Identifying the Minimum CO_2 Sensing Infrastructure Required to Enable Demand-Controlled Ventilation: A Preliminary Analysis

Brodie W. Hobson[1], H. Burak Gunay[1], Scott Shillinglaw[2]
[1]Department of Civil and Environmental Engineering, Carleton University, Ottawa, Canada
[2]Construction Research Centre, National Research Council Canada, Ottawa, Canada

Abstract

Maintaining CO_2 sensors for demand-controlled ventilation (DCV) requires frequent and costly recalibration. Therefore, identifying the minimum CO_2 sensing infrastructure needed to enable DCV can lower costs while improving indoor air quality (IAQ). This study simulates the floor of a building under ten stochastic occupancy scenarios to help inform how sensor placement should be prioritized. It was found that smaller multi-occupant zones disproportionately experienced high CO_2 concentrations, while core zones and single-occupant offices with lower ventilation to seating capacity ratios also experienced elevated concentrations. This study highlights the benefits of pairing stochastic occupancy scenarios with building simulation to inform sensor deployment.

Introduction

Carbon dioxide- (CO_2) based demand-controlled ventilation (DCV) is an example of occupant-centric control (OCC), specifically occupancy-centric control (i.e., controls which adjust the indoor environment based on the number of occupants present) (Park *et al.*, 2019). The occupancy data used for these controls are implicit (Melfi *et al.*, 2011); rather than measuring occupancy directly with proprietary sensors, the impact occupants have on CO_2 levels due to exhalation is measured, and ventilation is increased to reduce the concentration. While CO_2 concentrations are used for control, the resulting ventilation provided to spaces also ensures that bio-effluents and volatile organic compounds (VOCs) (Fisk and De Almeida, 1998), and – as has been the renewed focus of many studies given the COVID-19 pandemic – infectious aerosols are controlled in a space. CO_2 concentration has been shown to be a good proxy for the viral load in zones (Zivelonghi and Lai, 2021); CO_2-based DCV will therefore play a role in the safe return to work during and/or after the COVID-19 pandemic. Indeed, the pandemic will likely have a long-lasting impact on the occupancy patterns of office buildings (Deng, Morissette and Messacar, 2020), which are likely to have even more flexible working hours and sporadic/partial occupancies that began trending before the advent of COVID-19 (e.g., (U.S. DOT, 2000; Golden, 2001)). The reduction in ventilation provided to under- or unoccupied zones via CO_2-based DCV has been shown in numerous studies to reduce energy use (e.g., (Mysen, Rydock and Tjelflaat, 2003; Mysen *et al.*, 2005; Sun, Wang and Ma, 2011; Zhang *et al.*, 2013)), resulting in reduced energy costs and greenhouse gas (GHG) emissions with the magnitude dependant on the climate and energy source.

In order to provide adequate indoor air quality (IAQ), the breathing zone outdoor airflow (V_{bz}) can be increased by a proportional-integral-derivative (PID) controller to change air damper positions at certain CO_2 concentrations (Liu, Dasu and Zhang, 2012); ASHRAE 62.1 (2019) prescribes no more than 600 ppm above atmospheric CO_2 levels (i.e., ~1000 ppm), while Charles and Veitch (2002) found that indoor CO_2 concentrations at or below a maximum of 800 ppm were associated with improved human-factor outcomes. For the PID controller to modulate the damper, CO_2 readings must be taken. However, the CO_2 concentrations of a space are non-homogenous due to diffusion and mixing mechanics, occupant distribution and exhalation rates, and other transient conditions. Therefore, the location where CO_2 is monitored may not adequately reflect the prevailing IAQ conditions occupants are subjected to.

Consider the banquet hall shown in Figure 1(a). Ventilation is provided by a dedicated air handling unit (AHU) which provides air through supply grilles at the ceiling level. The return air CO_2 concentration is monitored in the AHU and used by the PID controller to control the outdoor air damper position. Five other CO_2 sensors were placed in the space to monitor localized CO_2 concentrations. On the day of an early morning event, see Figure 1(b), the peak CO_2 concentration of the AHU return air was in the high 500 ppm range. However, CO_2 concentrations at sensor 3 were as high as ~875 ppm, and in the mid to high 600 ppm range for sensors 1, 2, and 5. Therefore, the only sensor that reflected similar readings to the AHU was sensor 4.

By using the return air CO_2 concentration for DCV, the outdoor airflow did not increase adequately to address locally high CO_2 concentrations where, implicitly, the highest number of occupants are located. This problem not only has implications within zones, but across floorplates and entire buildings, as the distribution of occupants is non-homogenous. This non-homogeneity is contrary to the

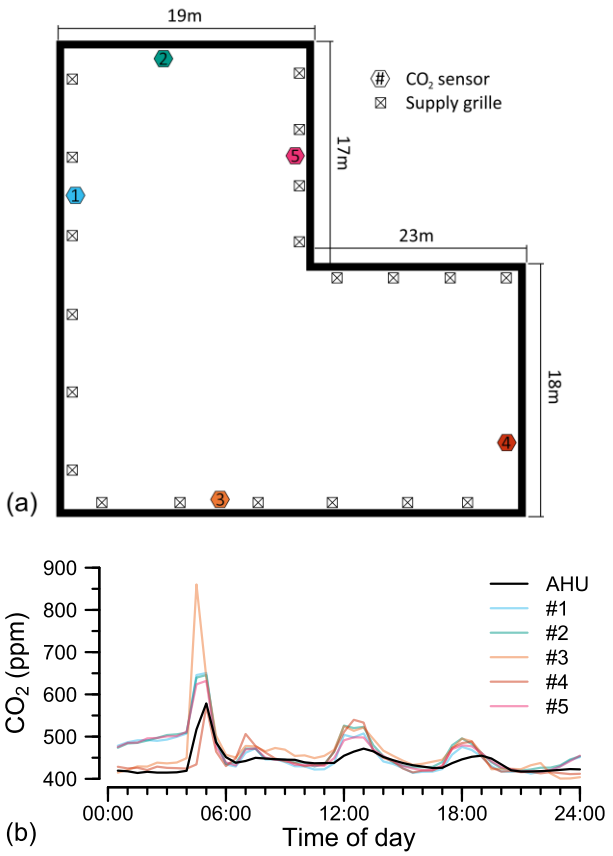

Figure 1: (a) Banquet hall configuration and (b) daily CO_2 concentrations.

default occupancy assumptions used in ASHRAE 62.1 (2019). Therefore, using the CO_2 concentration at the AHU-level alone was not insufficient to enable effective DCV in the case of this banquet hall under realistic occupancy conditions. The question remains: how much CO_2 sensing infrastructure is required to enable effective DCV? Especially in existing construction, it may be cost-prohibitive to install a dense CO_2 sensor grid like in Figure 1(a) in every zone across an entire floorplate. While the capital cost of these sensors and their installation is non-trivial, these sensors also require recalibration every five years (ASHRAE, 2019); this demands an exponential amount of time and effort from an operations team (or a contracted third-party) as the density of the sensing infrastructure increases.

Therefore, the aim of this study is to explore the minimum CO_2 sensing infrastructure needed to enable DCV for a floor of an office building via a simulation-based investigation. The two competing motivations for this analysis are the need to increase CO_2 sensing beyond the system level to allow for effective DCV, was also ensuring the sensor grid density does not become too costly or burdensome for operations personnel. These preliminary findings may be used to explore which zones and circumstances yield the highest CO_2 concentrations and inform how CO_2 sensor grids should be configured so that adequate IAQ can be provided with minimal sensing infrastructure and capital cost by focusing on these critical zones.

Methodology

An EnergyPlus 9.5 model of a 662 m^2 floor of an academic office building located in ASHRAE Climate Zone 6A was developed for the purpose of this study. The model was simulated ten times with stochastic occupant counts and schedules in each zone. These schedules were derived from sensor data from 44 zones on the same campus serving similar occupancy types. Hourly CO_2 concentrations from each zone were extracted from the simulation results. For each of the ten simulations, k-means clustering was performed on these CO_2 data to identify zones that exhibited similar CO_2 trends and/or critical zones with outlying CO_2 levels.

Building and energy model

The study building is currently under construction; the floor studied will be highly instrumented to support the use of it as a 'living lab'. Tender drawings were used to extract information related to the floor's geometry, envelope properties, and heating, ventilation, and air conditioning (HVAC) system. The model geometry was approximated such that zones served by the same variable air volume (VAV) terminal unit and with thermally similar exterior conditions were treated as a single thermal zone. Note that this floor is the top floor of the building with an identical floorplate below, and the south and east façades abut an existing structure; the perimeter surfaces of the zones along these façades, as well as the floor, were treated as adiabatic. An overview of the model geometry is provided in Figure 2. Note that the floor height is 3.96 m. This geometry was created in SketchUp 2017 using the Euclid 0.9.4.3 plug-in and imported into EnergyPlus.

The floor is served by a dedicated AHU, which supplies air for ventilation and conditioning purposes. Zone-level airflow is controlled by a single-ducted VAV terminal unit in each of the thermal zones outlined in Figure 2. Perimeter hydronic units are used to provide additional heating to the exterior perimeter zones (i.e., the zones along the north and west façades). Heating is available in the building from early September until mid-May, at which point seasonal switchover to cooling occurs. The AHU schedule for the building during workdays and weekends is 5 am to 7 pm and 6 am to 7 pm, respectively, with the unit not operating during holidays. The zone air temperature setpoint was set to 22°C during operating hours, with a setback to 17°C and 27°C during unoccupied hours in the heating and cooling seasons, respectively. The lighting power density for each zone was set to 8 W/m^2 (Bennet and O'Brien, 2017); as lighting was not the focus of this study, the use of artificial lighting was simplified such that occupants turned on the lights upon arrival in each zone, and lighting was switched off within one hour of occupants departing (i.e., a common control scheme for motion detector-based lighting

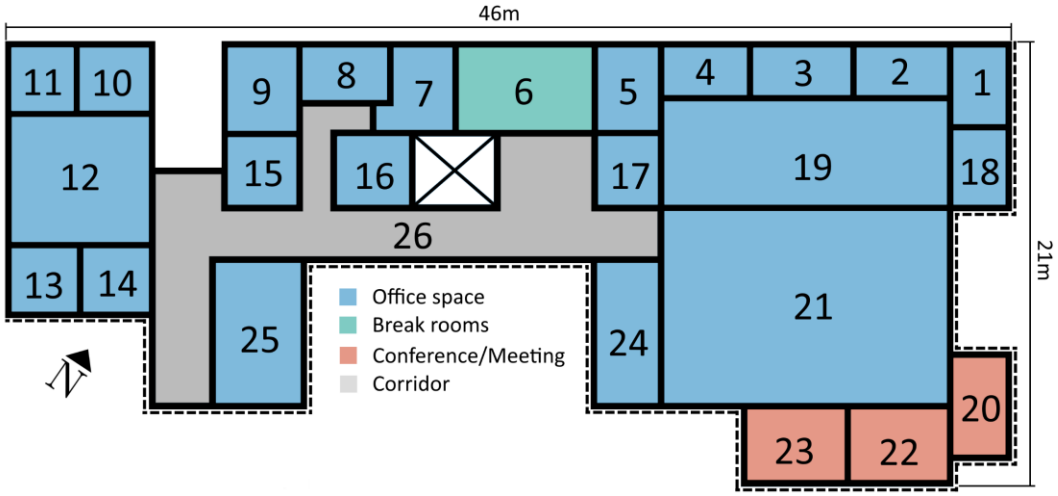

Figure 2: Overview of energy model geometry. Note that the dashed line represents an adiabatic boundary condition.

Table 1: Zone seating and ventilation capacity.

Zone	Seating (persons)	Ventilation* (persons)	Ventilation Seating
1	1	2.6	2.6
2	1	4.2	4.2
3	1	4.2	4.2
4	1	4.2	4.2
5	1	2.6	2.6
6	4	22.9	5.7
7	1	2.6	2.6
8	1	4.2	4.2
9	1	2.6	2.6
10	1	4.2	4.2
11	1	5.2	5.2
12	4	10.3	2.6
13	1	4.2	4.2
14	1	1.6	1.6
15	1	1.6	1.6
16	1	2.6	2.6
17	1	1.6	1.6
18	1	1.6	1.6
19	3	7.3	2.4
20	6	6.4	1.1
21	8	10.5	1.3
22	4	6.4	1.6
23	4	6.4	1.6
24	2	3.1	1.6
25	2	3.1	1.6
26	0	NA	NA

automation). The plug-in equipment electricity use was set to 100 W per person based on zones' typical occupancy, which was reduced to 75 W per person after-hours on workdays and 50 W per person during weekends (Gunay *et al.*, 2016).

Occupancy scenarios

Occupancy in each zone was altered by multiplying the zone's default occupancy by a schedule object containing hourly occupancy factors (i.e., 0 for unoccupied, <1 for partial occupancy, 1 for fully occupied, and >1 for over-occupied). For each of the ten scenarios, each of the 26 zones was randomly assigned an hourly occupancy schedule containing these occupancy factors; note that duplicates were not allowed within the same scenario (i.e., each of the 26 schedules within each scenario is unique). The default occupancy of each zone was based on the installed seating capacities provided in the tender drawings, see Table 1. The ventilation capacity in persons for each VAV is also provided for comparison.

The occupancy factor schedules used for the single-occupant offices were based on concurrent hourly-measured motion detector data from 35 single-occupant offices located on the same campus from midnight on May 1st, 2019, until 11 pm on December 31st, 2019. These data were binary in nature (i.e., occupancy factors of 0 for unoccupied if no motion, or one if motion was observed at any point during a timestep). Occupancy factor schedules used for multi-occupant zones (n = 9) were based on concurrent hourly occupant-count estimates developed in a previous study for ten multi-occupant zones located on the same campus for the same period as the single-occupant offices. Briefly, these estimates were based on a sensor fusion model that included CO_2, motion detector, and Wi-Fi access point data that was trained using manually collected ground truth data (R^2 = 83%) (Hobson *et al.*, 2019). These occupant counts were divided by the seating capacity of their respective zones to determine their occupancy factor schedules. Unlike the single-occupant offices, the occupancy factors for the multi-occupant zones ranged from 0 when unoccupied to 4 (i.e., four-times the intended occupancy of the zone), capturing situations such as brief periods of over-occupancy (e.g., an excessive amount of occupants in a small conference room for a meeting) or

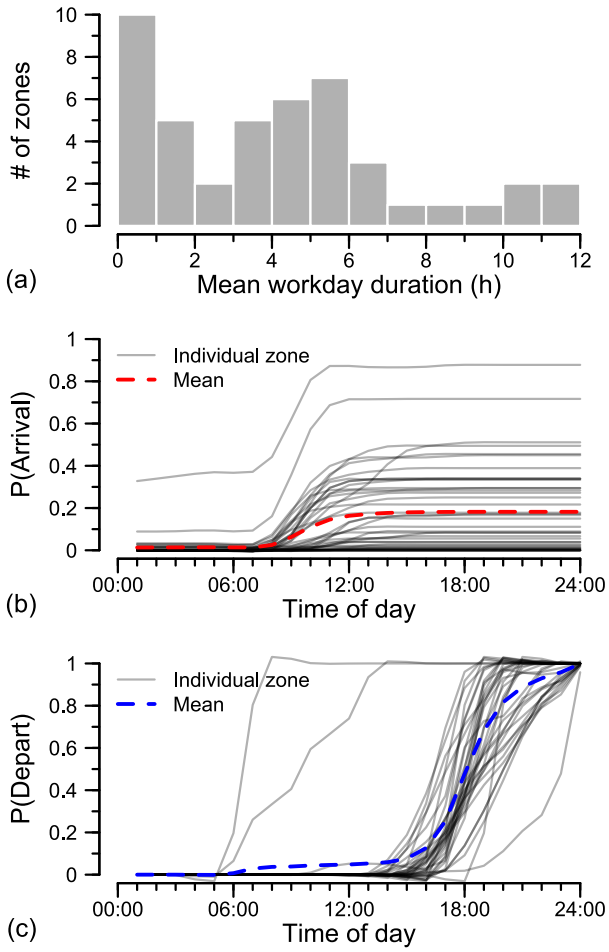

Figure 3: (a) Workday durations, probability of (b) first arrival and (c) last departure for occupancy data.

prolonged periods of over-occupancy (e.g., chronic use of a zone in an unintended fashion (Abuimara *et al.*, 2021)). These data reflected a diverse range of occupancy patterns, see Figure 3.

Note that the ten schedules selected do not capture the full breadth of occupancy patterns in the measured data (i.e., there are billions of possible combinations of data from the 44 zones) or broader occupant behaviours in general; they are by no means exhaustive. These select scenarios are meant to serve as a cursory example of how different occupancy patterns within a family of zones can impact CO_2 readings and the implications this has on sensor grid configuration. It is assumed that occupancy of the floor of the building considered in this study will be similar to the occupancy data used due to the buildings being co-located (i.e., on the same campus), the similar occupancy types between the buildings, and the demographics (i.e., the simulated floor belongs to the same faculty as the offices used for the occupancy schedules).

Once these stochastic occupancy factor schedules were selected for the ten simulation scenarios, the simulations were run for May 2019 to December 2019 period, where occupancy data were available. Note that actual meteorological year (AMY) weather data for the same period were used. The zone-level CO_2 concentrations, as well as the return air CO_2 concentrations at the AHU, were then extracted from the simulation results. It should be noted that EnergyPlus assumes a CO_2 generation rate of 0.0046 L/s per person corresponding to occupants with an average metabolic rate of 1.2 met (i.e., light office work) (United States Department of Energy, 2020). Additionally, EnergyPlus uses the well-mixed assumption when calculating and reporting CO_2 concentrations. These default assumptions were accepted for each simulation.

K-means clustering

The time-series CO_2 data for each simulation scenario underwent unity-based normalization. K-means clustering was selected as the preferred clustering algorithm as previous studies with time-series data have shown that the Caliñski-Harabasz (CH) index is higher with this method compared to agglomerative hierarchical clustering techniques (e.g., (Hobson *et al.*, 2020; Nweye and Nagy, 2021)). Clustering was performed based on a subset of data which contained CO_2 concentrations during occupied hours; afterhours, weekends, and holidays were excluded from these data. As the number of clusters for each scenario is not inherently known, the CH index for *k* clusters between one and five was computed for each simulation scenario in R using the *clues* (Qiu *et al.*, 2016) package. K-means clustering with the appropriate number of clusters was then performed in R.

Results and discussion

The cluster membership for each of the zones for the ten simulation scenarios are provided in Figure 4; zones consistently grouped in the same cluster have similar trends in CO_2 concentrations (i.e., peak and shape) subject to the stochastic occupancy scenarios. Thus, monitoring of CO_2 in a minimum of one zone belonging to each cluster is necessary to capture the diversity of CO_2 trends observed across the floor of the case study building. Note that all but one simulation scenario had two clusters; scenario five had three. Six of the ten simulation scenarios, including scenario five, had a cluster with a single outlying zone. Across simulation scenarios, zone 20 (i.e., a conference room) had outlying CO_2 patterns in four scenarios (e.g., C1 in Figure 5 (a)) and belonged to the smallest cluster in eight scenarios. In both cases, zone 20 belonged to the clusters with the highest peak CO_2 concentrations (e.g., C1 in Figure 5(a) and Figure 5(c)). This zone also had the lowest ventilation to seating capacity ratio out of the 26 zones studied, recall Table 1. Similarly, zone 22 (i.e., a meeting room) had outlying CO_2 patterns and belonged to the cluster with the highest peak CO_2 concentrations in two and six of the simulation scenarios, respectively (e.g., Figure 5(b)). Zone 23, an identical zone with the same seating capacity and ventilation rates, only experienced higher than average CO_2 concentrations once; this is because the occupancy factors

Figure 4: Cluster membership for each zone for scenarios (a) 1 through (j) 10.

for zone 23 happened to be lower than zone 22 (i.e., an average and maximum occupancy factor of 0.85 and 1.3 for zone 23, respectively, compared to 1.32 and 4 for zone 22, respectively). Both zones were tied for the third-lowest ventilation to seating capacity ratio, see Table 1. It should be noted that the peak CO_2 concentrations in zone 22 were typically lower than those in zone 20 due to this factor (i.e., out of the six scenarios where it belongs to the highest cluster, two of these have peak concentrations below the 800 ppm concentrations associated with improved human-factor outcomes (Charles and Veitch, 2002)). In fact, three

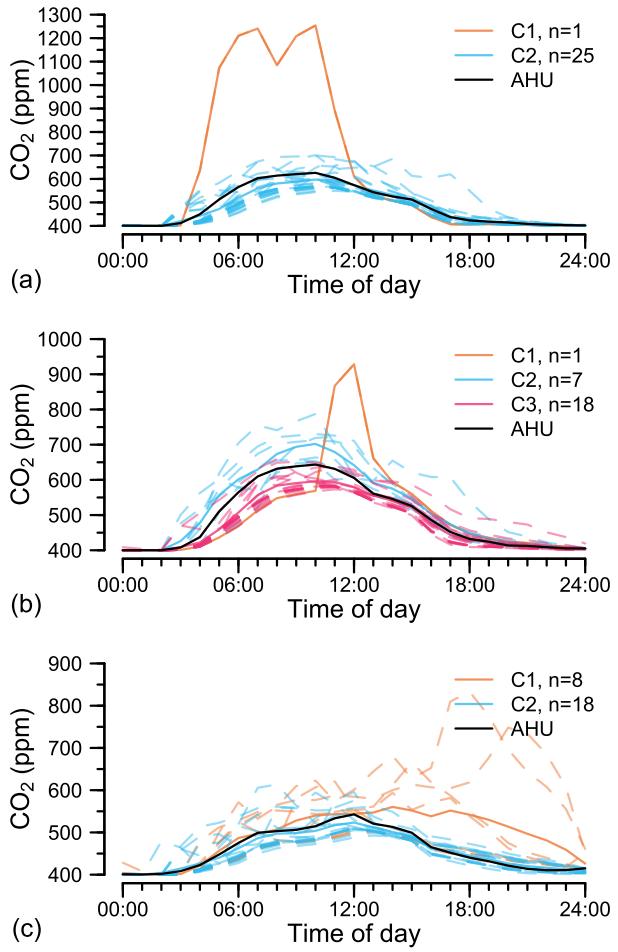

Figure 5: Example daily CO_2 concentrations for zones with respect to their cluster for (a) scenario 1, (b) scenario 2, and (c) scenario 9.

of the scenarios had peak CO_2 concentrations below 800 ppm across all zones (namely scenarios 3, 4, and 8).

Within the remaining seven scenarios where peak CO_2 concentrations routinely exceeded 800 ppm, 12% of zones (i.e., n = 21 out of 7 scenarios × 26 zones) belonged to clusters with the highest average CO_2 concentrations. Of this 12%, 52.3% of these zones (i.e., n = 11) were multi-occupant zones, despite multi-occupant zones only representing nine of the 26 zones studied (i.e., 36%). The remaining 47.7% of zones (i.e., n = 10) with consistently high peak CO_2 concentrations were, therefore, single-occupant zones, despite the fact that the occupancy factor schedules for these zones never exceeded one and the ventilation to seating capacity ratio for all zones was above one as well, see Table 2. Of these ten zones, eight of which were unique, four (i.e., 50%) had no exterior-facing surface area (i.e., zones 14, 15, 16, and 18). The four remaining zones that did have exterior-facing surface area had the lowest exterior-facing surface area to floor area ratios of the perimeter zones on the floor studied. This means that all single-occupant zones which experienced CO_2 concentrations exceeding 800 ppm had lower rates of outdoor air exchange due to decreased infiltration through the façade relative to their perimeter-facing counterparts.

Based on these preliminary results, it appears that zones with ventilation capacities that closely match their designed seating capacities intuitively have higher CO_2 concentrations. This is exacerbated when considering scenarios where the occupancy within these zones exceeds the design capacity. This was disproportionally true in multi-occupant zones, specifically for conference/meeting type occupancies (e.g., zone 20 and 22). However, multi-occupant office spaces (i.e., zones 12, 19, 21, 24, and 25), despite having the same diverse occupant schedules and occupancy factors, exhibited lower peak CO_2 concentrations compared to their conference/meeting type occupancy counterparts. This is in spite of the fact that zone 21 had the second lowest ventilation to seating capacity ratio and that these zones were predominately core spaces. These five zones are all larger than the conference/meeting room zones on this floorplate, which likely means that the accumulation of CO_2 would take longer under the same occupancy load and could account for this discrepancy. Considering these factors, rooms where ventilation to seating capacity ratios are close should be considered critical zones (i.e., below 2.65 would capture 95% of zones that experienced high CO_2 concentration at some point over the study period across all stochastic occupancy scenarios), although monitoring of seating capacities in zones is rarely conducted which may hinder the use of this metric. Single-occupant zones with little to no exterior surface area should also be prioritized as critical zones; monitoring a single one of these zones had an 80% chance of capturing representative values (i.e., four out of the five of the single-occupant core zones routinely encountered peak CO_2 levels in excess of 800 ppm).

Table 2: Single-occupant zones with peak CO_2 concentrations exceeding 800 ppm.

# single-occupant zones	Zone number CO_2 > 800 ppm	Ventilation Seating
1	10(i)	4.20
2	1(i), 5(h)	2.64
3	7(c), 7(i), 16(c)	2.62
4	14(i), 15(c), 18(e), 18(i)	1.56

Conclusion

In this study, a floor of an academic office building was simulated using ten stochastic occupancy schedules, and zone-level CO_2 data were recorded. These data were analysed to explore preliminary recommendations for minimum CO_2 sensing infrastructure to enable DCV under realistic occupancy scenarios by identifying critical zones. Small multi-occupant zones with the lowest ventilation to seating capacity ratios (i.e., conference/meeting rooms) were found to have high peak CO_2 concentrations across the majority of the scenarios studied. Single-occupant core

zones were also disproportionately vulnerable to high peak CO_2 concentrations due to reduced outdoor airflow from a lack of infiltration. Large multi-occupant zones, however, including core zones, did not experience the same high CO_2 concentrations, likely due to their larger volumes. Therefore, across the 26 thermal zone case study floor, three CO_2 sensors may be sufficient to capture representative CO_2 trends across the floorplate: one sensor for the critical conference/meeting rooms, one sensor for the critical single occupant core zones, and a remaining sensor for a representative perimeter private office or larger multi-occupant zone.

This methodology highlighted the benefit of using stochastic occupancy data for EnergyPlus simulations when developing recommendations for sensor deployment and OCC, as diverse occupancy across simulation scenarios within a single zone can reveal trends that might not otherwise be apparent with traditional static occupancy. In future work, further simulations will be performed to evaluate how monitoring CO_2 and controlling DCV at different granularities (e.g., individual zone-level, a collection of zones using a single VAV, or at the system level) impacts the efficacy of DCV in terms of energy use and IAQ. Different critical zones for CO_2 monitoring will be identified and tested systematically across additional stochastic occupancy scenarios via simulation. The findings from this work will be used to inform controls implementations in the case study building once construction is complete; different DCV control schemes' impacts on energy savings and IAQ will be evaluated during these real-world implementations. Different or additional stochastic occupancy data (e.g., (Chen, Hong and Luo, 2018)) in archetypical building performance simulation models (e.g., DOE reference buildings (Deru *et al.*, 2011)) could broaden the insights of research topic beyond the case study building in question.

Acknowledgement

This research is supported by research funding provided by Natural Research Council (NRC) Canada, Grant No. 980444. The authors acknowledge the IEA EBC Annex 79 researchers, as we have greatly benefitted from the associated discussions.

References

Abuimara, T. *et al.* (2021) 'Current state and future challenges in building management: practitioner interviews and a literature review', *Journal of Building Engineering*, 41(May), p. 102803. doi: 10.1016/j.jobe.2021.102803.

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

Bennet, I. E. and O'Brien, W. (2017) 'Office building plug and light loads: Comparison of a multi-tenant office tower to conventional assumptions', *Energy and Buildings*, 153, pp. 461–475. doi: 10.1016/j.enbuild.2017.08.050.

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

Chen, Y., Hong, T. and Luo, X. (2018) 'An agent-based stochastic Occupancy Simulator', *Building Simulation*, 11(1), pp. 37–49. doi: 10.1007/s12273-017-0379-7.

Deng, Z., Morissette, R. and Messacar, D. (2020) *Running the economy remotely: Potential for working from home during and after COVID-19*. Ottawa.

Deru, M. *et al.* (2011) *U.S. Department of Energy commercial reference building models of the national building stock*, *NREL*. Golden, CO. Available at: http://digitalscholarship.unlv.edu/renew_pubs/44.

Fisk, W. J. and De Almeida, A. T. (1998) 'Sensor-based demand-controlled ventilation: a review', *Energy and Buildings*, 29(1), pp. 35–45. doi: https://doi.org/10.1016/S0378-7788(98)00029-2.

Golden, L. (2001) 'Flexible work Time: Correlates and consequences of work scheduling', *American Behavioral Scientist*, pp. 1157–1178.

Gunay, H. B. *et al.* (2016) 'Modeling plug-in equipment load patterns in private office spaces', *Energy and Buildings*, 121, pp. 234–249. doi: 10.1016/j.enbuild.2016.03.001.

Hobson, B. W. *et al.* (2019) 'Opportunistic occupancy-count estimation using sensor fusion: A case study', *Building and Environment*, 159(January), p. 106154. doi: 10.1016/j.buildenv.2019.05.032.

Hobson, B. W. *et al.* (2020) 'Clustering and motif identification for occupancy-centric control of an air handling unit', *Energy and Buildings*, 223, p. 110179. doi: 10.1016/j.enbuild.2020.110179.

Liu, G., Dasu, A. R. and Zhang, J. (2012) 'Review of Literature on Terminal Box Control, Occupancy Sensing Technology and Multi-zone Demand Control Ventilation', *PNNL Report PNNL - 21281*, (March).

Melfi, R. *et al.* (2011) 'Measuring building occupancy using existing network infrastructure', *2011 International Green Computing Conference and Workshops, IGCC 2011*. doi: 10.1109/IGCC.2011.6008560.

Mysen, M. *et al.* (2005) 'Occupancy density and benefits of demand-controlled ventilation in Norwegian primary schools', *Energy and Buildings*, 37(12), pp. 1234–1240. doi: 10.1016/j.enbuild.2005.01.003.

Mysen, M., Rydock, J. P. and Tjelflaat, P. O. (2003) 'Demand controlled ventilation for office cubicles—can it be profitable?', *Energy and Buildings*, 35(7), pp. 657–662. doi: 10.1016/S0378-7788(02)00212-8.

Nweye, K. and Nagy, Z. (2021) 'MARTINI: Smart Meter Driven Estimation of HVAC Schedules and Energy Savings

Based on WiFi Sensing and Clustering', pp. 0–13. Available at: http://arxiv.org/abs/2110.08927.

Park, J. Y. *et al.* (2019) 'A critical review of field implementations of occupant-centric building controls', *Building and Environment*, 165(May), p. 106351. doi: 10.1016/j.buildenv.2019.106351.

Qiu, W. *et al.* (2016) 'Clustering method based on local shrinking. Package "clues". The R project for statistical computing.'

Sun, Z., Wang, S. and Ma, Z. (2011) 'In-situ implementation and validation of a CO2-based adaptive demand-controlled ventilation strategy in a multi-zone office building', *Building and Environment*, 46(1), pp. 124–133. doi: 10.1016/j.buildenv.2010.07.008.

U.S. DOT (2000) *Department of Transportation and Related Agenices Appropriations, 2001*. US.

United States Department of Energy (2020) 'EnergyPlus Version 9.3.0 Documentation: The Encyclopedic Reference to EnergyPlus Input and Output'. U.S. Department of Energy.

Zhang, J. *et al.* (2013) *Energy Savings for Occupancy-Based Control (OBC) of Variable- Air-Volume (VAV) Systems*, *Pnnl-22072*. Richland, WA. doi: 10.2172/1063080.

Zivelonghi, A. and Lai, M. (2021) 'Mitigating aerosol infection risk in school buildings: the role of natural ventilation, volume, occupancy and CO2 monitoring', *Building and Environment*, 204(July), p. 108139. doi: 10.1016/j.buildenv.2021.108139.