeter zones; both the wall- and ceiling-mounted PIR sensors were confirmed to be functional at the time of testing, and the CO_2 response of the NDIR sensors were verified to fall within manufacturer specified tolerances (i.e., ± 30 ppm + 3 % of sensor reading). It was noted that some sensor drift was experienced with the CO_2 sensors since the buildings commissioning, namely a gradual increase in the ambient CO_2 concentrations measured during afterhours. In order to correct this sensor drift, an automatic background calibration (ABC) script was created in each controller to reset the CO_2 sensors ambient reading down to atmospheric CO_2 concentration (i.e., approximately 420 ppm) based on the lowest weekly CO_2 concentration reading from each sensor. The functionality of the hydronic baseboard radiator valves was also verified using a thermal camera; one leaky valve was identified through this process and reported to facilities personnel for repairs. The flow of each VAV terminal unit was also measured with its damper at the fully opened and fully closed position using a micromanometer. The maximum measured flowrate of each VAV terminal unit was verified to match their designed maximum flowrate, however, five of the 11 units had non-negligible volumes of airflow (i.e., greater than 10 % of their rated airflow capacity) when the dampers were fully closed, indicating possible leaky dampers. These leaky dampers were also identified and reported to facilities personnel.

3.2. Implemented sequences of operation

The sequences of operation implemented over the course of this study are summarized below. Table 2 provides a brief overview of the major implementation phases and milestones. Note that limited trials of the RP-1747 control logic were conducted in the summer preceding its implementation to allow for testing and bug fixing to occur.

Table 2
Major phases and milestones for the implementation.

Phase	Milestone	Date
Before	Data collection begins	February 1st, 2023
	Seasonal switchover to cooling	May 12th, 2023
After RP-1747	RP-1747 implementation starts	August 15th, 2023
	Seasonal switchover to heating	October 13th, 2023
	Probabilistic setback trial phase start	November 15th, 2023
After probabilistic setback	Probabilistic setback full implementation starts	November 28th, 2023
	SAT reset fixed	February 15th, 2024
	End of data collection	April 30th, 2024

3.2.1. RP-1747 demand-controlled ventilation

The completed ASHRAE RP-1747 is described in full detail in O'Neill et al. [18] and is summarized in O'Neill et al. [19]. Briefly, this approach uses trim and respond (T&R) logic [20,21] to dynamically adjust VAV terminal units' minimum primary airflow setpoint (V_{pz-min}) based on ASHRAE Standard 62.1-2022 [12] ventilation requirements. It is worth noting that it is the *minimum setpoint* which is adjusted; because space conditioning and ventilation are inherently linked in a VAV AHU system, the zone primary airflow (V_{pz}) is free to increase beyond the minimum position by opening the VAV terminal units' dampers if needed to meet conditioning loads.

In order to determine whether V_{pz-min} should be increased or decreased at each timestep, the logic first determined what the breathing zone outdoor airflow (V_{bz}) requirements were for each zone. V_{bz} varies depending on which sensors (if any) are available in each zone. For zones with CO_2 sensors and motion detectors where occupied-standby mode was allowed (i.e., the perimeter private offices, or zones 6–8, 10–12, 16, 17, 19, 20, and 22 in Fig. 1), as per Hobson et al. [22]:

$$V_{bz} = min\begin{cases} S\left(P_z R_p + A_z R_a\right) \\ \dfrac{max[0,\,(C_{bz} - C_s)\,]E_z V_{pz} R_p}{k \bullet m} + A_z R_a \end{cases} \quad (1)$$

where S is the binary zone occupied status based on the motion detector (i.e., 0 is unoccupied, 1 is occupied); P_z and A_z are the zone population and area, respectively (recall Table 1); R_p and R_a are the people and area outdoor airflow rates, respectively, per Table 6-1 of ASHRAE Standard 62.1-2022 [12]; C_{bz} and C_s are the breathing zone CO_2 concentrations and supply air CO_2 concentrations leaving the AHU, respectively; k is the CO_2 generation rate per occupant (assumed to be 0.00383 L/s/met/person); and m is the metabolic rate (assumed to be 1.2 met for office work).

For zones with motion detectors where occupied-standby mode was allowed (i.e., zones 1, 3–5, 9, 13–15, 18, 21, 23–28 in Fig. 1), as per Appendix A of RP-1747 [19]:

$$V_{bz} = S\left(P_z R_p + A_z R_a\right) \quad (2)$$

For all other zones (i.e., the washroom, or zone 2 in Fig. 1) as per ASHRAE Standard 62.1-2022 [12]:

$$V_{bz} = P_z R_p + A_z R_a \quad (3)$$

The zone outdoor airflow (V_{oz}) was then determined using each zone's air distribution effectiveness (E_z) per Table 6-4 of ASHRAE Standard 62.1-2022 [12] at the given timestep (i.e., Equation 6–2 of ASHRAE Standard 62.1-2022 [12]):

$$V_{oz} = V_{bz}/E_z \quad (4)$$

Once V_{oz} is determined, the zone's primary outdoor air fraction (Z_{pz}) can be calculated by dividing V_{oz} by V_{pz} (i.e., Equation A-3 from ASHRAE Standard 62.1-2022 [12]):

$$Z_{pz} = V_{oz}/V_{pz} \quad (5)$$

Z_{pz} can then be divided by the maximum zone primary outdoor air fraction (Z_p) to determine the criticality ratio for each zone (CR_z) [19]:

$$CR_z = Z_{pz}/Z_p \quad (6)$$

which essentially represents how much outdoor air is entering a zone, relative to ASHRAE Standard 62.1-2022 [12] ventilation requirements and the maximum amount the VAV can provide at the AHUs current OAR. The OAR was determined using the OAD position and the following empirically derived formula for the case study building's AHU [36]:

$$OAR = \frac{5.69}{1 + e^{(-3.01 \times OAD + 4.61)}} \quad (7)$$

The zone criticality ratio (CR_z) were then used to determine whether V_{pz-min} should be increased or decreased at each timestep. If CR_z was high (i.e., above 80 %), requests (R) were sent to the VAV controller for additional airflow, until CR_z dropped below 70 %. The number of requests increased depending on the CR_z and the importance multiplier (M) associated with that zone. M was used to ensure rogue zones (i.e., zones that, either due to inappropriately sized equipment or sensor errors, were chronically generating requests due to elevated CR_z) were ignored (i.e., by setting M equal to zero). The threshold where a zone was considered a rogue zone was if requests were generated for over 70 % of operating hours [20]. The default M for all zones is one. The number of requests R was equal to M (i.e., one by default, zero for a rogue zone) when CR_z was less than 90 %; above 90 %, the number of requests R was equal to 2 M (i.e., two by default, zero for a rogue zone). When requests are generated, V_{pz-min} will 'respond' by increasing (i.e., by an amount $(R - I) \times SP_{res}$) to allow higher ventilation in the zone, until CR_z reaches below 70 %. I represented the number of ignored requests, which could be adjusted for zones as needed to reduce equipment cycling. By default, all zones had an I equal to zero (i.e., all requests were considered). The maximum amount V_{pz-min} could respond by at each timestep was governed by $SP_{res-max}$. If CR_z was low (i.e., below 70 %), V_{pz-min} would 'trim' (i.e., by an amount SP_{trim}) to reduce ventilation to the zone, until CR_z reached above 80 %. V_{pz-min} was bounded by a maximum (i.e., SP_{max}, equal to one) and minimum (i.e., SP_{min}, equal to zero) value based on the zones' VAV capacities. This logic is summarized in Fig. 3. See Table 3 for normalized T&R parameters used in the case study building. Each normalized parameter is multiplied by the VAV terminal units' maximum zone primary airflow rates (V_{pz-max}), recall Table 1.

At the same time, the AHU will dynamically determine the required OAR using the ventilation rate procedure (VRP) outlined in ASHRAE Standard 62.1-2022 [12]. This resulted in a decrease in V_{pz-min} when the primary air was rich with outdoor air (i.e., when the AHU was economizing), which lowers reheat provided by terminal devices (i.e., baseboard units) and fan energy use; when the OAR was low (i.e., when the AHU was mechanically heating or cooling), V_{pz-min} could be increased to the most critical zones, reducing the conditioning load associated with the incoming outdoor air with a marginal increase in reheat, see Fig. 4. This complex self-balancing process occurred at each timestep, thus reducing energy use while maintaining adequate ventilation in occupied zones. Further details on RP-1747 are available in the formative literature [17–19].

The RP-1747 DCV logic was fully implemented in the case study building on August 15th, 2023, and ran for the duration of the study until

Table 3
Trim and respond parameters.

Parameter	Value
SP_{min}	0
SP_{max}	1
SP_{trim}	0.08
SP_{res}	0.12
$SP_{res-max}$	0.2
I	0
M	1

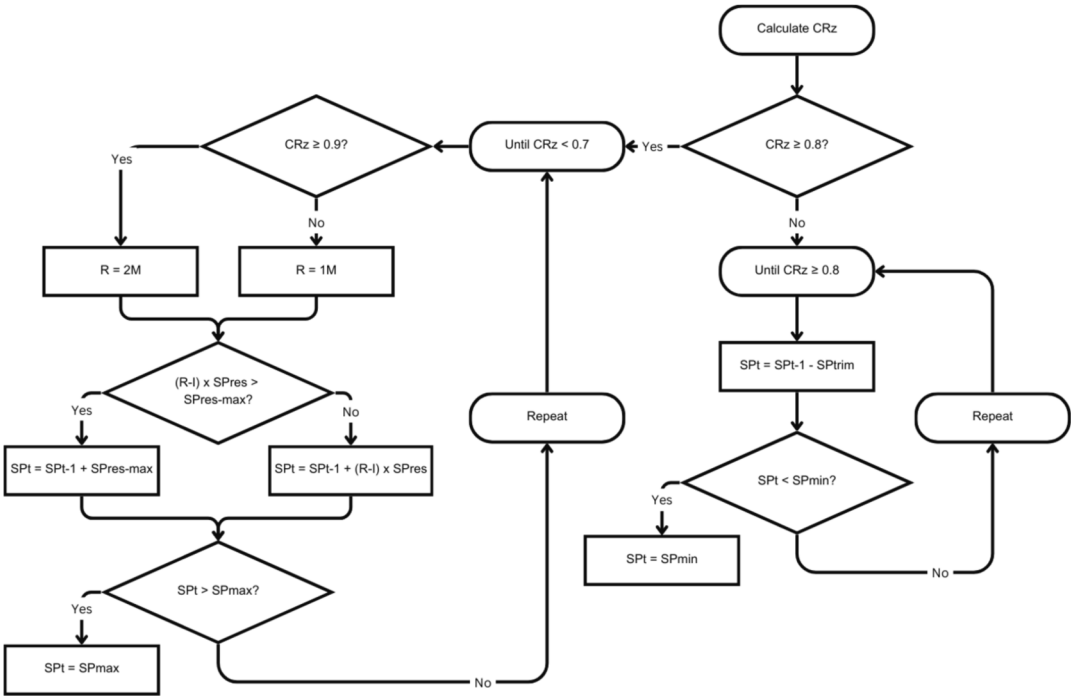

Fig. 3. Trim and respond (T&R) control logic for each VAV terminal unit.

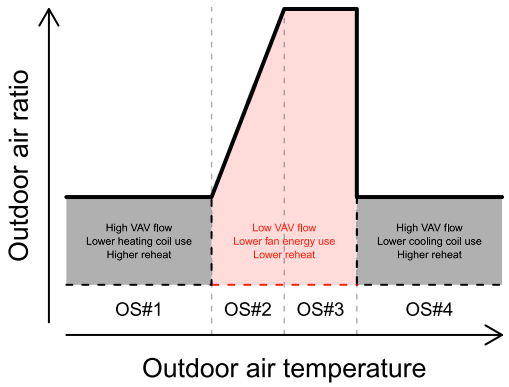

Fig. 4. Conceptual diagram of an air handling unit's (AHU) outdoor air ratio (OAR) traditionally (solid black line) versus when RP-1747 demand-controlled ventilation (DCV) is enabled. The shaded regions represent areas where the OAR and resulting energy use can be reduced.

April 30th, 2024. As the seasonal switchover to heating occurred on October 13th, 2023, the RP-1747 implementation occurred over approximately two months of the cooling season and six and a half months of the heating season.

3.2.2. *Probabilistic zone temperature setback*

Recall that, while ZAT setpoint setbacks and setups occurred after-hours during the heating and cooling season, respectively, a constant 22 °C default ZAT setpoint was utilized in all 28 zones of the case study building during operating hours, regardless of the presence of occupants. However, the presence of occupants within each zone in the case study building was highly variable. Using binary occupancy data from the motion detectors in each zone (excluding the washroom, recall Fig. 1), the first and last motion detection events in each zone can be leveraged to estimate occupants' arrivals and departures each day. Due to the case study buildings occupancy (i.e., institutional office), these data from September 1st, 2023, until November 15th, 2023 (i.e., after the lower-occupancy summer period and at the beginning of an academic semester) were utilized to create cumulative probability distributions functions (CDFs) for arrival and departure during workdays (i.e., Monday through Friday) during this period, as shown in Fig. 5(a) and Fig. 5(b), respectively. Fig. 5(a) shows that the cumulative probability of arrival for zones in the case study building was below 60 % during weekdays (i.e., less than three out of five days in a given work week), with a large spread of arrival and departure times; probability of arrival and arrival/departure times for each zone are summarized in Table 4.

Given this variability in zone-level occupant presence, a probabilistic temperature setback/setup scheme was devised based on observed motion detection patterns and occupants' predicted arrival times within the case study building. For this purpose, a zone is considered occupied if motion is detected for six or more continuous timesteps to avoid cleaning staff or brief visits (e.g., an occupant stopping by before or after off-site activities to drop off or pick up belongings) being erroneously

Table 4
Arrival and departure times for individual zones based on the cumulative probability distribution functions (CDFs) presented in Fig. 5.

Zone	P(Arrival)	Earliest arrival	Latest arrival	Latest departure
1	95 %	5:20	9:00	23:50
2	NA	NA	NA	NA
3	82 %	6:40	10:20	18:40
4	95 %	7:40	9:00	19:50
5	67 %	7:40	14:00	19:40
6	47 %	8:50	13:30	18:40
7	56 %	10:20	15:00	19:10
8	47 %	10:30	14:00	17:40
9	7 %	8:00	13:20	14:40
10	40 %	10:30	11:50	17:10
11	13 %	9:00	19:30	20:30
12	80 %	8:50	13:30	21:40
13	26 %	11:10	16:10	20:10
14	93 %	7:50	8:40	21:50
15	11 %	11:10	16:20	17:30
16	56 %	8:40	10:10	19:20
17	44 %	12:10	16:10	21:10
18	91 %	8:10	10:00	21:40
19	6 %	11:20	17:40	18:40
20	62 %	8:10	11:30	18:30
21	38 %	10:20	17:30	22:50
22	44 %	9:10	12:30	17:40
23	58 %	6:30	9:10	23:50
24	89 %	7:50	14:00	19:10
25	76 %	8:50	13:30	18:40
26	78 %	7:50	14:00	18:30
27	80 %	8:20	9:40	21:30
28	36 %	8:40	17:20	21:50

counted as occupancy. Generally, occupants had relatively consistent schedules from week to week: for example, an occupant might work from home on Mondays and Fridays, and work from their office from Tuesday through Thursday. If these patterns are persistent over a prescribed number of weeks, then the confidence in the probability that an occupant will or will not work from their office on any given day of the week increases, and ZAT setbacks/setups can be implemented in zones predicted to be unoccupied with increasing magnitude. Note that, while this behaviour is characteristic of the occupancy observed in the case study building (i.e., an institutional office building), occupancy patterns vary significantly for different buildings and occupancy types. The following analysis is therefore applicable to the case study building or other buildings with similar occupancy patterns.

Fig. 6 shows the frequency of each error type when using an increasing number of previous days of the week to predict if a zone will be considered occupied – note that 'occupancy' in this context refers to whether or not an occupant will occupy the space daily and not at hourly or sub-hourly intervals. When only considering the previous week (e.g., using the occupancy of the previous Monday to predict the occupancy that Monday), whether a given zone is occupied or not can be determined with approximately 45 % accuracy. For another 45 % of instances, occupancy is predicted to occur, but the room remains unoccupied (i.e., false positive), and for the remaining 10 % of instances,

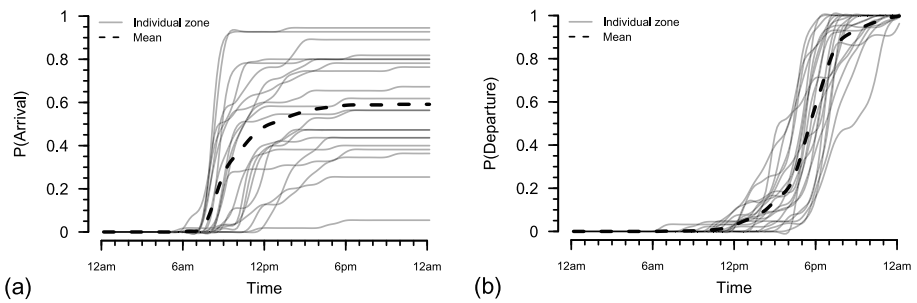

Fig. 5. Cumulative probability distribution functions (CDFs) for (a) arrival and (b) departure for zones in the case study building.

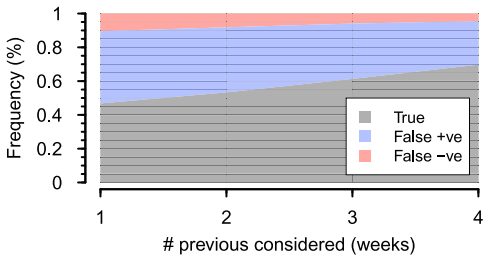

Fig. 6. Frequency of error types when predicting occupied days of week based on the same day of week from previous weeks.

occupancy is not predicted to occur, but the room becomes occupied (i.e., false negative). A false positive represents a missed opportunity for ZAT setbacks/setups during operating hours, while a false negative represents instances where a setback/setup would have been implemented while the occupant would be present, potentially causing thermal discomfort. In order to minimize the likelihood of thermal discomfort, the number of previous weeks considered was extended to four weeks (e.g., the past four Mondays), after which the likelihood of a false negative was below 5 %. Additionally, the magnitude of the setback/setup was increased as the number of weeks without occupancy on that day of the week increased, see Table 5. Conversely, in order to decrease the impact of the relatively large amount of false positives on the energy savings potential, a small ZAT setback/setup (i.e., 0.5 °C) was implemented in zones if no motion was detected after the latest expected arrival for each zone (recall Table 4) on days where occupancy was otherwise predicted; these times were hard-coded based on the latest expected arrival time determined during the training period for simplicity. These hard-coded values will differ for other buildings and zones and must be recalculated for each zone when major changes in occupancy occur. In all cases, if motion was detected in the zone, the ZAT setback/setup would be relinquished, and control of the ZAT setpoint would return to the occupant. This logic is summarized in Fig. 7.

The probabilistic setback logic was fully implemented in the case study building on November 28th, 2023, and ran for the duration of the study until April 30th, 2024. Given the seasonal switchover to heating occurred on October 13th, 2023, an estimate of the energy savings from the RP-1747 implementation alone during the heating season can be estimated, followed by the combined savings of the two interventions running simultaneously.

3.2.3. Sub-optimal sequences

As previously mentioned in Section 3.1, a sub-optimal SAT reset scheme was employed in the case study building during commissioning. This scheme compares the ZAT to the ZAT setpoints in each of the 28 zones. The difference between the ZAT and the *largest* of either the ZAT setpoint or 22.5 °C was taken for each zone, and the average difference of the eight zones with the smallest differences was taken to determine the heating demand (expressed as a negative number). Simultaneously, the difference between the ZAT and the *smallest* of either the ZAT setpoint or 22.5 °C was taken for each zone, and the average difference of the eight zones with the largest differences was taken to determine the cooling demand (expressed as a positive number). If the average difference was less than −0.5 °C, then the SAT was increased by 1 °C at each 10-minute timestep, to a maximum of 27 °C. If the average difference was greater than 0.5 °C, then the SAT was decreased by 1 °C at each 10-minute timestep, to a minimum of 13 °C.

Several issues exist with the SAT reset logic as originally programmed. In order to calculate the heating demand, the largest of the ZAT setpoint or 22.5 °C was subtracted from the ZAT. Recall that the default ZAT setpoint was 22 °C; therefore, at each timestep, the ZAT setpoint was typically lower than 22.5 °C. This logic resulted in a typical value of −0.5 °C for each zone, even if the zone was at the 22 °C ZAT setpoint. This led to many zones generating heating demand despite them being at the specified setpoint, causing the SAT to frequently increase when no actual heating demand was present, even at high OATs. This is especially suboptimal for the VAV AHU system configuration, which relies on the SAT to be low enough to provide cooling year-round due to a lack of terminal cooling devices. To remedy this, the heating demand logic was removed from the SAT reset scheme. Instead, only the cooling demand would be calculated to prioritize cooling. The cooling demand calculation was changed from the difference between the ZAT and the *smallest* of either the ZAT setpoint or 22.5 °C to the difference between the ZAT and the *largest* of either the ZAT setpoint or 20 °C (i.e., the lowest temperature the occupant could set their zone to using their thermostat), plus a 0.5 °C deadband. These changes ensured that occupants who purposefully had their ZAT setpoints above 22.5 °C, or rooms at the default 22 °C setpoint, didn't erroneously generate cooling demand as with the previous logic. Instead of the average cooling demand for the eight zones with the highest cooling demand, the number of zones which demanded cooling was used to control the SAT. If more than three zones had some cooling demand (i.e., their ZAT exceeded the largest of their ZAT or 20 °C by more than 0.5 °C), then the SAT was decreased by 1 °C at each 10-minute timestep to a minimum of 13 °C. When no cooling demand was present, the SAT was increased by 1 °C at each 10-minute timestep to the maximum of 20 °C specified during the design. This improved scheme ensured that the SAT was maximized when no cooling demand was present, but ensured the SAT did not increase so high that the supply air itself would cause overheating. It should be noted that this reduction in overheating potential would come at the cost of increased cooling load to reach the lower SATs used by this scheme. Cooling needs could then be met for the handful of zones requiring cooling, with the terminal baseboard hydronic units providing reheat in zones with heating demand.

In addition to the suboptimal SAT reset initially employed in the building, the minimum OAD position was also programmed in a suboptimal fashion during the building's commissioning. While mechanical drawings and specifications indicated that a minimum outdoor airflow rate of 965 L/s was to be utilized, the 20 % OAD position initially used in the building yielded outdoor airflow of approximately 520 L/s (recall Equation (7)). It is possible that a programming error during the building's commissioning resulted in the minimum OAR (i.e., approximately 20 % given the minimum outdoor airflow rate of 965 L/s and the total AHU capacity of 5147 L/s) being confused with the minimum OAD, or a linear relationship between the OAR and OAD was used. However, given the known relationship between the OAR and OAD in the building, the OAD was increased to 40 % to ensure that the 965 L/s specified during design was delivered during operation, see Fig. 8. Note that the empirical relationship derived in Equation (7) is within the AMCA range [37,38] for opposed-blade damper operation (i.e., the same type of damper used in the case study building's AHU). While necessary to ensure compliance with ASHRAE 62.1–2019 [16] ventilation requirements, it should also be noted that increasing the minimum OAD position will result in higher volumes of outdoor air requiring conditioning during OS#1 and OS#4 (i.e., low and high outdoor air temperatures, recall Fig. 4), reducing the savings potential of the control interventions during these periods.

Table 5
Temperature setbacks and setups for zones based on past weeks' occupancy.

Number of weeks since occupancy on this day of week	Setback/Setup (°C)
1	±0.5
2	±1
3	±2
≥4	±4

3.3. Supervisory control framework

In order to facilitate data collection, analysis, archiving, and

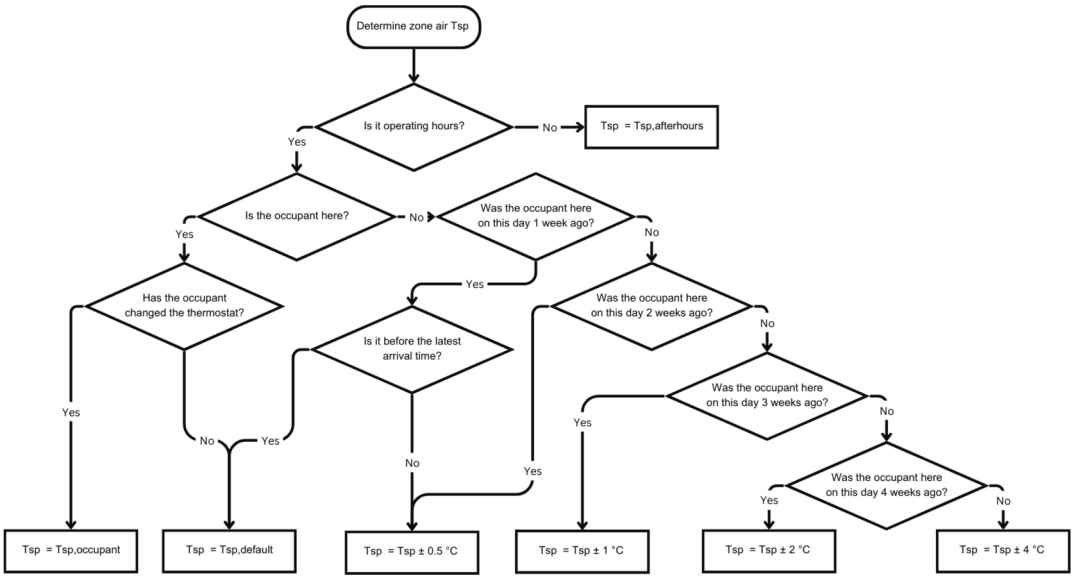

Fig. 7. Probabilistic setback/setup logic implemented in each zone.

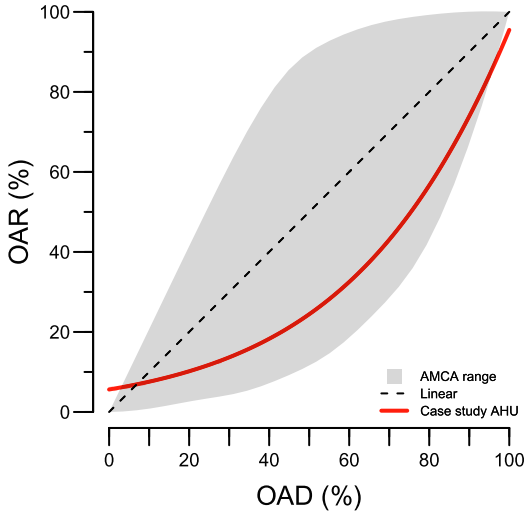

Fig. 8. Outdoor air damper (OAD) and outdoor air ratio (OAR) relationship for air handling units.

supervisory control, a custom server developed by Markus et al. [39] with bi-directional (i.e., read/write) access to the BACnet [40] protocol-based BAS was utilized. This provided several advantages for implementation. Traditionally, to implement such control changes within a BAS, programs on each of the 28 VAV controllers and the AHU controller would need to be added or altered (i.e., modification or addition of 57 programs, or two on each of the 28 VAV controllers for RP-1747 logic and probabilistic setback, and one on the AHU controller). The T&R parameters for RP-1747 DCV would need to be created as variables within each of these controllers. Similarly, the probabilistic setback would require the largest number of continuous binary motion detection events each day to be created as a variable, along with the latest arrival times for each zone and the setback/setup thresholds. Trend logs would need to be created to record parameters of interest beyond the limited buffer provided by the controller. This represents a significant burden to both control integrators and operators. Rather than providing a new program and creating additional variables within each individual controller for the control logic, calculations could be performed on the Python-based server. Variables needed to perform the calculations for RP-1747 DCV and the probabilistic temperature setback were pulled from the BAS at each 10-minute timestep via the bi-directional server. This process reduced the memory use on the controllers and allowed for rapid changes to be performed; simply put, changing the control logic in a centralized Python-based environment was significantly faster than altering the control logic in each of the 29 controllers (i.e., 28 VAVs and one AHU) individually.

The RP-1747 DCV and probabilistic setback logic each had a single Python-based script to handle the calculations for all zones. Once the Python-based scripts calculated the new V_{pz-min} and ZAT setpoints with the appropriate setback for each zone, these setpoints were written back to the existing minimum airflow rate setpoints and ZAT setpoints within each controller at 10-minute intervals. The server internally stored all additional variables required for the RP-1747 and probabilistic setback logic (i.e., without having to create these variables within the BAS or read and write them back and forth at each timestep, thus reducing network traffic) including CR_z, M, R and cumulative requests, V_{bz} requirements, the previous V_{pz-min} setpoints, the T&R parameters, the total and the largest continuous number of binary motion detection events, and the zone-level latest arrival times. These programs used the popular *Pandas* [41,42] and *NumPy* [43] libraries. However, it should not be discounted that future or other controls interventions could take full advantage of the innumerable open-source Python-based libraries to perform more advanced machine learning (ML) or other statistical analyses in real-time, which Markus et al. demonstrated in their research [39]. This far exceeds the capabilities of traditional controllers used in commercial HVAC systems.

It should be noted that the use of this type of supervisory control framework comes with its own unique challenges. One such challenge was that the V_{pz-min} and ZAT written to each controller by the server automatically inherited the lowest level of the BACnet priority array within the BAS. While this is important to ensure that the written values did not override the BAS (e.g., in a life-safety event), this also meant that the existing setpoint had higher priority than the incoming setpoint. To remedy this, a single program was added to the AHU controller to promote the incoming minimum setpoints to one level above their existing counterparts. This additional program allowed facilities personnel to simply override the RP-1747 DCV or probabilistic setback implementation, allowing for the BAS to return to its default state. As discussed in Markus et al. [39], another obstacle to the use of this framework was that manual mapping of BAS points was required to enable the retrieval of the sensor data, which required expert knowledge

of the case study building's BAS. The importance of industry-wide adoption of a standardized semantic data tagging scheme to reduce this manual and burdensome process – such as those proposed by Project Haystack [44], Brick [45], or the emerging ASHRAE Standard 223P [46] – cannot be overstated.

In summary, this use of this server allowed for an implementation that could be rapidly prototyped and was minimally invasive within the BAS, lessening the burden on facilities personnel and local controllers.

3.4. Quantification of energy use and occupant comfort

As discussed in Section 3.1, energy valves were only available for the hydronic baseboard units in the 11 exterior perimeter offices, while the radiator valve positions (S_{rad}) were available for all 28 zone. In order to determine the energy use of the hydronic baseboard units in all zones, a logarithmic function was trained using non-linear regression to describe the radiator power output as a function of the radiator valve position. The average hourly radiator power output was normalized by the radiator capacity for the 11 exterior zones (recall Table 1) where energy valves were available. 10-fold cross-validation was utilized with a test-train split of three and eight for each fold, respectively, and the model with the highest average coefficient of determination (R^2) across all folds was selected, see Fig. 9. This logarithmic function was then used to convert the radiator valve positions for the remaining 17 zones into the normalized radiator power output, which was multiplied by each individual radiator's capacity taken from the design specifications (recall Table 1). The total energy use by the hydronic baseboard units at each hourly timestep was then summed and added to the heating coil energy use to determine the buildings total steam load.

The steam and chilled water data and the OAT data from the BAS were used to train ASHRAE Guideline 14-compliant [47] three-parameter univariate changepoint models for the building before the sequences of operation were altered. These models were then used to predict these loads during the period after implementation had RP-1747 and the probabilistic setback not been employed. The predicted loads were compared to the measured data to quantify the change in steam and chilled water load attributable to the sequences of operation implemented.

In addition to energy use, a longitudinal comfort survey following the methodology developed by Berquist, Ouf, and O'Brien [48] was implemented in the case study building to monitor occupant-reported satisfaction with their perceived temperature and air quality on a five-point Likert scale for the duration of the implementation. Two electronic occupant survey devices were installed next to the main stairwell to solicit feedback from occupants as they exited the building for this purpose (i.e., one for temperature and one for air quality). Signage was posted above the electronic survey devices which asked, "How satisfied were you with the temperature in [the case study building] today?" and "How satisfied were you with the air quality in [the case study building] today?", with the response options provided in Fig. 10(a) and (b), respectively. Additionally, signage was placed beside the devices with a QR code to allow occupants to provide comments and feedback in an open-response format. At secondary exits, signage with a different QR code was provided which allowed occupants to digitally select the response options provided by the electronic survey devices in addition to the open feedback to ensure full coverage of the building.

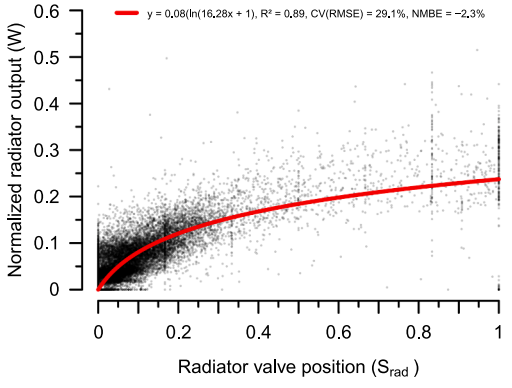

Fig. 9. Average radiator valve position versus average normalized radiator power output for the 11 zones with baseboard metering in the case study building.

4. Results and discussion

This section presents the results from the implemented sequences of operation. First, the impact of the sequences on steam and chilled water energy use are discussed. Then, the behaviour of the supply fan after implementation is briefly explored. Several example dates from zone 12 are examined to explore the impact of select control changes on the operation of zone-level equipment. The accuracy of the occupancy prediction approach used by the probabilistic setback is examined, and impacts on occupant comfort are discussed. Finally, the limitations and unresolved issues from this study are detailed.

4.1. Steam and chilled water energy use

The three-parameter univariate changepoint models for the steam and chilled water loads before the implementation of RP-1747 and the probabilistic setback are presented in Fig. 11. Note that these changepoint models fell within the calibration range for ASHRAE Guideline 14–2014 [47] for sub-hourly data (i.e., a CV(RMSE) and NMBE below 30 % and ±10 %, respectively). Using the 95 % confidence interval (CI) for the slope and intercept of each changepoint model, an estimated 2 ± 6 % reduction in the chilled water load and 30 ± 2 % reduction in steam load was found after RP-1747 was implemented. The addition of the probabilistic setback increased the estimated steam savings to 37 ± 2 % compared to the period before implementation, for a total estimated savings of 36 ± 2 % between the time RP-1747 was first implemented and the end of the study period, see Table 6.

In the case of the chilled water load, the estimated energy savings from RP-1747 were less than expected, with a non-negligible chance that the chilled water load may have increased by up to 4 % after the introduction of RP-1747 compared to the baseline period. The magnitude of this reduction (or lack thereof) can be attributed to several phenomena that occurred in the case study building, namely the aforementioned erroneous minimum OAD position as well as the change in occupancy patterns between the period before RP-1747 implementation and after.

Regarding the erroneous minimum OAD position, recall that the building was operated with a 20 % OAD position compared to the 40 % OAD position as designed during the period before implementation. This was corrected during the implementation of RP-1747. Additionally, the new minimum OAD position achieved when using the VRP was 16 %; this resulted in a nearly negligible decrease in the outdoor airflow rate given the non-linearity of the damper (recall Fig. 8); the minimum OAR was only reduced from approximately 10 % at the old 20 % minimum OAD to approximately 9 % at the new 16 % minimum OAD. Given this already low minimum OAR, many timesteps in both the heating (OS#1) and mechanical cooling (OS#4) modes saw the OAD position increase above its original 20 % position to meet ventilation demand, see Fig. 12. Although this resulted in a potential increase in chilled water load during certain timesteps, this was desirable as at the erroneous 20 % OAD position, some zones were unable to meet their ventilation requirements when fully occupied, even when operating at their maximum primary airflow rate. It is therefore expected that higher savings would have been observed in the cooling season if the building was operating at the 40 % minimum OAD position originally specified during design, as the new minimum OAR would have decreased from

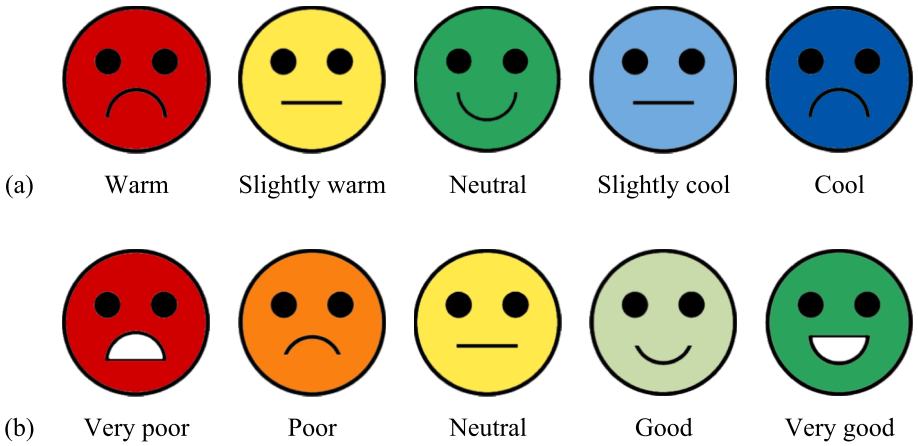

Fig. 10. Likert-scale response options for (a) temperature and (b) thermal comfort.

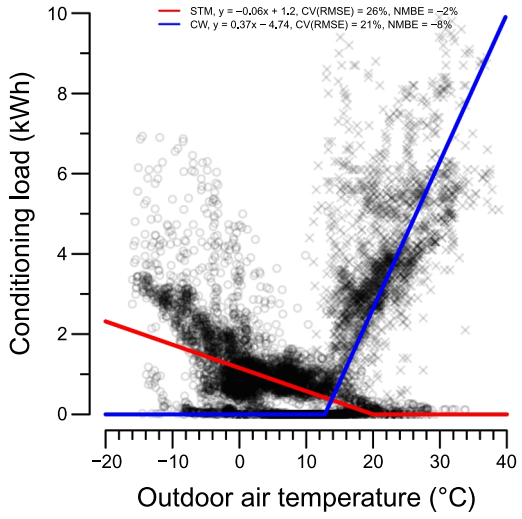

Fig. 11. Three-parameter univariate changepoint models for steam (STM) and chilled water (CW) data during the pre-implementation period.

Table 6
Measured and estimated energy usage after RP-1747 implementation.

Parameter	After measured (kWh)	After estimated* (kWh ± 95 % CI)	Estimated savings (%)
CW	6009	6136 ± 343	2 ± 6
STM (pre-setback)	2809	3995 ± 89	30 ± 2
STM (post-setback)	11,646	18,502 ± 269	37 ± 2
STM (total)	14,455	22,497 ± 380	36 ± 2

*using changepoint model.

approximately 20 % to 9 % (i.e., less than half of the volume of outdoor air would have required conditioning when ventilation needs were satisfied).

Regarding the occupancy of the case study building before and after RP-1747 implementation, it should be noted that the cooling season period before the implementation (i.e., from seasonal switchover to cooling on May 12th, 2023, up to implementation on August 15th, 2023, recall Table 2) overlapped entirely with the low-occupancy summer period in the case study building (i.e., from May through August inclusive). Subsequently, the period of RP-1747 implementation during the cooling period (i.e., from implementation on August 15th, 2023, up

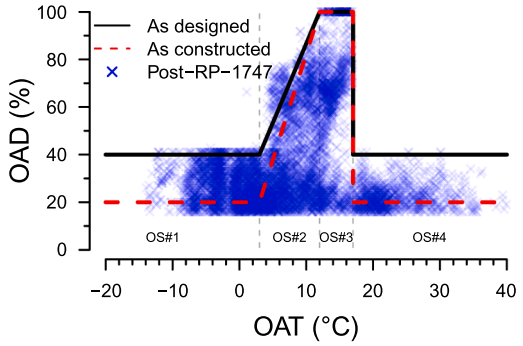

Fig. 12. Outdoor air damper (OAD) positions relative to the outdoor air temperature (OAT) for the case study building's air handling unit (AHU) after RP-1747 implementation.

to seasonal switchover to heating on October 13th, 2023, recall Table 2) overlapped largely with the relatively higher occupancy academic semester (i.e., from September through December inclusive). The difference in occupancy between these two periods is highlighted in Fig. 13. It is therefore expected that the higher occupancy for the majority of the post-RP-1747 implementation period would have contributed to higher internal gains, resulting in an increase in cooling demand which is not captured by the univariate (i.e., temperature-based) changepoint model.

Due to the heating season overlapping with the higher occupancy academic semesters, a higher magnitude of savings was determined with higher certainty than the chilled water data. As previously stated, the estimated reduction in steam load was 30 ± 2 % after RP-1747 implementation. While the minimum OAD increased above the original 20 %

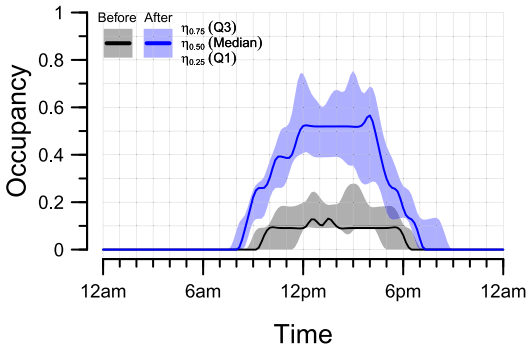

Fig. 13. Workday occupancy profiles at each quantile for the periods before and after implementation.

position during heating mode (OS#1) frequently to meet ventilation demand during occupied hours, the overall system primary airflow was reduced by approximately 8 % (i.e., from an average of 53 ± 20 % before implementation to 45 ± 21 % after implementation) during operating hours, see Fig. 14(a). This can be attributed to the minimum zone primary airflow rate setpoint reducing to 0 L/s during unoccupied hours in zones with motion detection (i.e., all zones except zone 2, recall Fig. 1) where occupied-standby mode was allowed. Before implementation, no zones had a near-zero zone minimum primary airflow rate during operating hours, whereas after implementation, an average of 34.8 % of zones had a near-zero minimum primary airflow rate (i.e., were considered unoccupied) during operation. This aligns with the occupancy patterns in the building presented in Fig. 5, where the average cumulative probability of arrival for any given zone in the case study building was below 60 % on any given workday. It should be noted that, despite the reduction in the system primary airflow, the fraction of zones with zone primary airflow rates above 70 % of the VAVs capacity increased from 8.9 % to 17.3 % after RP-1747 implementation. Therefore, the overall system primary airflow was reduced due to occupied-standby mode in unoccupied zones and lower minimum zone primary airflow rates in occupied zones during timesteps where ventilation governed, however, zones with higher ventilation demands received higher volumes of ventilation air when needed. This resulted in an increased zone outdoor airflow rate for occupied zones, see Fig. 14(b), despite an overall decrease in system primary airflow. Simply put, RP-1747 allowed for a higher volume of outdoor air to reach zones with ventilation demand at lower overall system airflow rates, improving ventilation while reducing conditioning loads, especially during the heating season, as well as fan energy use.

4.2. Supply fan energy use

An unintended consequence of the higher fraction of VAVs with elevated minimum primary airflow rates was a reduction in the duct static pressure. The AHU in the case study building was programmed to operate at a duct static pressure between 400 Pa and 150 Pa, which decreased as the number of open VAV dampers increased. Recall that, before RP-1747 implementation, each VAV had a hard-coded static minimum primary airflow rate setpoint in the range of 10 % to 30 % of their maximum capacity, see Table 1. These minimum airflow setpoints resulted in relatively low and constant VAV damper positions across the case study building, with the duct static pressure remaining at the maximum 400 Pa. However, after RP-1747 implementation, the number of VAV with elevated minimum primary airflow rate setpoints (and therefore dampers positions) increased (recall RP-1747 aims to keep the CR_z between 70 % and 80 % for zones) as previously mentioned. The duct static pressure was subsequently reduced to approximately 250 Pa, resulting in a significant reduction in fan energy use. Consider the rearranged form of the affinity laws:

$$\frac{HP_2}{HP_1} = \left(\frac{SP_2}{SP_1}\right)^{\frac{3}{2}} = \left(\frac{250\,Pa}{400\,Pa}\right)^{\frac{3}{2}} = 49.4\% \quad (8)$$

where HP and SP are the horsepower and static pressure, respectively. Therefore, nearly 50 % of the electricity demand of the fan can be reduced by operating at the lower duct static pressure, potentially representing the largest change in the fan energy use before and after RP-1747 implementation, though not directly attributable to the RP-1747 control logic itself. Fig. 15 shows the frequency of loads experienced by the supply fan's VFD before and after the implementation of RP-1747 at each 10-minute timestep, with the fan load switching from a left-skewed distribution to a right-skewed distribution (i.e., a higher proportion of timesteps spent at lower fan load after implementation of RP-1747). The average reduction in the supply fan energy use was 35 ± 26 %.

4.3. Zone-level equipment operation

The impact of RP-1747 and the probabilistic setback on zone-level equipment operation will be explored through the example of zone 12. Typical daily operation of this zone during the heating season before any changes to the sequences of operation are depicted in Fig. 16. Note that, from 6 AM to 10 PM (i.e., the operating hours during weekdays), V_{pz-min} remains constant, even before occupancy was detected, despite the ZAT remaining within the deadband of the setpoint. As the zone becomes occupied, the occupant and their activities (e.g., artificial lighting usage, computer or other plug-in equipment usage) begin to generate heat, causing the ZAT to rise above the temperature deadband. The VAV responds by increasing zone primary airflow to the zone, allowing the supply air to cool the space in response to the cooling demand. Similarly, as the ZAT falls below the temperature deadband, the hydronic baseboard terminal unit responds by opening the radiator valve to heat the zone in response to the heating demand. In this configuration, the zone primary airflow rate is driven entirely by the cooling demand of the zone. Both the VAV and hydronic baseboard unit will respond to thermal loads when the ZAT is outside the deadband, regardless of occupancy.

After the implementation of RP-1747, the zone minimum primary

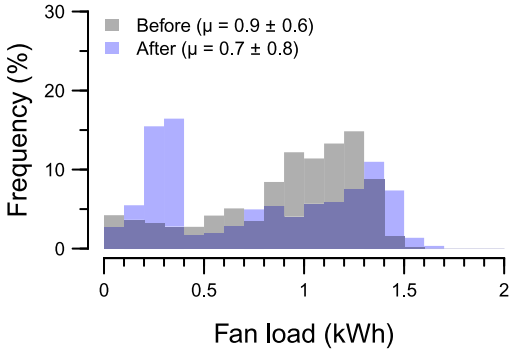

Fig. 15. Frequency of loads experienced by the supply fan before and after RP-1747 implementation.

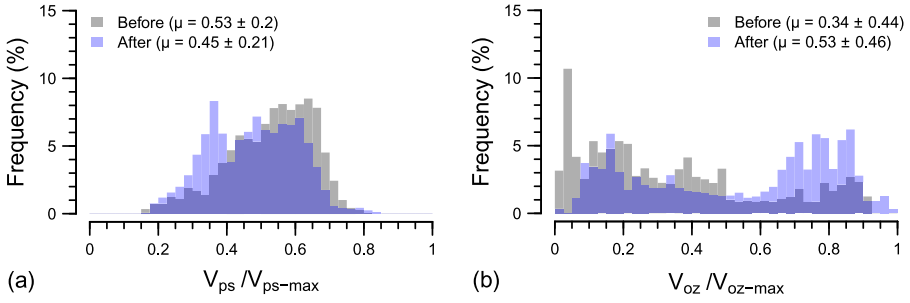

Fig. 14. The proportion of (a) system primary airflow (V_{ps}) and (b) zone outdoor airflow (V_{oz}) before and after RP-1747 implementation.

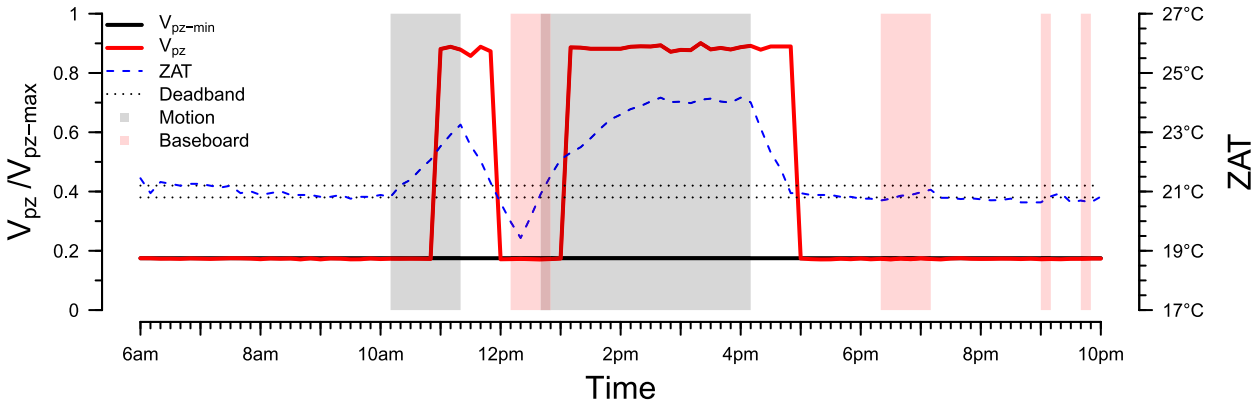

Fig. 16. Zone primary airflow (V_{pz}) ratio, zone air temperature (ZAT), motion detection, and baseboard actuation for zone 12 on Friday, February 24th, 2023 (i.e., before interventions).

airflow rate is no longer static, see Fig. 17. During operating hours, but before the zone is occupied, the zone minimum primary airflow can be outright eliminated as there is no ventilation demand in the zone (i.e., motion is not detected, and the CO_2 concentrations within the zone have not increased enough to cause ventilation demand, recall Equation (1)). This occurs until the zone becomes occupied. As this zone is equipped with zone-level CO_2 sensing, there is a lag of several timesteps between the generation of the CO_2 in the space upon the occupant's arrival, and the detection of the CO_2 by the wall-mounted BAS-integrated thermostat, including the time it takes for the CO_2 concentration to increase sufficiently such that a demand for ventilation is determined. As these concentrations increase, the required zone outdoor airflow rate (V_{oz}) increases (recall Equation (4)), which in turn causes the zone primary outdoor air fraction (Z_{pz}) to increase (recall Equation (5)), which will increase the criticality ratio of the zone (CR_z) if the OAR is unchanged (recall Equation (6)). As the CR_z exceeds 90 %, two requests (R) are generated, and the V_{pz-min} responds by increasing by $SP_{res-max}$ (i.e., 20 % of the VAVs capacity, recall Table 3). By 10 AM, the CR_z drops below 90 %, and a single R is generated. V_{pz-min} responds by increasing by SP_{res} (i.e., 12 % of the VAVs capacity, recall Table 3). By this time of day, enough zones have generated requests for ventilation that the OAR has begun to increase, and the CR_z of the zone begins to decrease at the same zone primary airflow. V_{pz-min} then begins to trim by SP_{trim} (i.e., 8 % of the VAVs capacity, recall Table 3). Shortly thereafter, the occupant vacates the space, and V_{pz-min} continues to trim down as the required breathing zone outdoor airflow rate (V_{bz}) for the zone becomes 0 L/s. However, while V_{pz-min} continues to trim down all the way, the zone primary outdoor airflow rate (V_{pz}) increases in response to a thermal load in the space, returning to V_{pz-min} once the thermal demand is satisfied. The same process repeats itself as the occupant returns to their zone for another 90-minute period of occupancy. Upon their departure, increased solar gains from the afternoon sun result in a cooling demand being generated in the space, and V_{pz} is increased to provide cooling during this time. Note that V_{pz-min} remains at 0 L/s during this period.

In both Fig. 16 and Fig. 17, the zone primary airflow rate (V_{pz}) increases in response to cooling demands regardless of the minimum zone primary airflow rate setpoint (V_{pz-min}), and the hydronic baseboard unit's radiator valve opens in response to heating demand. However, further energy savings are possible by providing the probabilistic setback to the zone, thus allowing for the ZAT to increase or decrease above the zone's default setpoint without causing further cooling or heating demand, respectively. In the case of the former, this results in V_{pz} remaining at V_{pz-min} for a larger proportion of operating hours, thus reducing the overall volume of outdoor air requiring conditioning and contributing to further energy savings. Consider Fig. 18, which depicts a relatively warm spring day. Zone 12 had a ZAT setpoint of 21 °C on this day. With the typical 21.5 °C ZAT used by the occupant in this zone, cooling demand would have caused V_{pz} to increase above V_{pz-min} at approximately 1:40 PM for a large portion of operating hours. However, with the 2.5 °C temperature setup (i.e., the setup temperature used when less than an hour of consecutive occupancy had been observed on the past three weeks on the same day of the week, or Monday in the case of Fig. 18), the ZAT is allowed to increase above the default temperature setpoint as no occupancy is expected; as a result, V_{pz} remains at V_{pz-min} for the duration of the day, eliminating ventilation air to the zone and thus reducing the volume of outdoor air requiring conditioning as well as the system primary airflow.

Recall that the probabilistic setback would be relinquished if an

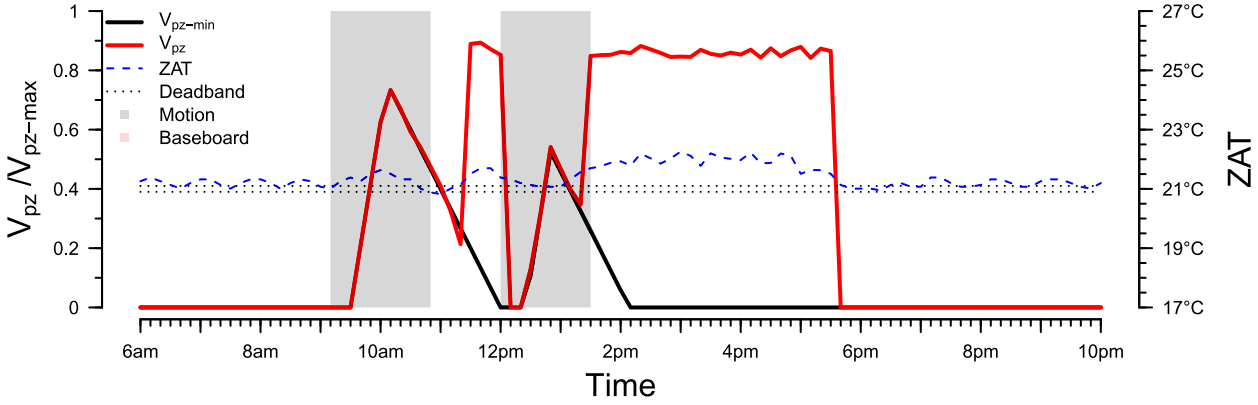

Fig. 17. Zone primary airflow (V_{pz}) ratio, zone air temperature (ZAT), motion detection, and baseboard actuation for zone 12 on Friday, November 17th, 2023 (i.e., after initial RP-1747 implementation).

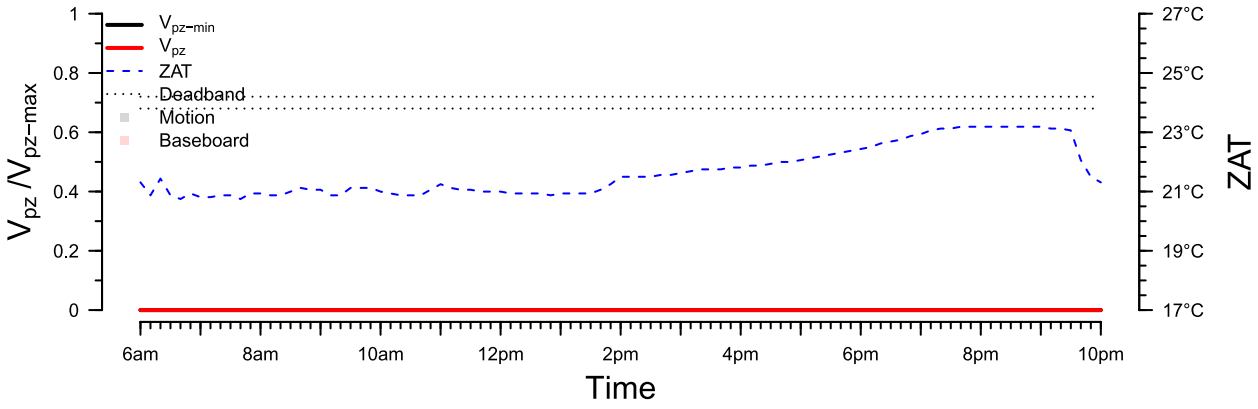

Fig. 18. Zone primary airflow (V_{pz}) ratio, zone air temperature (ZAT), motion detection, and baseboard actuation for zone 12 on Monday, April 30th, 2024 (i.e., an unoccupied day after probabilistic setback implementation).

occupant entered a zone. An example of this behaviour is provided in Fig. 19. The ZAT setpoint of 22 °C was setback by 2.5 °C during operating hours. In the early morning, the hydronic baseboard unit responds to heating demands as the ZAT falls below the setback temperature; note that the baseboard would have had to run for a longer duration or at a higher capacity to increase the ZAT to the default setpoint compared to the setback. At approximately 12 PM, an occupant is detected in the zone, and the ZAT and deadband return to their default of 22 °C and 0.1 °C, respectively. There is now simultaneous demand for V_{pz} to increase as V_{pz-min} and CR_z increase, as well as heating demand to increase the ZAT to the new setpoint. Note that the increased volume of ventilation air (i.e., at an SAT that is at or below the ZAT) would result in an additional heating load on the hydronic baseboards. This tradeoff is to be expected with RP-1747 DCV; by providing increased zone primary airflow, while the zone-level terminal reheat devices may experience a small increase in their load, this is less energy intensive than increasing the outdoor airflow rate for the entire system to meet the critical zone's ventilation demand. These changes in heating loads are highlighted in Table 7, which shows that the proportion of the heating load provided by the hydronic baseboards increased from 3.4 % before the introduction of RP-1747 to 11.1 % after its introduction, with a further 1.1 % increase after the introduction of the probabilistic setback. Once the ventilation demand is met and CR_z begins to decrease, V_{pz-min} can then begin to trim down to 0 L/s as the room remains unoccupied. At this point, because the probabilistic setback has been relinquished due to an occupancy event, the hydronic baseboard continues to actuate to achieve the default ZAT setpoint, even after the occupant has left. Simply put, the addition of the probabilistic setback increased the number of hours where V_{pz} followed V_{pz-min}; before the introduction of the probabilistic setback, the proportion of operating hours where V_{pz} equaled V_{pz-min} was 54.5 %, whereas after the introduction of the probabilistic setback, this proportion increased to 72.9 %. This resulted in reduced zone and system primary airflow to zones during periods where cooling would have been otherwise requested at the default ZAT and deadband.

Table 7
Distribution of steam load between air handling unit (AHU) heating coil and zone-level terminal devices (i.e., hydronic baseboard units).

Phase	Proportion of total STM load (%)	
	Heating coil	Baseboard
Before	96.6	3.4
After RP-1747	88.9	11.1
After setback	87.8	12.2

4.4. Probabilistic zone temperature setback accuracy

It should be noted that the ZAT in zone 12 in Fig. 19 took approximately 90-minutes to increase back to the desired ZAT setpoint after the setback, with the ZAT remaining below the ZAT setpoint for nearly the entire duration of the occupant's presence in the zone. This occurrence represents an instance of a false negative occupancy prediction by the probabilistic setback. Overall, the total proportion of true predictions was 56.4 %, below the ~70 % true prediction accuracy expected (recall Fig. 7). False positives accounted for 33.1 % of predictions, whereas false negatives accounted for 10.5 % of predictions. The number of false negatives was higher than the 5 % anticipated. Of the false negatives,

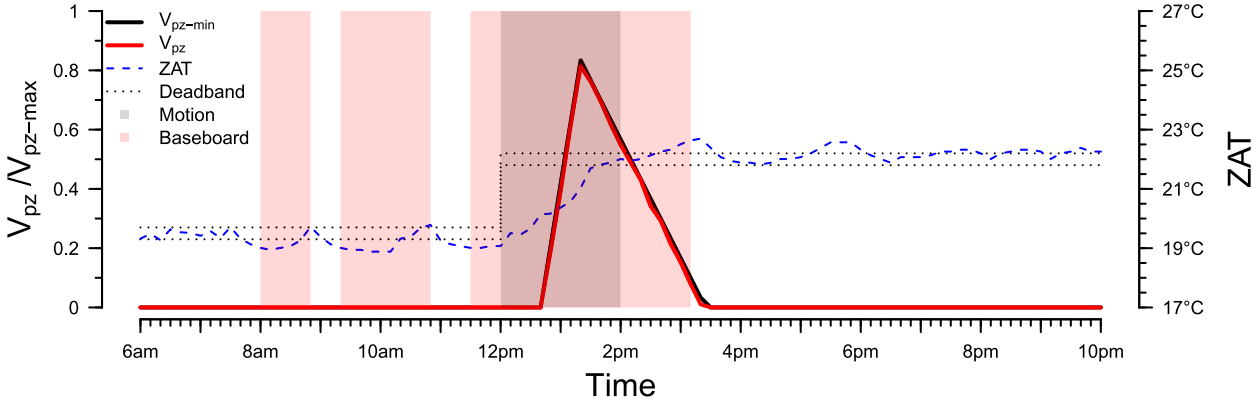

Fig. 19. Zone primary airflow (V_{pz}) ratio, zone air temperature (ZAT), motion detection, and baseboard actuation for zone 12 on Wednesday, January 17th, 2024 (i.e., after probabilistic setback implementation).

2.7 % were at the 4 °C setback, 1.0 % were at the 2 °C setback, 1.9 % were at the 1 °C setback, and 4.9 % were at the 0.5 °C setback. The proportion of false negatives may have been higher due to changes in occupants' patterns between semesters (i.e., from the fall semester at the end of 2023 and the beginning of the winter semester at the start of 2024), as it would take the probabilistic setback scheme approximately four weeks to regain the required confidence to introduce the full 4 °C setup/setback. Additionally, holidays (such as lieu days on Mondays for holidays which fell on weekends) caused an elevated number of false negatives at the 0.5 °C setback range, which accounted for nearly 50 % of all false negatives. As such, the probabilistic setback scheme could be improved by directly accounting for holidays in the control logic, as well as by resetting the probabilities during major changes in occupancy within the building (e.g., in the case of the case study building, between semesters).

Despite the lower-than-expected accuracy of the probabilistic setback scheme, the proportion of occupied timesteps spent at or around the ZAT setpoint remained relatively unchanged in the case study building, with an average deviation of 0.3 ± 1.4 °C before the implementation of the probabilistic setback versus an average deviation of 0.2 ± 1.3 °C after the implementation, see Fig. 20. Therefore, considering the 7 % increase in heating energy savings using this approach (i.e., from 30 ± 2 % with RP-1747 only compared to 37 ± 2 % after the introduction of the probabilistic setback, recall Table 6), the benefits of a setup/setback scheme pair with the changes to zone-level ventilation rates are evident.

4.5. Impact on occupant comfort

While the longitudinal comfort survey data were meant to be used in the assessment of whether the occupants' perception of their thermal comfort and the IAQ changed over the course of the implementation, occupant participation in the longitudinal thermal comfort survey was minimal. Over the course of the study, only three comfort survey responses were received via the QR code interface, and only ten individual interactions occurred with the electronic survey devices. No occupants provided feedback via the open-response format, and all interactions occurred within the first month of the installation of the signage and electronic survey devices. Possible reasons for the lack of interactions include survey fatigue or occupants' apathy towards the operation of the building. Another mitigating factor may have been occupants' concerns with other aspects of the building's operation which overshadowed the changes made over the course of the implementation. These include persistent problems with lighting controls on the floor (e.g., inappropriate lighting on/off behaviour, light flickering, etc.), acoustic properties of the floor (e.g., the open floor plan in zone 24, insufficient soundproofing around ductwork, etc.), and water drainage issues on the roof which caused leakage through the ceiling. Simply put, occupants' perceived comfort may have been more severely impacted by these issues than by the changes introduced by either RP-1747 DCV or the probabilistic setback. In any event, insufficient data was gathered from the longitudinal comfort survey to draw any meaningful conclusions about the occupants' perceived thermal comfort or IAQ as a result of the RP-1747 implementation.

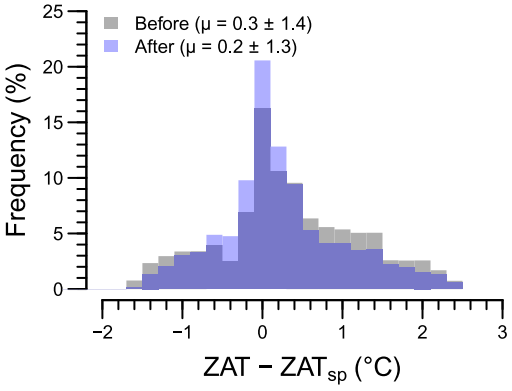

Fig. 20. Differential between zone air temperature (ZAT) and setpoint (ZAT_{sp}) during occupied hours before and after implementation.

While data on occupants' perceived thermal comfort and IAQ were not available, the operational data from the BAS considered over the course of this study can be examined to determine how RP-1747 DCV and the probabilistic setback impacted the environmental conditions in the building using the zone air temperatures and outdoor airflow rates. Recall in Fig. 20 that the differences between the zone air temperature and zone air temperature setpoints were minimal before and after implementation (i.e., 0.3 ± 1.4 °C versus 0.2 ± 1.3 °C, respectively), which indicates that the overall thermal conditions – and presumably the occupants' level of comfort under these thermal conditions – remained relatively unchanged. At the same time, the proportion of the zone air which was composed of outdoor air increased from an average of 34 ± 44 % to 53 ± 46 % before and after implementation, respectively (recall Fig. 14(b)), meaning per person ventilation was improved in the occupied zones after implementation. Therefore, while qualitative and quantitative data were not available to assess occupants' perceived thermal comfort and IAQ, the BAS data suggest that the building's ability to meet setpoint was unchanged and ventilation rates were improved after the implementation of RP-1747 DCV and the probabilistic setback.

4.6. Limitations and unresolved issues

The results presented throughout this study should be considered with knowledge of the overarching limitations and unresolved issues, namely:

- *How can the impact on occupant comfort be quantified?*

The intended methodology to assess occupant comfort over the course of this study was largely unsuccessful as discussed in Section 4.5. Beyond the mitigating factors of the building's conditions, improvements can be made to the methodology to increase the frequency of occupants' responses in future studies. These include providing compensation to encourage occupant participation, or having detailed discussions with the occupants at the onset of the survey so that they are explicitly aware of and feel invested in the survey process. Additionally, occupants' feedback can be actively solicited through wearable devices (e.g., [49,50]) – which has been shown to reduce survey fatigue [49] – although ethical and privacy implications must be considered. While the operational data gathered over the course of this study can be used as a crude approximation of the impact of these interventions on environmental conditions which affect occupant comfort, their limitations should be acknowledged. These include the neglect of other considerations for thermal comfort, such as air movement, occupants' position (e.g., proximity to radiating devices, proximity to supply grilles, the use of sit/stand desks, etc.), adaptive actions taken by the occupants, as well as occupants' diverse preferences when it comes to their thermal environment. Similarly, CO_2 concentrations and per person ventilation rates within zones do not capture other aspects of IAQ such as the concentrations of VOCs and other pollutants. Ideally, these qualitative and quantitative data should be combined to increase the richness of the dataset with respect to occupant comfort and behaviour [51].

- *What buildings are these results applicable to?*

The findings presented herein are applicable to the case study building described in Section 3.1, located in ASHRAE Climate Zone 6A. This building had granular sensing infrastructure and thermal zoning and was composed primarily of private offices. A decrease in the density

of thermal zones (e.g., more offices controlled by a single VAV) would likely reduce savings from interventions such as RP-1747 DCV and the probabilistic setback due to occupied-standby mode being less likely to occur, as the likelihood of occupancy being observed on any given day would be higher when considering multiple occupants. This is demonstrated by the average higher probability of arrival in the open offices with multiple occupants (i.e., 95 % in zone 4 and 89 % in zone 24) compared to the individual private offices (i.e., an average of 49 %). A decrease in the granularity of the sensing structure would also correspond to reduced savings due to less zones using occupied-standby mode in the case of motion detectors, or less zones using a lower person fraction ($P_z R_P$) in the case of CO_2 sensors. This phenomenon was explored in Hobson et al. [52]. For 11 zones with CO_2 sensing and all zone with motion detection (i.e., the configuration used in the real case-study building), they estimated that median HVAC energy savings would be approximately 31 % after the implementation of RP-1747 DCV in their simulation-based study. These results are in line with the findings of this study, which estimated HVAC energy savings from 28 to 32 % after the initial RP-1747 DCV implementation. They estimated that median HVAC savings with other sensor configurations could range from as little as 3.5 % with a single CO_2 sensor and no motion detection, all the way to nearly 50 % if CO_2 sensors and motion detectors were installed in every single zone. While neither of these studies assessed the impact of different climate zones on the magnitude of these savings, recall that O'Neill et al. [18] simulated RP-1747 DCV in ASHRAE Climate Zones 1A, 3A, 3B, and 5A and found savings increased from 9 % up to 33 % from the warmest to coldest climate zone. However, further implementations of RP-1747 DCV in buildings with different geometries, zoning, sensor configurations, and across different climate zones will help to better characterize the impact of these properties on the efficacy of this approach.

- *What is the impact of CO_2 sensor precision and accuracy on savings potential?*

In its current form, RP-1747 is a CO_2-based DCV approach. However, the placement of CO_2 sensors within the zone, location of CO_2 generating sources (e.g., occupants) and supply and return grilles, as well as the layout of doors, windows, and furniture within the zone will all influence the CO_2 concentrations at the sensor location due to the non-homogeneity of the zone air [53]. Due to this, there can be a lag time between the onset of CO_2 generation and detection by the sensor; this is partially reflected in the delay between motion detection in the zone and V_{pz-min} beginning to increase in zones where CO_2 sensors are utilized for DCV, recall Fig. 17 and Fig. 19. Additionally, while the CO_2 sensing infrastructure was verified to be functional and within manufacturer-specified calibration ranges at the onset of this study, CO_2 sensors are notorious for quality control issues which impact accuracy as well as decreasing precision due to a lack of calibration over prolonged periods [53–56]. In their study, Lu et al. [57] found that typical errors in airflow and CO_2 sensors could reduce the annual HVAC energy savings from RP-1747 DCV by as much as 17 %. Therefore, it is critical that calibration practices [58,59] for these sensors are followed in order for the full potential of RP-1747 DCV to be realized.

- *Can other sensors be used for DCV purposes?*

Given that the CO_2 concentrations from the zone-level sensors are meant to account for the person fraction of the ($P_z R_P$) of the breathing zone outdoor airflow rate (V_{bz}, recall Equation (1)), there exists potential for the data from CO_2 sensors to be replaced with infrastructure which count occupants directly for DCV purposes. Current occupant-counting sensing technologies (e.g., people-counting cameras, Wi-Fi or Bluetooth-enabled device counts [60–62], etc.) may be cost-prohibitive to deploy at the granularity needed to achieve zone-level DCV. Other sources, such as badge in/out data (e.g., [63]) may be ill-suited for buildings with more transient or otherwise uncontrolled occupant access. In both cases, these data might come from sources that are disparate from the BAS, such as IT or security networks and infrastructure. However, the supervisory control approach used in this study opens the door to a variety of possibilities for using zone-level occupant counts in DCV directly. The bi-directional access to and existence outside of the BAS allows for data from other sources to be accessed through the custom server used in this study (e.g., Wi-Fi-enabled device counts from the IT infrastructure). Given the server's ability to use Python and the open-source libraries it provides, these data can be amalgamated using a variety of ML approaches which rely on sensor fusion (e.g., [64–67]) that would not be possible with a traditional BAS. In short, the use of a similar supervisory control approach undertaken in this study may allow for the occupancy of zones to be measured or deduced in real-time by taking advantage of data and/or computational resources that exist outside of a traditional BAS, thus supplanting the need for CO_2 sensing infrastructure and the mass-balance approach employed in Equation (1) to be used directly.

- *What energy savings can be expected for RP-1747 during the cooling season?*

The reduction in chilled water load after RP-1747 was implemented was estimated to be between −4 and 8 %. In reality, the reduction in chilled water load is expected to be higher had the occupancy of the case study building been similar in the periods before and after the RP-1747 implementation. While RP-1747 was able to reduce the V_{pz-min} in some cases during the cooling season, the suboptimal SAT reset strategy originally used in the building as well as a lack of zone temperature setbacks in zones resulted in frequent occurrences of V_{pz} having to exceed V_{pz-min} to meet conditioning loads. It was shown that fixing the SAT reset scheme in tandem with the probabilistic setback reduced this frequency from 45.5 % of operating hours to only 27.1 % of operating hours during the heating season, however, the impact on the cooling season is unknown. Similarly, the pre-implementation data are based on the erroneous 20 % minimum OAD in the AHU which was fixed during the implementation of RP-1747 to match the case study building's design. The impact of these limitations on assessing cooling savings are discussed further in Hobson et al. [23]. Data from the pre-implementation period where the SAT reset logic has been improved and the appropriate minimum OAD position was used are needed to remove the influence of these suboptimal sequences on the equipment loads before and after RP-1747 implementation.

5. Conclusions and recommendations

The implementation of RP-1747 DCV and its subsequent interactions with other BAS programming in a 626 m^2 institutional office building located in ASHRAE Climate Zone 6A – the first real-world implementation of RP-1747 DCV in an occupied building to date – reduced the chilled water load by an estimated 2 ± 6 %, the steam load by 30 ± 2 %, and the supply fan energy by 36 ± 26 % during operating hours, compared to predicted values for the same period had RP-1747 not been implemented. The magnitudes of these savings were in line with the simulation-based studies conducted by O'Neill et al. [18] (i.e., 29 % annual savings for a building located in ASHRAE Climate Zone 5A) and Hobson et al. [52] (i.e., 31 % annual savings for a building with the same sensor configuration in ASHRAE Climate Zone 6A). During the course of the RP-1747 implementation, suboptimal sequences of operation were noted and corrected in the building, namely the SAT reset logic and minimum OAD position used by the AHU in the heating and mechanical cooling modes.

The addition of a probabilistic temperature setback in zones further increased the predicted steam load reduction to 37 ± 2 %. These savings were largely attributable to zones entering occupied-standby mode,

where ventilation could be outright eliminated to these unoccupied zones. The probabilistic setback ensured that a higher proportion of timesteps during operating hours were spent at the minimum zone primary airflow setpoint, with an increase from 54.5 % of timesteps to 72.9 % of timesteps. However, the occupancy prediction logic used in the probabilistic setback had lower accuracy than expected, with a false negative rate of 10.5 % compared to the targeted 5 %. Despite the lower overall system primary airflow during the implementation of RP-1747, the zone outdoor airflow rates increased in the building after the implementation, thus improving the per person delivery rates of ventilation to occupied zones. The energy savings and ventilation improvements were achieved without impacting the building's ability to maintain ZAT setpoints during occupied hours, even with the higher-than-expected false negative rate of the probabilistic setback. Further data on the impact of the implementation on occupants' thermal comfort and perceived IAQ were sought via an electronic survey, however, occupants of the case study building had insufficient participation for these insights to be derived. The lack of responses was likely due to a handful of other more pressing issues faced by the occupants in the building, unrelated to the operation of the HVAC system, which drew their attention away from the electronic survey.

The supervisory control approach employed in this study helped to address the ballooning complexity of ASHRAE Guideline 36 by allowing for the implementation of RP-1747 DCV and the probabilistic setback on a custom Python-based server with bi-directional access to the BAS, thus avoiding the laborious process of creating dozens of new programs and hundreds of new variables and trend logs within the BAS itself. This flexible and centralized approach reduced the potential for programming errors to go unnoticed on individual controllers, lessened the complexity and burden on facilities personnel, and provided them with a quick and robust method for enabling or disabling the experimental controls as needed. The potential of this platform to employ advanced controls and analytics beyond the capabilities of typical commercial BASs has [39] and will continue to be explored in future work.

While RP-1747 CO_2-based DCV is poised for adoption by forthcoming versions of ASHRAE Guideline 36, future work should consider the use of non-CO_2-based sources for determining the person fraction of the ventilation demand in buildings where sufficient controls and sensing infrastructure are available. Cloud- or server-based supervisory controls approaches are likely to play a key role as they can enable a number of established occupant-count estimation methodologies that depend on data from disparate or non-BAS-integrated sources, and have sufficient computational power to employ advanced statistical or ML techniques to generate insights into occupancy patterns by leveraging existing data sources – without the need for a dense network of dedicated occupant-counting technologies.

CRediT authorship contribution statement

Brodie W. Hobson: Writing – original draft, Software, Methodology, Formal analysis, Conceptualization. **Andre A. Markus:** Writing – review & editing, Software, Methodology, Data curation. **Jayson Bursill:** Writing – review & editing, Software, Methodology. **H. Burak Gunay:** Writing – review & editing, Supervision, Resources, Methodology, Funding acquisition, Conceptualization. **Darwish Darwazeh:** Writing – review & editing, Supervision, Resources, Funding acquisition. **Zheng O'Neill:** Writing – review & editing, Methodology.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

This work was supported by a research contract with the National Research Council Canada under Contract No. 996635. The authors thank Scott Macdonald, Gavin Symonds, and their team at Carleton University's Facilities Management and Planning for their support. The authors acknowledge the RP-1547 and RP-1747 team members for their significant research contributions to the field. The authors also acknowledge the IEA EBC Annex 79 researchers, as we have greatly benefitted from the associated discussions.

Data availability

Data will be made available on request.

References

[1] W. O'Brien, A. Wagner, M. Schweiker, A. Mahdavi, J. Day, M.B. Kjærgaard, S. Carlucci, B. Dong, F. Tahmasebi, D. Yan, T. Hong, H.B. Gunay, Z. Nagy, C. Miller, C. Berger, Introducing IEA EBC Annex 79: Key challenges and opportunities in the field of occupant-centric building design and operation, Build. Environ. (2020) 106738, https://doi.org/10.1016/j.buildenv.2020.106738.

[2] H.B. Gunay, B.W. Hobson, M.M. Ouf, Z. Nagy, C. Miller, Design of sequences of operation for occupant-centric controls, in: W. O'Brien, F. Tahmasebi (Eds.), Occupant-Centric Simul.-Aided Build. Des. Theory Appl. Case Stud., 1st ed., Routledge, 2023: pp. 312–346.

[3] H.B. Gunay, Z. Shi, A. Ashouri, G.R. Newsham, Development of a clustering-based morning start time estimation algorithm for space heating and cooling, in: Proc. 6th ACM Int. Conf. Syst. Energy-Effic. Build. Cities Transp., Association for Computing Machinery, 2019: pp. 297–305. https://doi.org/10.1145/3360322.3360840.

[4] H.B. Gunay, G.R. Newsham, A. Ashouri, I. Wilton, Deriving sequences of operation for air handling units through building performance optimization, J. Build. Perform. Simul. 13 (2020) 501–515, https://doi.org/10.1080/19401493.2020.1793221.

[5] Y. Benezeth, H. Laurent, B. Emile, C. Rosenberger, Towards a sensor for detecting human presence and characterizing activity, Energy Build. 43 (2011) 305–314, https://doi.org/10.1016/j.enbuild.2010.09.014.

[6] W. Zhang, C. Zhang, Maximize thermal comfort in open-plan offices by occupant-oriented control based on individual thermal profile, ASHRAE Trans. (2019) 167–175.

[7] T. Labeodan, K. Aduda, W. Zeiler, F. Hoving, Experimental evaluation of the performance of chair sensors in an office space for occupancy detection and occupancy-driven control, Energy Build. 111 (2016) 195–206, https://doi.org/10.1016/j.enbuild.2015.11.054.

[8] J.Y. Park, M.M. Ouf, H.B. Gunay, Y. Peng, W. O'Brien, M.B. Kjærgaard, Z. Nagy, A critical review of field implementations of occupant-centric building controls, Build. Environ. 165 (2019) 106351, https://doi.org/10.1016/j.buildenv.2019.106351.

[9] W.J. Fisk, A.T.D. Almeida, Sensor-based demand-controlled ventilation: a review, Energy Build. 29 (1998) 35–45, https://doi.org/10.1016/S0378-7788(98)00029-2.

[10] J. Goins, M. Moezzi, Linking occupant complaints to building performance, Build. Res. Inf. 41 (2013) 361–372, https://doi.org/10.1080/09613218.2013.763714.

[11] W. Shen, G.R. Newsham, H.B. Gunay, Leveraging existing occupancy-related data for optimal control of commercial office buildings: a review, Adv. Eng. Inform. 33 (2017) 230–242, https://doi.org/10.1016/j.aei.2016.12.008.

[12] American Society of Heating Refrigerating and Air-Conditioning Engineers, ASHRAE Standard 62.1-2022: Ventilation and Acceptable Indoor Air Quality, American Society of Heating Refrigerating and Air-Conditioning Engineers, Inc., 2022.

[13] S.H. Kim, H.J. Moon, Case study of an advanced integrated comfort control algorithm with cooling, ventilation, and humidification systems based on occupancy status, Build. Environ. 133 (2018) 246–264, https://doi.org/10.1016/j.buildenv.2017.12.010.

[14] American Society of Heating Refrigerating and Air-Conditioning Engineers, ASHRAE Guideline 36-2018: High-Performance Sequences of Operation for HVAC Systems, American Society of Heating Refrigerating and Air-Conditioning Engineers, Inc., 2018.

[15] American Society of Heating Refrigerating and Air-Conditioning Engineers, ASHRAE Guideline 36-2021: High-Performance Sequences of Operation for HVAC Systems, American Society of Heating Refrigerating and Air-Conditioning Engineers, Inc., 2021.

[16] American Society of Heating Refrigerating and Air-Conditioning Engineers, ASHRAE Standard 62.1-2019: Ventilation for Acceptable Indoor Air Quality, American Society of Heating Refrigerating and Air-Conditioning Engineers, Inc., 2019.

[17] J. Lau, X. Lin, G. Yuill, RP-1547 - CO_2-based demand controlled ventilation for multiple zone HVAC systems, (2013).

[18] Z.D. O'Neill, Y. Li, H.C. Cheng, X. Zhou, S.T. Taylor, Energy savings and ventilation performance from CO_2-based demand controlled ventilation: simulation results

from ASHRAE RP-1747, Sci. Technol. Built Environ. 26 (2020) 257–281, https://doi.org/10.1080/23744731.2019.1620575.
[19] Z. O'Neill, Y. Li, X. Zhou, S. Taylor, H. Cheng, RP-1747 - Implementation of RP-1547 CO_2-based demand controlled ventilation for multiple zone HVAC systems in direct digital control systems, (2017) 1–181.
[20] M. Hydeman, B. Eubanks, RP-1455 - Advanced control sequences for HVAC Systems: Phase I Air Distribution and Terminal Systems, (2014).
[21] S.T. Taylor, Resetting setpoints using trim & respond logic, ASHRAE J. 47 (2015) 52–57.
[22] B.W. Hobson, H.B. Gunay, S. Shillinglaw, Exploring CO_2 sensor grid configurations and the implications for demand-controlled ventilation, in: Proc. 5th Int. Conf. Build. Energy Environ., Springer, Montreal, CA, 2022: pp. 1855–1864. https://doi.org/10.1007/978-981-19-9822-5_194.
[23] B.W. Hobson, A.A. Markus, J. Bursill, H.B. Gunay, F. Rizvi, Z. O'Neill, Implementation of RP-1747 demand-controlled ventilation in an office building: Preliminary results from the cooling season, in: ASHRAE Trans., American Society for Heating, Refrigeration and Air Conditioning Engineerings Inc, Chicago, USA, 2024, pp. 77–85.
[24] K. Zhang, D. Blum, H. Cheng, G. Paliaga, M. Wetter, J. Granderson, Estimating ASHRAE guideline 36 energy savings for multi-zone variable air volume systems using Spawn of EnergyPlus, J. Build. Perform. Simul. 15 (2022) 215–236, https://doi.org/10.1080/19401493.2021.2021286.
[25] M. Wetter, P. Ehrlich, A. Gautier, M. Grahovac, P. Haves, J. Hu, A. Prakash, D. Robin, K. Zhang, OpenBuildingControl: digitizing the control delivery from building energy modeling to specification, implementation and formal verification, Energy 238 (2022), https://doi.org/10.1016/j.energy.2021.121501.
[26] M.M. Ardehali, T.F. Smith, J.M. House, C.J. Klaassen, Building energy use and control problems: an assessment of case studies, ASHRAE Trans. 109 (2003) 111–121.
[27] S. Katipamula, R.M. Underhill, N. Fernandez, W. Kim, R.G. Lutes, D. Taasevigen, Prevalence of typical operational problems and energy savings opportunities in U. S. commercial buildings, Energy Build. 253 (2021), https://doi.org/10.1016/j.enbuild.2021.111544.
[28] American Society of Heating Refrigerating and Air-Conditioning Engineers, Chapter 43: Supervisory control strategies and optimization, in: ASHRAE Handb. - HVAC Appl., American Society for Heating, Refrigeration and Air Conditioning Engineers, Inc., Atlanta, USA, 2023: p. 43.1-43.50.
[29] S. Wang, Z. Ma, Supervisory and optimal control of building HVAC systems: a review, HVAC R Res. 14 (2008) 3–32, https://doi.org/10.1080/10789669.2008.10390991.
[30] C. Anuntasethakul, D. Banjerdpongchai, Design of supervisory model predictive control for building HVAC system with consideration of peak-load shaving and thermal comfort, IEEE Access 9 (2021), https://doi.org/10.1109/ACCESS.2021.3065083.
[31] A. Aswani, N. Master, J. Taneja, D. Culler, C. Tomlin, Reducing transient and steady state electricity consumption in HVAC using learning-based model-predictive control, Proc. IEEE 100 (2012) 240–253, https://doi.org/10.1109/JPROC.2011.2161242.
[32] A. Afram, F. Janabi-Sharifi, Theory and applications of HVAC control systems – A review of model predictive control (MPC), Build. Environ. 72 (2014) 343–355, https://doi.org/10.1016/j.buildenv.2013.11.016.
[33] B. Dong, K.P. Lam, A real-time model predictive control for building heating and cooling systems based on the occupancy behavior pattern detection and local weather forecasting, Build. Simul. 7 (2014) 89–106, https://doi.org/10.1007/s12273-013-0142-7.
[34] J. Bursill, An approach to data-driven sensing and predictive supervisory control for commercial buildings with in-situ evaluation, Carleton University, 2020. PhD thesis.
[35] Canadian Commission on Building and Fire Codes, National Energy Code of Canada for Buildings: 2020, Natural Research Council of Canada, 2020.
[36] A.A. Markus, B.W. Hobson, H.B. Gunay, S. Bucking, A framework for a multi-source, data-driven building energy management toolkit, Energy Build. 250 (2021) 111255, https://doi.org/10.1016/j.enbuild.2021.111255.
[37] Air Movement and Control Association, AMCA Publication 502-06 (R2009): Damper Application Manual for Heating, Ventilating, and Air Conditioning, Air Movement and Control Association, Inc., 2006.
[38] American Society of Heating Refrigerating and Air-Conditioning Engineers, 2019 ASHRAE Handbook – HVAC Application (SI edition), American Society of Heating Refrigerating and Air-Conditioning Engineers, Inc., 2019.
[39] A.A. Markus, B.W. Hobson, J. Bursill, H.B. Gunay, Deployment of real-time building automation system-integrated inverse-model-based fault detection and diagnostics algorithms, Sci. Technol. Built Environ. 30 (2024) 134–152, https://doi.org/10.1080/23744731.2023.2290976.
[40] ASHRAE, ASHRAE Standard 135-2020: BACnet™ - A Data Communication Protocol for Building Automation and Control Networks, (2020).
[41] The pandas development team, pandas-dev/pandas: Pandas, (2020). https://doi.org/10.5281/zenodo.3509134.
[42] W. McKinney, Data Structures for Statistical Computing in Python, in: S. van der Walt, J. Millman (Eds.), Proc. 9th Python Sci. Conf., 2010: pp. 56–61. https://doi.org/10.25080/Majora-92bf1922-00a.
[43] C.R. Harris, K.J. Millman, S.J. van der Walt, R. Gommers, P. Virtanen, D. Cournapeau, E. Wieser, J. Taylor, S. Berg, N.J. Smith, R. Kern, M. Picus, S. Hoyer, M.H. van Kerkwijk, M. Brett, A. Haldane, J.F. del Río, M. Wiebe, P. Peterson, P. Gérard-Marchant, K. Sheppard, T. Reddy, W. Weckesser, H. Abbasi, C. Gohlke, T.E. Oliphant, Array programming with NumPy, Nature 585 (2020) 357–362, https://doi.org/10.1038/s41586-020-2649-2.
[44] D.R. Karger, D. Quan, Haystack: a user interface for creating, browsing, and organizing arbitrary semistructured information, in: CHI 04 Ext. Abstr. Hum. Factors Comput. Syst., Association for Computing Machinery, New York, NY, USA, 2004: pp. 777–778. https://doi.org/10.1145/985921.985931.
[45] B. Balaji, A.A. Bhattacharya, G. Fierro, J. Gao, J. Gluck, D. Hong, A. Johansen, J. Koh, J. Ploennigs, Y. Agarwal, M. Berges, D. Culler, R. Gupta, M.B. Kjærgaard, M. Srivastava, K. Whitehouse, Brick: Towards a Unified Metadata Schema For Buildings, in: Proc. 3rd ACM Int. Conf. Syst. Energy-Effic. Built Environ., Association for Computing Machinery, 2016: pp. 41–50. https://doi.org/10.1145/2993422.2993577.
[46] ASHRAE's BACnet Committee, ASHRAE 223P - Designation and Classification of Semantic Tags for Building Data, ASHRAEs BACnet Comm. Proj. Haystack Brick Schema Collab. Provide Unified Data Semantic Model. Solut. (2018). http://web.archive.org/web/20181223045430/https://www.ashrae.org/about/news/2018/ashrae-s-bacnet-committee-project-haystack-and-brick-schema-collaborating-to-provide-unified-data-semantic-modeling-solution (accessed October 9, 2024).
[47] American Society of Heating Refrigerating and Air-Conditioning Engineers, ASHRAE Guideline 14-2014: Measurement of Energy, Demand, and Water Savings, American Society of Heating Refrigerating and Air-Conditioning Engineers, Inc., 2014.
[48] J. Berquist, M.M. Ouf, W. O'Brien, A method to conduct longitudinal studies on indoor environmental quality and perceived occupant comfort, Build. Environ. 150 (2019) 88–98, https://doi.org/10.1016/j.buildenv.2018.12.064.
[49] P. Jayathissa, M. Quintana, T. Sood, N. Nazarian, C. Miller, Is your clock-face cozie? A smartwatch methodology for the in-situ collection of occupant comfort data, J. Phys. Conf. Ser. 1343 (2019) 012145, https://doi.org/10.1088/1742-6596/1343/1/012145.
[50] F. Tartarini, M. Frei, S. Schiavon, Y.X. Chua, C. Miller, Cozie Apple: an iOS mobile and smartwatch application for environmental quality satisfaction and physiological data collection, J. Phys. Conf. Ser. 2600 (2023) 142003, https://doi.org/10.1088/1742-6596/2600/14/142003.
[51] J.K. Day, W. O'Brien, Oh behave! Survey stories and lessons learned from building occupants in high-performance buildings, Energy Res. Soc. Sci. 31 (2017) 11–20, https://doi.org/10.1016/j.erss.2017.05.037.
[52] B.W. Hobson, A.A. Markus, H.B. Gunay, F. Rizvi, Minimum sensor grid density and configuration to enable CO_2-based demand-controlled ventilation in an office building, Energy Build. 298 (2023) 113536, https://doi.org/10.1016/j.enbuild.2023.113536.
[53] X. Lu, Z. Pang, Y. Fu, Z. O'Neill, Advances in research and applications of CO_2-based demand-controlled ventilation in commercial buildings: a critical review of control strategies and performance evaluation, Build. Environ. 223 (2022), https://doi.org/10.1016/j.buildenv.2022.109455.
[54] J. Berquist, Z. Xiong, H.B. Gunay, M. Vuotari, Investigation of the accuracy of BAS-grade CO_2 sensors for measuring infiltration rates, J. Build. Eng. 80 (2023), https://doi.org/10.1016/j.jobe.2023.108064.
[55] Z. Xiong, J. Berquist, H.B. Gunay, C.A. Cruickshank, An inquiry into the use of indoor CO_2 and humidity ratio trend data with inverse modelling to estimate air infiltration, Build. Environ. 206 (2021), https://doi.org/10.1016/j.buildenv.2021.108365.
[56] X. Lu, Z. Pang, Y. Fu, Z. O'Neill, The nexus of the indoor CO_2 concentration and ventilation demands underlying CO_2-based demand-controlled ventilation in commercial buildings: a critical review, Build. Environ. 218 (2022), https://doi.org/10.1016/j.buildenv.2022.109116.
[57] X. Lu, Z. O'Neill, Y. Li, F. Niu, A novel simulation-based framework for sensor error impact analysis in smart building systems: a case study for a demand-controlled ventilation system, Appl. Energy 263 (2020), https://doi.org/10.1016/j.apenergy.2020.114638.
[58] American Society of Heating Refrigerating and Air-Conditioning Engineers, ASHRAE Standard 189.1-2022: Design of High-Performance, Green Buildings Except Low-Rise Residential, American Society of Heating Refrigerating and Air-Conditioning Engineers, Inc., 2022.
[59] California Energy Commission, 2019 Building Energy Efficiency Standards for Residential and Nonresidential Buildings, California Natural Resources Agency, 2018.
[60] B. Balaji, J. Xu, A. Nwokafor, R. Gupta, Y. Agarwal, Sentinel: occupancy based HVAC actuation using existing Wi-Fi infrastructure within commercial buildings, in: Proc. 11th ACM Conf. Embed. Networked Sens. Syst., Association for Computing Machinery, 2013: pp. 1–14. https://doi.org/10.1145/2517351.2517370.
[61] J.Y. Park, T. Dougherty, Z. Nagy, A Bluetooth based occupancy detection for buildings, in: Proc. SimBuild 2022, 2018. https://publications.ibpsa.org/conference/paper/?id=simbuild2018_C111.
[62] W. Wang, J. Chen, T. Hong, N. Zhu, Occupancy prediction through Markov-based feedback recurrent neural network (M-FRNN) algorithm with Wi-Fi probe technology, Build. Environ. 138 (2018) 160–170, https://doi.org/10.1016/j.buildenv.2018.04.034.
[63] S. Salimi, A. Hammad, Critical review and research roadmap of office building energy management based on occupancy monitoring, Energy Build. 182 (2019) 214–241, https://doi.org/10.1016/j.enbuild.2018.10.007.
[64] B.W. Hobson, D. Lowcay, H.B. Gunay, A. Ashouri, G.R. Newsham, Opportunistic occupancy-count estimation using sensor fusion: a case study, Build. Environ. 159 (2019) 106154, https://doi.org/10.1016/j.buildenv.2019.05.032.
[65] B.W. Hobson, T. Abuimara, A. Ashouri, H.B. Gunay, Disaggregating building-level occupancy into zone-level occupant counts using sensor fusion, in: Proc. 5th Int.

Conf. Build. Energy Environ., Springer, Montreal, CA, 2022: pp. 1913–1924. https://doi.org/10.1007/978-981-19-9822-5_201.

[66] Z. Chen, M.K. Masood, Y.C. Soh, A fusion framework for occupancy estimation in office buildings based on environmental sensor data, Energy Build. 133 (2016) 790–798, https://doi.org/10.1016/j.enbuild.2016.10.030.

[67] J.P. Roselyn, R.A. Uthra, A. Raj, D. Devaraj, P. Bharadwaj, S.V.D.K. Kaki, Development and implementation of novel sensor fusion algorithm for occupancy detection and automation in energy efficient buildings, Sustain. Cities Soc. 44 (2019) 85–98, https://doi.org/10.1016/j.scs.2018.09.031.