



Modeling and multiobjective optimization of indoor airborne disease transmission risk and associated energy consumption for building HVAC systems

Michael J. Risbeck^{a,*}, Martin Z. Bazant^b, Zhanhong Jiang^a, Young M. Lee^a, Kirk H. Drees^a, Jonathan D. Douglas^a

^a Johnson Controls International plc, Milwaukee, WI, USA

^b Departments of Chemical Engineering and Mathematics, Massachusetts Institute of Technology, Cambridge, MA, USA

ARTICLE INFO

Article history:

Received 9 August 2021

Revised 17 September 2021

Accepted 20 September 2021

Available online 23 September 2021

Keywords:

Airborne infection risk

Energy cost

Building modeling

Multiobjective optimization

ABSTRACT

The COVID-19 pandemic has renewed interest in assessing how the operation of HVAC systems influences the risk of airborne disease transmission in buildings. Various processes, such as ventilation and filtration, have been shown to reduce the probability of disease spread by removing or deactivating exhaled aerosols that potentially contain infectious material. However, such qualitative recommendations fail to specify *how much* of these or other disinfection techniques are needed to achieve acceptable risk levels in a particular space. An additional complication is that application of these techniques inevitably increases energy costs, the magnitude of which can vary significantly based on local weather. Moreover, the operational flexibility available to the HVAC system may be inherently limited by equipment capacities and occupant comfort requirements. Given this knowledge gap, we propose a set of dynamical models that can be used to estimate airborne transmission risk and energy consumption for building HVAC systems based on controller setpoints and a forecast of weather conditions. By combining physics-based material balances with phenomenological models of the HVAC control system, it is possible to predict time-varying airflows and other HVAC variables, which are then used to calculate key metrics. Through a variety of examples involving real and simulated commercial buildings, we show that our models can be used for monitoring purposes by applying them directly to transient building data as operated, or they may be embedded within a multi-objective optimization framework to evaluate the tradeoff between infection risk and energy consumption. By combining these applications, building managers can determine which spaces are in need of infection risk reduction and how to provide that reduction at the lowest energy cost. The key finding is that both the baseline infection risk and the most energy-efficient disinfection strategy can vary significantly from space to space and depend sensitively on the weather, thus underscoring the importance of the quantitative predictions provided by the models.

© 2021 Elsevier B.V. All rights reserved.

1. Introduction

With the recognition that airborne exchange of respiratory aerosols may be the primary means of transmission for COVID-19 and other diseases [1–4], it is crucial to understand how HVAC systems can affect this infection route [5]. In order to reduce the aerosol concentration and associated infection risk in buildings, HVAC systems can deliver clean air flow via a combination of high-efficiency filtration, extra fresh-air ventilation, and other sources. Unfortunately, each option leads to increased energy

use, and their effect on indoor-air quality can vary significantly based on outdoor-air conditions. In this paper, we propose a model-based optimization framework to determine the combination of clean-air sources that can achieve a given reduction in infection risk in the most cost-effective manner. The optimization can recommend a spectrum of possible operational profiles ranging from minimum infection risk to minimum energy, providing predicted energy consumption and infection risk for each. This information can empower building managers to make informed decisions about how to operate their buildings given current cost and disinfection priorities.

Underpinning the proposed optimization is a set of dynamic models for zone temperature, humidity, and the expected

* Corresponding author.

E-mail address: michael.james.risbeck@jci.com (M.J. Risbeck).

transmission rate, which are parameterized in terms of physical and usage characteristics of the zone (e.g., floor area, ceiling height, design occupancy, HVAC equipment sizes) that are representative of typical commercial buildings. These models can then be used to predict the effect of various design and operational variables on energy and infection risk, thus allowing actionable and informed recommendations to be made. We begin by reviewing the relevant background for airborne transmission modeling in buildings along with the associated energy analysis, and then we present the main goals of this paper.

1.1. Airborne disease transmission

For airborne diseases, the primary transmission route is via viruses or bacteria that are entrained in aerosolized particles released into the air by exhalation of an actively infectious individual [3,4]. These expiratory aerosols are partially dried droplets of respiratory fluids, such as saliva [6] or airway mucus [7], and they are composed of aqueous solutions of proteins and salts, along with potentially infectious microbes. The concentration and size distribution of the particles covers several orders of magnitude [8,9], with significant dependence on vocalization volume [2,10] and respiratory mode [11,12]. The largest of these droplets, produced primarily by coughs and sneezes [13] can be up to 1 mm in diameter. However, because they settle rapidly to the ground, these large droplets generally do not pose a risk for long-range airborne transmission [14,15]. By contrast, the smallest droplets are sized between 0.1 and 10 μm and are produced in high concentrations by standard respiration. Because these particles settle much more slowly, they tend to remain suspended in the air for long periods of time that exceed the time scales for turbulent mixing and air change in an indoor space [4]. Such aerosol particles are of primary concern for airborne transmission, as they can spread over long distances due to the air circulation provided by natural convection and HVAC systems [16–18], thus leading to new infections if inhaled by susceptible individuals.

Although infectious aerosols can remain airborne for long periods, their concentration does decay over time due to a variety of removal mechanisms [4]. For example, the particles naturally deposit onto surfaces (with a size-dependent rate [19,20]) as fomites, thus rendering them unlikely to transmit the disease [21]. In addition, the particles may simply be removed from the air via outdoor-air ventilation (with particle-laden air exhausted to the ambient and replaced with a similar volume of outdoor air free from respiratory particles) or filtration (by which some fraction of the particles become trapped in a filter) [22–24]. In addition, the infectious material in the particles can become deactivated, e.g., due to natural decay (the rate of which depends slightly on temperature and humidity) or deliberate UV irradiation [25–28]. Thus, these mechanisms result in a gradual decay of the airborne infectious particle concentration over time. Ultimately, the concentration of infectious aerosols in the space is determined by the balance of generation and removal mechanisms. However, because conditions are often unsteady throughout the course of the day (e.g., with generation rates varying in accordance with the occupancy cycle and removal rates determined by the transient operation of the HVAC system), it is important to account for dynamic effects when considering infection risk.

Due to the significant likelihood and associated negative consequences of indoor disease transmission, there is growing interest in mitigating the associated risk by changing the design or operation of building HVAC systems. Many such strategies are aimed at reducing the generation rate of infectious particles within the space and reducing the exposure of susceptible individuals to those particles. For example, by requiring occupants to wear face masks, the particle concentrations released by exhaling infectors and

received by inhaling susceptibles are reduced in accordance with the filtration efficiency of the mask [4,29]. Masks also largely eliminate the momentum of respiratory flows, such as turbulent plumes from speech [30] or coughs and sneezes [14], which would otherwise enhance the risk of short-range airborne transmission [31] relative to the long-range airborne transmission in a well mixed space [4]. Furthermore, screening occupants for high temperatures or other symptoms prior to entering a building can significantly reduce the expected number of infectors in the space (and thus also the infectious particle generation rate), although the effectiveness of this strategy is attenuated for diseases such as COVID-19 for which asymptomatic spread is prevalent [6,32]. By contrast, other commonly employed measures may have little to no effect on airborne transmission. For example, the installation of plexiglass dividers and enforcement of social distancing can reduce the likelihood of short-range transmission associated with respiratory jets [4,6,31] or large droplets [14,15], but once the smaller particles have mixed with the surrounding air, these strategies no longer provide protection to susceptible occupants. These observations indicate that focusing on the long-range airborne transmission route is a more effective means of reducing infection rates. Indeed, recent analyses of COVID-19 in schools found no significant change in transmission rates resulting from 3-foot versus 6-foot social distancing requirements [33] or whether physical barriers were installed to separate desks, but reduced transmission was observed from mask usage, increased ventilation, and improved filtration [34].

Although these findings indicate possible strategies for reducing airborne transmission, and even which among them are likely to be most effective, the value of such *qualitative* guidance is somewhat limited. Specifically, knowing that a particular action provides *relative* reduction in infection risk does not indicate whether it is necessary in an *absolute* sense. Given the significant diversity of building operation and occupant behavior, the expected rate of indoor infection can vary over several orders of magnitude for different buildings or spaces [4]. Thus, in order to guide decision-making, it is important to provide *quantitative* guidance for various strategies.

1.1.1. Modeling framework

Mathematical models to estimate the rate of airborne disease transmission have been known for over 50 years, with seminal contributions from Wells [22] and Riley et al. [35]. Within the Wells-Riley framework, the hypothetical concentration of infectious aerosols is modeled using material balances, from which the received dose and infection probability of susceptible individuals in the space can be calculated. This basic approach has been extended in many directions since its original proposal [38,19,20,23,24,36,37,4,39–41]. In general terms, the key relationships for one susceptible individual can be expressed as follows:

$$\begin{aligned} \text{Dose} &= \int \alpha \times \text{Concentration } dt \\ \text{Infection Probability} &= 1 - \exp(-\beta \times \text{Dose}) \\ &\approx \beta \times \text{Dose} \end{aligned}$$

That is, the individual's received dose (measured in some unit of amount, e.g., number or mass of virions) is equal to the time integral of the airborne concentration of infectious particles (as modeled by the material balance) multiplied by an exposure rate α (generally equal to the individual's breathing rate). That individual's probability of infection is then given by an exponential relaxation to certainty proportional to the total dose, in which the proportionality constant β reflects the infectivity or hazard rate of the dose. From Wells [22], it is common to measure each amount of infectious particles in units of "infectious quanta", where one "quantum" (abbreviated q) represents the amount of pathogen

exposure that results in an infection probability of $1 - \exp(-1) \approx 0.63$. (In these units, the proportionality constant β thus has a value of 1 q^{-1} .) More generally, infection quanta may be defined in terms of the mean rate of transmission between each infected-susceptible pair [4], without regard to the microscopic pathogen concentration.

In order to determine the resulting infectious particle concentration, a common assumption is to assume that the particles are well-mixed (i.e., spatially uniform in concentration) within the air. Indeed, the same assumption is commonly made when modeling other mass and energy balances within buildings, and it can be justified both theoretically [4] and empirically via simulation [18] and experiment. Where the effects are relevant, the well-mixed models can be corrected to account for short-range respiratory jets [6,31] or fluctuations in airflow [16], but it is nevertheless a useful first approximation to assume well mixed indoor air. Another complication is particle-size dependence of various processes, primarily the filtration efficiency and deposition rate, but also infectivity and decay rate [4]. These effects can be accounted for, but it is common to include them only via size-averaged variables in existing models [4,20,42]. Additional non-uniformity can result when a space is separated into multiple rooms that share the same supply air stream, but modeling in such cases is straightforward, if the room dimensions, occupancy, and airflow are known [43].

1.1.2. HVAC mitigation strategies

As alluded to above, the operation of building HVAC systems can have a significant impact on the airborne concentration of infectious particles. The primary sources of removal and deactivation are ventilation, filtration, and UV disinfection. Some past studies have assessed which HVAC strategies are most effective and energy-efficient for infection risk reduction [37,40], but findings can vary significantly based on climate and HVAC system design. A key focus from ASHRAE [44] and others has been on increasing ventilation rates. Indeed, increased ventilation does reduce the effective concentration of infectious particles, and in addition the associated reduction in indoor CO_2 concentration is beneficial for general health and wellness [45]. However, if the additional ventilation is provided only by increasing the outdoor air *fraction* in the supply air stream, the supplemental reduction in infectious particle concentration can be small to negligible if that stream also passes through a filter with a high MERV rating [44]. Similarly, switching to a higher-efficiency filter will remove more infectious particles from the air, but if the volume of air passing through that filter is low, then absolute benefits may not be significant [37]. Thus, it is important to consider all the disinfection strategies available and account for how they interact when making decisions. In particular, to provide meaningful reduction in infection risk, it is often necessary to increase the *total* airflow provided by the HVAC system so as to more rapidly dilute and remove infectious particles present in the air.

A key caveat when considering HVAC infection mitigation strategies is that due to system design, it is often not possible to directly manipulate the desired variables. For example, a practitioner may want to double the ventilation and total airflow provided to a zone. However, in constant-volume systems, the supply flow generally cannot be increased above its design value, and so the only way to increase total flow is to install standalone filtration devices that operate on their own dedicated stream of air. Alternatively, in variable-volume systems, the flow delivered to a given room is determined by the action of the VAV controller for the purposes of temperature regulation, and so absolute flow can only be adjusted indirectly (e.g., by manipulating supply-air temperatures). Finally, if the system does not have an airside economizer, it may not be possible to increase ventilation rates above their minimum values, and even when present, the maximum ven-

tilation rate may be further restricted based on weather conditions and heating/cooling capacities. The main observation is that when considering HVAC mitigation strategies for airborne infection risk, it is not sufficient to consider just the fundamental physical variables, and instead one must step further back to use the setpoints exposed by the HVAC system and predict their effect on the relevant physical variables.

1.2. Energy consumption

Although HVAC systems can be operated so as to reduce the risk of airborne transmission of disease, unfortunately the actions taken will almost always *increase* energy consumption relative to baseline operation. Thus, when choosing between two strategies, it is important to consider which is most energy-efficient. A common basis to quantify the disinfection efficacy of a given mechanism is its removal rate of infectious particles, which is best expressed as the equivalent volumetric flow of particle-free air that would provide the same removal rate. This concept is referred to as the “clean air delivery rate” [46] or the “equivalent outdoor air” rate [44], and two devices that have the same value are thus equivalent from an infection-risk perspective and can be compared in terms of energy consumption. However, when combining multiple clean-air mechanisms in series (e.g., using an air-filter downstream from the outdoor-air intake), it is important to consider their *combined* efficacy, as there may be only incremental benefits to clean-air delivery but a significant change in energy consumption.

Given these associated energy costs and the societal emphasis on energy efficiency and sustainability, it is important to make informed decisions about which strategies to apply. This sentiment has been echoed by the ASHRAE Epidemic Task Force [47], which has recommended “select[ing] control options, including standalone filters and air cleaners, that provide desired exposure reduction while minimizing associated energy penalties”. In order to make this determination, it is necessary to estimate the energy consumption associated with possible courses of action. This knowledge will enable building managers to make informed design decisions while adapting operational decisions regularly to account for changing weather or infection-prevalence conditions.

1.2.1. Filtration and fan power

A key technology for removing infectious particles from a space is filtration. The overall clean-air delivery rate associated with a filter is the product of the volumetric flow of air through the filter multiplied by its filtration efficiency [44]. This observation suggests two different ways to utilize filters, which have different impacts on operation and energy consumption.

The first way to increase the clean-air delivery provided by a filter is to increase the amount of air flowing through it. Physically, this increased flow is provided by increasing fan speed, which thus requires additional electricity consumption that can be calculated as the product of volumetric flow and pressure rise divided by overall fan efficiency. Because the pressure rise and efficiency each vary with flow, the result is essentially a static nonlinear model in terms of flow [48]. The simulation package EnergyPlus [49] supports various mathematical forms for variable-speed fan energy models, including low-order polynomials and a power law. Although model accuracy is best when calibrated against actual operating data, reasonable accuracy can be achieved by scaling standard models with nameplate design flow and power consumption of the fans [48].

As an alternative to increasing total flow, particle removal rates can be increased by switching to a filter with higher efficiency. The MERV rating standard specifies required filtration efficiencies for three different ranges of particle sizes, and thus for a given particle size distribution, these values can be used to calculate effective the

overall filtration efficiency of each MERV rating [44]. Unfortunately, this increase in efficiency restricts the flow of air through the filter, thus increasing the required pressure difference to produce the same total flow rate. In general, the higher pressure difference results in higher energy consumption, although the magnitude can vary based on the fan control strategy that is applied [50]. To a reasonable approximation, the increased energy consumption can be estimated by considering the change in pressure drop associated with the new filter type and appropriately scaling the baseline energy consumption [37].

1.2.2. Ventilation and coil loads

When supplying additional ventilation, there is a change the heating or cooling load necessary to maintain indoor air temperature, which thus results in a change in energy consumption. Past studies have quantified the extra costs associated with increased ventilation rates [45,51,52], and a key finding is that there is significant variation with climate. The impact of day-to-day variation can perhaps be minimized via the use of an economizer (which will automatically increase ventilation rates when outdoor conditions make it energetically favorable), but the commonly employed rule-based strategies are often insufficient to capture all the available benefits [53]. In Azimi and Stephens [37], the energy costs of extra ventilation were explicitly compared to that of higher-efficiency filtration for various climates and filter types, finding that filtration is generally the more energy-efficient option overall. However, results can change seasonally with weather conditions, and the study did not consider the energy consumption associated with dehumidification. In general, it is recognized that ventilation is a costly component of HVAC system operation, and so [54] specifies *minimum* ventilation rates that must be provided and allows technologies like demand-controlled ventilation to further reduce the amount of ventilation that must be supplied. Energy-recovery technologies can be applied to reduce the energy impact of ventilation [55], but they are not commonly employed in existing buildings.

The energy costs associated with ventilation ultimately derive from the fact that the outdoor air may have to be heated or cooled to bring it to suitable conditions of temperature and humidity for indoor comfort. Thus, the energy consumption should be modeled by considering the thermal loads at the appropriate coils. For cooling coils, effectiveness models can capture heat-transfer limitations and thus provide an accurate estimate of energy consumption [56–58]. However, simplified contact-mixture models can also be used and have been shown to match experimental data despite their relative simplicity [59]. For heating coils, the modeling principles are similar but generally simpler because there is no associated change in humidity that must be considered as a latent load. The calculated coil loads can then be converted to the appropriate energy source (generally electricity or gas) via the appropriate coefficient of performance or efficiency factor [37].

The key differentiator between ventilation and filtration for the purposes of reducing airborne transmission is that for a given level of airflow or particle removal, the energy cost of filtration is essentially constant, whereas the energy cost of ventilation varies significantly (and can even become negative) based on outdoor-air conditions. Thus, optimal energy efficiency requires operating strategies to be adjusted seasonally (or more frequently), perhaps with filtration preferred in the extreme heating and cooling months and ventilation used during milder transition seasons.

1.2.3. Other disinfection devices

For completeness, we note that infectious particles can also be removed or deactivated by auxiliary disinfection devices, including UV irradiation zones in ducts or standalone filters and air cleaners that are placed in the space. For UV disinfection, the radiation

intensity to achieve a given level of disinfection can be computed from expected flow rates [28]. However, once the UV lamps have been installed, their electricity consumption is essentially constant. Thus, for energy analysis, it is only necessary to know the total rated power of the UV lamps (and of course also whether they are on or off). For filtration devices, theoretical models have been developed to estimate and optimize overall particle removal rate for a given electricity consumption [60], but as before, the main value relevant for energy analysis is the device's rated power. Where variable operation is possible (e.g., for a device that can operate at different fan speeds or flow rates), it may be necessary to fit simple models to account for variation, possibly with data obtained from manufacturer specifications.

1.3. Overview of paper and contributions

Given the importance of managing both indoor infection risk and HVAC energy consumption [47], the goal of this work is to present dynamic models that can be used for monitoring and optimization of both variables in buildings. The models for indoor temperature, humidity, and hypothetical infectious particle concentration are based on mass and energy balances that follow from standard physics. By contrast, the model of the regulatory control system is phenomenological so as to match general HVAC system behavior (including time variation in airflow resulting from variation in thermal loads) while remaining computationally efficient. Both sets of models are parameterized in terms of readily available data, including space dimensions, equipment capacities, weather conditions, and occupancy profiles. Where values are not specifically known, they can be substituted with standard values provided in ASHRAE standards [54,55]. The main contribution is that the models are coupled, and thus they predict simultaneously the infection rate, energy consumption, and indoor conditions that result from commonly available HVAC setpoints and design variables. Thus, these models can provide *actionable* guidance that can be readily applied in real buildings. In particular, they facilitate Pareto analysis to assess the optimal tradeoff between the two primary metrics so that both aspects can be considered.

The remainder of this paper is structured as follows. In Section 2, we present the full set of dynamic models for airborne infection, thermal dynamics, and the regulatory control system. For brevity, we present only the models themselves (i.e., equations and descriptions of variables), but additional detailed discussion is provided in the appendices. In Section 3, we then discuss how the models can be used for monitoring and optimization purposes. These calculations could be used by building managers to evaluate the infection risk in their buildings as operated and determine how infection rate could be reduced or energy efficiency can be improved. We then present examples in Section 4 to illustrate the insights and operational benefits that can be provided by application of the modeling framework. Finally, Section 5 provides conclusions and a discussion of future directions.

2. Modeling

In this section, we present various dynamic models that can be used to estimate and predict infection risk and HVAC energy consumption within an indoor zone. All models are defined in the standard continuous-time state-space form

$$\frac{dx}{dt} = f_{\theta}(x, u), \quad y = h_{\theta}(x, u)$$

for states x , inputs u , outputs y , and static parameters θ . (Note the distinction between u and θ is that values of u vary throughout the simulation period, whereas values of θ are held constant.) Mod-

els are simulated using the standard 4th order explicit Runge–Kutta method assuming a piecewise-constant hold on u and calculating y as the *average* value over the integration interval, which for our purposes is $\Delta = 15$ min.

Performance metrics are all calculated for each timestep as functions $\ell_\theta(y, u)$ that are then integrated as appropriate over the simulation interval. In various contexts, they are used as optimization objectives or constraints.

2.1. Airborne disease transmission

In order to predict the transmission risk within each zone, we model the hypothetical concentration N of airborne infectious particles that would be produced by infectious individuals in the space. Consistent with the infection probability distribution suggested by the Wells–Riley equation [37], we model this concentration as “quanta” per volume (unit q/m^3) where one “quantum” (unit q) of infectious particles is defined as the average amount of infectious particles that will lead to infection in a susceptible person. For brevity, we refer to the concentration of airborne infectious particles as the “quanta concentration”. From this concentration, we calculate the hypothetical exposure rate for susceptible individuals which then produces an estimate of the number of transmission events during the evaluation period.

With this general approach, the main dynamic model for the airborne infectious quanta concentration C is as follows:

$$\begin{aligned} V \frac{dC}{dt} &= q_I - f_{\text{clean}} C \\ f_{\text{clean}} &= k_{\text{natural}} V + \eta_{\text{AHU}} f_{\text{AHU}} + \sum_j a_{\text{aux},j} f_{\text{aux},j}^{\text{max}} \\ \eta_{\text{AHU}} &= 1 - (1 - x_a)(1 - a_{\text{UV}} \eta_{\text{UV}})(1 - \eta_{\text{filter}}) \\ q_I &= (1 - \mu_I) N_I C f_{\text{breathe},I} \\ q_S &= (1 - \mu_S) N_S C f_{\text{breathe},S} \end{aligned} \quad (1)$$

From these variables, we then calculate the primary metrics via

$$\begin{aligned} N_T &= \int_0^\tau \beta_q q_S(t) dt \\ E_{\text{aux}} &= \int_0^\tau a_{\text{UV}} W_{\text{UV}} + \sum_j a_{\text{aux},j} W_{\text{aux},j} dt \end{aligned} \quad (2)$$

where N_T is the expected number of transmission events that occur and E_{aux} the energy consumption associated with UV and auxiliary disinfection, both over the time period $t \in [0, \tau]$. Descriptions and units of all variables are given in Table 1.

Although similar to other dynamic formulations of the Wells–Riley equation, the main contribution of this model is that it represents *all* infectious-particle removal sources as an effective clean-air flow f_{clean} . This value gives the volumetric flow of clean (i.e., $C = 0$) air that would displace infectious particles at the same total rate as the mechanisms being considered. (As mentioned before, similar concepts have been referred to as “clean air delivery rate” [46] and “equivalent outdoor air” [44], but those terms generally refer only to a subset of the various removal mechanisms.) The key benefits of this quantity are that its definition as a volumetric flow makes it more intuitive for HVAC engineers, and that (in the cases considered) its value is independent of the underlying infectious particle concentration C , allowing it to be calculated from readily available building data. In addition, by linearizing the standard exponential infection probability for each individual, we can thus aggregate across occupants and estimate the expected number of transmissions specifically within the specific space being considered. This *space-centric* formulation is critical to guide decision-making, as it allows identification of high-risk spaces and separate optimization of disinfection strategies tailored to

Table 1
Variables in the infectious particle concentration model.

Symbol	Description	Unit
<i>States</i>		
C	airborne concentration of infectious particles	q/m^3
<i>Inputs</i>		
f_{AHU}	supply and return airflow from and to the AHU	m^3/s
x_a	outdoor air intake fraction	–
a_{UV}	activation fraction for UV disinfection	–
$a_{\text{aux},j}$	activation fraction for in-zone disinfection device j	–
N_I	number of infectious individuals in the zone	–
N_S	number of susceptible individuals in the zone	–
<i>Outputs</i>		
f_{clean}	total effective clean-air flow	m^3/s
q_I	total infection quanta generation rate across all infectious occupants	q/s
q_S	total infection quanta dose rate received across all susceptible occupants	q/s
η_{AHU}	total effective disinfection efficiency for air passing through the AHU	–
<i>Performance Metrics</i>		
N_T	expected number of infection transmissions to susceptibles over the evaluation period	#
E_{aux}	energy consumption associated with UV and auxiliary disinfection equipment	kJ
<i>Parameters</i>		
V	volume of air in the zone	m^3
C_I	infection quanta net concentration in the exhaled breath of infectors	q/m^3
k_{natural}	infection quanta deactivation rate due to natural processes (including deposition and decay)	s^{-1}
η_{filter}	AHU filter disinfection fraction	–
η_{UV}	AHU UV disinfection fraction at maximum activation	–
$f_{\text{aux},j}^{\text{max}}$	in-zone clean-air rate at maximum activation for device j	m^3/s
$f_{\text{breathe},I}$	average breathing rate of infectious individuals in the zone	m^3/s
$f_{\text{breathe},S}$	average breathing rate of susceptible individuals in the zone	m^3/s
μ_I	attenuation factor for infector generation rate (e.g., due to masks)	–
μ_S	attenuation factor for susceptible dose rate (e.g., due to masks)	–
τ	evaluation time for metrics	s
β_q	hazard rate for transmissions	q^{-1}
W_{UV}	power consumption of in-duct UV at maximum activation	kW
$W_{\text{aux},j}$	power consumption of in-zone disinfection device j at maximum activation	kW

each space. Specific details about the model are discussed in Appendix A, including the clean-air delivery calculation, the number-of-transmissions equation, and values for key parameters and inputs.

2.2. Thermal dynamics and energy consumption

To model the evolution of temperature and humidity within the zones, we use straightforward (thermal) energy and (water vapor) mass balances respectively. Although there are some pressure fluctuations in the space, they are generally small enough to allow a constant-density assumption for air within the zones. Consistent with the previous infection model, the volume within a zone is assumed to be well-mixed and thus have uniform temperature and humidity. Note that throughout this discussion, we use “humidity” to mean the “absolute humidity ratio” (commonly denoted ω) unless otherwise specified.

ODE equations for the mass and energy balances are as follows:

$$\begin{aligned}\rho_a c_a V \frac{dT_z}{dt} &= \rho_a c_a f_{AHU} (T_{AHU} - T_z) \\ &\quad + Q_h + k_{az} (T_a - T_z) + Q_z \\ \rho_a V \frac{d\omega_z}{dt} &= \rho_a f_{AHU} (\omega_{AHU} - \omega_z) + w_z\end{aligned}\quad (3)$$

The AHU outlet conditions and loads are calculated from simple heating and cooling coil models with the following functional form (and described in more detail in Section B.2):

$$\left(T_{AHU}, \omega_{AHU}, \Delta h_{AHU}^h, \Delta h_{AHU}^c \right) = AHU(T_z, \omega_z, x_a, T_a, \omega_a, T_{AHU}^{sp}, x_{bypass}) \quad (4)$$

The electricity consumption of the fan is given by a cubic polynomial (see Section B.3 for more details):

$$w_{fan} = c_1 \frac{f_{AHU}}{f_{AHU}^{max}} + c_2 \left(\frac{f_{AHU}}{f_{AHU}^{max}} \right)^2 + c_3 \left(\frac{f_{AHU}}{f_{AHU}^{max}} \right)^3 \quad (5)$$

Finally, from all of these outputs we the main metrics associated with energy consumption and thermal comfort in the space:

$$\begin{aligned}E_{HVAC} &= \int_0^\tau \rho_a f_{AHU} \left(\beta_h \Delta h_{AHU}^h - \beta_c \Delta h_{AHU}^c \right) + \beta_h Q_h + w_{fan} dt \\ \epsilon_T &= \int_0^\tau \max(T_z - T_z^{max}, 0) + \max(T_z^{min} - T_a, 0) dt \\ \epsilon_\omega &= \int_0^\tau \max(\omega_z - \omega_z^{max}, 0) + \max(\omega_z^{min} - \omega_z, 0) dt\end{aligned}\quad (6)$$

Descriptions and units of all variables in these equations can be found in Table 2.

For the most part, this model follows standard thermal modeling principles for buildings, with some simplifications to avoid

Table 2
Variables in the thermal dynamics and energy model.

Symbol	Description	Unit
<i>States</i>		
T_z	temperature of the air inside the zone	°C
ω_z	humidity ratio of the air inside the zone	-
<i>Inputs</i>		
f_{AHU}	supply and return airflow from and to the AHU	m ³ /s
x_a	ambient air intake fraction	-
Q_h	sensible heating from HVAC heating sources	kW
T_{AHU}^{sp}	AHU supply temperature setpoint	°C
T_a	ambient temperature	°C
ω_a	ambient humidity ratio	-
Q_z	sensible thermal load on the zone	kW
w_z	latent moisture load on the zone	kg/s
T_z^{max}	maximum comfortable zone temperature	°C
T_z^{min}	minimum comfortable zone temperature	°C
ω_z^{max}	maximum comfortable zone humidity	-
ω_z^{min}	minimum comfortable zone humidity	-
<i>Outputs</i>		
T_{AHU}	actual supply temperature of AHU	°C
ω_{AHU}	actual supply humidity of the AHU	-
Δh_{AHU}^h	enthalpy difference across the AHU heating coil	kJ/kg
Δh_{AHU}^c	enthalpy difference across the AHU cooling coil	kJ/kg
w_{fan}	electricity consumption at the fan	kW
<i>Performance Metrics</i>		
E_{HVAC}	total HVAC energy consumption	kJ
ϵ_T	total temperature constraint violation	°C.s
ϵ_ω	total humidity constraint violation	s
<i>Parameters</i>		
ρ_a	density of air	kg/m ³
V	volume of air in the zone	m ³
c_a	heat capacity of air (at average humidity)	kJ/kg.°C
k_{az}	zone/ambient heat transfer coefficient	kW/°C
x_{bypass}	coil model bypass airflow fraction	-
c_i	fan power curve coefficients	kW
f_{AHU}^{max}	maximum flow in the AHU	m ³ /s
β_c	energy conversion factor for cooling	-
β_h	energy conversion factor for heating	-

parameters and inputs that are not readily available in real buildings (e.g., accurate solar intensity forecasts or predictions of thermal mass dynamics). We consider both temperature and humidity to ensure that energy consumption is accurately calculated and that occupant comfort is thoroughly considered. Key features of this model are discussed in more detail in Appendix B.

2.3. Control layer

An important feature of the modeling framework is that we consider the actions of the regulatory control layer in response to supplied controller setpoints and the thermal evolution of the zone. This functionality is critical because two of the inputs to the infection model (airflow f_{AHU} and outdoor-air fraction x_a) that have a significant impact on clean-air delivery *cannot* be directly specified in real buildings and instead are determined as a side effect of the regulatory control layer's actions. Although the controllers in each piece of equipment operate according to explicitly defined control laws, those laws often consist of complicated rule-based logic that operates on much faster timescales than desired simulation timestep. Thus, rather than try to replicate the control laws directly, we instead choose a phenomenological approach in which we attempt to model the *outcomes* of the controllers' actions. For example, rather than trying to model the (heavily modified) PI control law within the VAV controllers to specify airflow as a function of setpoint tracking error, we instead assume that the controller is well-designed and thus maintains temperature setpoints within the space whenever possible. This overall approach is similar to the simulation strategy used by EnergyPlus [49], and it simplifies our modeling effort by allowing us to use information about internal system states and disturbances that would not be directly available to the controllers.

The main controller output variables are calculated using the following equations:

$$\begin{aligned}f_{AHU} &= \text{clip}(\tilde{f}_{AHU}, f_{vent}^{min}, f_{AHU}^{max}) \\ Q_h &= \text{clip}(\tilde{Q}_h, 0, Q_h^{max}) \\ x_a &= \begin{cases} x_a^{max} & \text{if } T_a \leq T_a^{econ} \\ x_a^{min} & \text{else} \end{cases} \\ a_{UV} &= \begin{cases} \bar{a}_{UV} & \text{if } \sigma = 1 \\ 0 & \text{else} \end{cases} \\ a_{aux,j} &= \begin{cases} \bar{a}_{aux,j} & \text{if } \sigma = 1 \\ 0 & \text{else} \end{cases}\end{aligned}\quad (7)$$

These calculations use the following intermediate variables:

$$\begin{aligned}T_{AHU}^{sp} &= \bar{T}_{AHU}^{sp} \\ f_{vent}^{min} &= \begin{cases} \bar{f}_{vent} & \text{if } \sigma = 1 \\ 0 & \text{else} \end{cases} \\ x_a^{min} &= f_{vent}^{min} / f_{AHU} \\ x_a^{max} &= \text{clip}\left(\frac{\max(T_z - T_{AHU}^{sp}, 0)}{\max(T_z - T_a, 0)}, x_a^{min}, 1\right)\end{aligned}\quad (8)$$

Note that the clip function is defined as $\text{clip}(x, a, b) = \min(\max(x, a), b)$ for $a \leq b$. Variables used in this model are described in Table 3.

Overall, the main functionality of the controller model is that it allows prediction of the HVAC control variables that ultimately determine infection risk and energy consumption in the space. In an ideal system, we would be able to manipulate (and thus optimize) the airflow f_{AHU} and outdoor-air fraction x_a directly, but in existing HVAC systems, these quantities reflect the response of the regulatory control system and thus can be adjusted only indi-

Table 3
Variables in the HVAC controller model.

Symbol	Description	Unit
<i>Inputs</i>		
T_z	temperature of the air inside the zone	°C
ω_z	humidity ratio of the air inside the zone	–
T_a	ambient temperature	°C
ω_a	ambient humidity ratio	–
Q_z	sensible thermal load on the zone	kW
T_z^{max}	maximum comfortable zone temperature	°C
T_z^{min}	minimum comfortable zone temperature	°C
f_{AHU}^{max}	maximum AHU air flow	m ³ /s
Q_h^{max}	maximum reheat rate	kW
σ	occupancy indicator variable	{0, 1}
<i>Outputs</i>		
f_{AHU}	supply and return airflow from and to the AHU	m ³ /s
x_a	outdoor air intake fraction	–
Q_h	sensible heating from HVAC heating sources	kW
T_{AHU}^{sp}	AHU supply temperature setpoint	°C
f_{vent}^{min}	minimum outdoor-air ventilation rate	m ³ /s
a_{UV}	activation fraction for in-duct UV disinfection	–
$a_{aux,j}$	activation fraction for in-zone disinfection device j	–
\tilde{f}_{AHU}	hypothetical value of f_{AHU} that drives T_z to T_z^{max} in one timestep (see Section C.1)	m ³ /s
\tilde{Q}_h	hypothetical value of Q_h that drives T_z to T_z^{min} in one timestep (see Section C.1)	kW
x_a^{min}	minimum outdoor-air fraction needed to provide minimum ventilation	–
x_a^{max}	maximum outdoor-air fraction that can meet supply temperature setpoint	–
<i>Parameters</i>		
T_{AHU}^{sp}	default AHU supply temperature setpoint	°C
T_a^{econ}	economizer threshold ambient temperature	°C
\tilde{f}_{vent}	occupied outdoor-air minimum ventilation rate	m ³ /s
\bar{a}_{UV}	maximum activation fraction for UV disinfection	–
$\bar{a}_{aux,j}$	maximum activation fraction for in-zone disinfection device j	–

rectly. Using this simple controller model, we thus bridge the gap between the controller settings we can manipulate and the HVAC variables required for infection-risk and energy modeling. Specifically, by coupling to the thermal model from Section 2.2, we can adjust the static controller parameters θ and predict how the HVAC control system will respond. The resulting metrics can then be used to optimize for energy consumption or infection risk depending on our current priorities. Specific controller model components are discussed in more detail in Appendix C.

3. Monitoring and optimization

The models presented in Section 2 have two primary use cases. First, where actual measurements (air flows and temperatures) are available, the models can be used in a monitoring capacity to estimate the clean-air delivery and resulting infection risk within a space. These calculations can be used to determine whether each space is or is not being provided adequate disinfection as actually operated. Second, for spaces with high infection risk or energy consumption, the models can be used for optimization to predict the effects of changing various design or operational parameters. By evaluating the tradeoff between infection risk and energy consumption in a space, decision-making can be guided in accordance with current priorities. We discuss these two use cases in the following sections.

3.1. Monitoring

When used for monitoring purposes, we assume that measurements of temperatures and airflows are available from the building

management system (BMS). Thus, we need not use the thermal and controller models from Sections 2.2 and 2.3, as the relevant output quantities are already known. We do, however, need to simulate the airborne infection model from Section 2.1 to assess infection risk in the space. The two primary outputs of interest are the expected number of transmissions N_T and the total clean-air delivery rate f_{clean} ; the former can be used to quantify the infection risk in the zone, while the latter indicates to what extent HVAC measures are reducing that risk. Providing both values is important to be able to assess whether HVAC measures are providing sufficient mitigation (as indicated by high f_{clean} and low N_T) or whether non-HVAC operational changes are needed (as indicated by high N_T despite high f_{clean}).

To simulate the infection model, we require values for the airflow inputs f_{AHU} and x_a , as well as the activation fractions a_{UV} and $a_{aux,j}$ for the relevant disinfection devices. The numbers of infectors and susceptibles N_S and N_I can be chosen from a typical occupancy profile as discussed in Section A.4. If the evaluation period begins during unoccupied hours, then the initial condition for the infection quanta concentration C can be set to zero; otherwise, a “warmup” period of roughly 5 h should be simulated to obtain a more accurate initial condition. From there, the hypothetical infectious particle concentration in the space can be simulated via Eq. (1), and the expected number of transmissions can be calculated using Eq. (2).

Where estimates of energy consumption are required, the AHU and fan models from Eqs. (4) and (5), can be used, since their inputs are generally all known. In addition to f_{AHU} and x_a , these models also require the zone conditions T_z and ω_z , the ambient conditions T_a and ω_a , and the AHU conditions T_{AHU} and ω_{AHU} (although ω_{AHU} can also be estimated from the AHU model Eq. (4), as its value is often not measured in real buildings). In-zone heating energy can also be included where Q_h is known, although its value generally cannot be directly measured and must be estimated from other sources (e.g., heating-coil activation commands). We discuss some additional use cases in the following sections.

3.1.1. Uncertainty quantification

As mentioned in Section 2.1, many of the key infection-model parameters are related to occupant behavior or disease characteristics, which means they are somewhat uncertain in value. For the monitoring procedure described above, we assume mean values of all parameters to simulate mean values of the key infection metrics. However, it may be useful to quantify the uncertainty associated with the model's predicted number of infection transmissions N_T .

Fortunately, for many of the key model parameters, uncertainty quantification is straightforward. First, we note that the infection quanta concentration ODE in Eq. (1) is *linear* in the product of parameters $(1 - \mu_I)f_{breathe,I}C_I$. Thus, the (time-varying) value of the infection quanta concentration C is an *affine* function of this group. We know also that if $C_I = 0$, then $C \equiv 0$, and thus we have that C is *directly proportional* to these parameters. Second, the performance metric N_T is proportional to the product of parameters $(1 - \mu_S)f_{breathe,S}$ that multiply C to determine q_S , and thus its value is *directly proportional* to them. Combining these two facts, we find the relationship

$$N_T \propto (1 - \mu_I)(1 - \mu_S)f_{breathe,I}f_{breathe,S}C_I$$

Thus, in order to obtain distributional information for N_T , we need only run a single simulation of Eq. (1) to find a baseline value of N_T and then sample its distribution by rescaling the baseline value using samples of the other five parameters.

Unfortunately, the remaining model inputs and parameters are either time-varying or enter the dynamic model nonlinearly. Thus,

estimating their impact on the distribution of N_T requires, e.g., Monte Carlo methods that would repeatedly simulate the system. For example, if the number of susceptible occupants N_I has uncertain intra-day variation, then separate simulations should be run for various samples of the N_I sequence. Given space limitations, we do not go into more detail here, but simply note that the variance of model predictions can be estimated from the parameter distributions at the cost of running additional simulations.

3.1.2. Disturbance inferencing

Although the thermal model Eq. (3) is not directly needed for monitoring purposes, it can nevertheless provide valuable insight into the operation of the space. Specifically, we note that the model includes inputs Q_z and w_z for the internal sensible and moisture loads in the space. Although reasonable estimates of these values can be obtained from occupancy profiles and typical energy densities (e.g., from ASHRAE standards [55]), the character of these disturbances can vary significantly from space to space. Thus, to better predict the airflow delivered to each space, it is desirable to have a more accurate forecast of Q_z and w_z , which ultimately determine the response of the HVAC system.

To infer the values of these parameters, we note that all other terms in Eq. (3) are generally known with higher accuracy, either via direct measurements from the building or based on simple physical characteristics of the space. In particular, the derivative terms dT_z/dt and $d\omega_z/dt$ can be estimated from data through appropriate filtering. Thus, we can back-calculate values on a pointwise basis as

$$\begin{aligned} Q_z &= \rho_a c_a f_{AHU} (T_{AHU} - T_z) \\ &\quad + Q_h + k_{az} (T_a - T_z) - \rho_a c_a V \frac{dT_z}{dt} \\ w_z &= \rho_a f_{AHU} (\omega_{AHU} - \omega_z) - \rho_a V \frac{d\omega_z}{dt} \end{aligned}$$

generally with additional filtering to minimize the impact of measurement noise. Once these disturbance trajectories have been calculated, future values can be predicted through various regression techniques [61]. These predicted inputs can then be used for optimization purposes as discussed in the next section.

3.2. Optimization

When used for optimization purposes, the dynamic models in Section 2 serve to predict values of key performance metrics that would result from certain decision variables. By coupling the HVAC controller model from Section 2.3 to the thermal model in Section 2.2, we can predict the airflow and energy HVAC energy consumption for a space. The predicted airflow values can then be used as inputs to the infection-risk model from Section 2.1 to estimate infection risk. The key realization is that there is an inherent tradeoff between energy consumption and infection risk in the space: almost all HVAC actions that reduce infection risk (by providing additional clean air to the space) also result in increased energy consumption. Since there are multiple different mechanisms by which clean air can be provided, the goal of the optimization is to identify the most efficient clean-air sources, which can vary from day to day based on ambient conditions. In addition, it may be the case that some spaces already provide adequate disinfection under standard operation, and thus there is no motivation to further increase clean-air delivery. In order to address these needs, we formulate a multiobjective optimization problem using the dynamic models to predict the necessary decision metrics. Mathematically, the overall structure of the optimization problem is as follows:

$$\begin{aligned} \min_{\theta \in \Theta} & \quad (z_1, \dots, z_{K_1}) \\ \text{s.t.} & \quad z_k^{\min} \leq z_k \leq z_k^{\max}, \quad k \in \{K_1 + 1, \dots, K_2\} \\ & \quad z_k = \sum_{t=0}^{\tau} \ell_k(y_t, u_t, \theta), \quad k \in \{1, \dots, K_2\} \\ & \quad x_{t+1} = f(x_t, u_t, \theta), \quad t \in \{0, \dots, \tau - 1\} \\ & \quad y_t = h(x_t, u_t, \theta), \quad t \in \{0, \dots, \tau\} \\ & \quad x_0 \text{ given} \\ & \quad u_t \text{ given}, \quad t \in \{0, \dots, \tau\} \end{aligned} \quad (9)$$

That is, we perform multiobjective optimization over a set of model parameters $\theta \in \Theta$ with each scalar objective function and constraint z_k , calculated from the inputs u_t and outputs y_t of the simulation model. Note that in this context, the models $f(\cdot)$, $h(\cdot)$, and $\ell_k(\cdot)$ represent the *coupled* models, i.e., the combination of infections quanta, thermal, and controller models as mentioned in the previous paragraph.

Owing to the multiple objectives, the goal is not to find just a single optional point θ^* but rather the *set* of non-dominated points

$$\Theta^* = \{\theta^* \in \Theta : \text{for all } \theta \in \Theta \text{ there exists } k^* \text{ such that } z_{k^*}(\theta^*) \leq z_{k^*}(\theta)\}$$

also referred to as the ‘‘Pareto set’’. All values in this set are efficient in the sense that it is not possible to improve the value of one objective function without making a different objective function worse. The specific values to implement can then be chosen from this set in accordance with current infection-risk and energy priorities.

The proposed optimization process is illustrated in Fig. 1. As will be discussed in Section 3.2.3, we choose a *finite* set of optimization cases specified by the values of their decision variables. A simulation is run for each optimization case, holding the other model parameters and inputs constant across all cases. Within the simulation, the controller and thermal models (from Sections 2.3 and 2.2 respectively) are coupled, with the controller model determining the HVAC control actions (primarily flow rate and outdoor air fraction) based on the previous zone state (i.e., temperature and relevant disturbances) and configured setpoints. The thermal model then evolves based on those control actions, and the process is repeated until the end of the simulation period. In parallel, the HVAC control actions are passed to the infection model so that infection risk and other objectives can be evaluated. Once the objectives and constraints have been simulated for each optimization case, a Pareto search is performed to characterize the feasibility and optimality of each case. The optimal cases can then be presented the user to select from in accordance with their current operating goals. We discuss the decision variables, performance metrics, and Pareto optimization algorithm in the following sections.

3.2.1. Decision variables

For the purposes of optimization, we split the decision variables into two different groups: operational variables that can be adjusted via setpoints sent to the BMS, and design variables that would require physical equipment to be installed or replaced. The key distinction is that operational variables can be reasonably be adjusted on a daily or faster basis, whereas changes to design variables would generally be changed only monthly or less frequently. Thus, the frequency of optimization (and the time period considered in the optimization) will typically depend on the variables being adjusted. Fig. 1 illustrates the general case in which the decision variables can affect both the HVAC control setpoints and the parameters of the infection model, but often only one

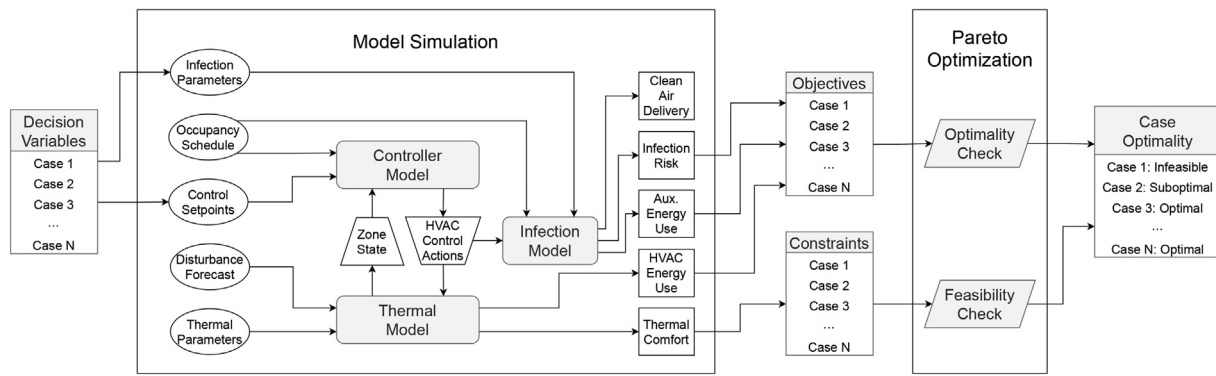


Fig. 1. Data flow diagram for optimization process. For each optimization case, the corresponding decision variables are simulated using the proposed models to calculate values of objective functions and constraints. Results are then post-processed via the Pareto Optimization algorithm to identify whether each case is infeasible, or optimal.

group (i.e., operational or design) of variables will be optimized at a time.

The primary operational variables to optimize are the supply temperature setpoint \bar{T}_{AHU}^{sp} , the minimum ventilation setpoint \bar{f}_{vent} , economizer threshold temperature \bar{T}_{econ} , and the activations \bar{a}_{UV} and $\bar{a}_{aux,j}$ for disinfection devices. Note that all of these variables are parameters of the controller model from Section 2.3. For simplicity, we assume that the operational variables are all held constant over the simulation interval (e.g., each day) so as to keep the optimization space reasonably small. Various extensions are possible that would adjust these variables more frequently (e.g., every hour), but care would be needed to account for the transient dynamics of the control layer after setpoint changes, as these effects are not included in the current models.

The primary design variables to optimize are the in-duct filter type, which implicitly adjusts η_{filter} and the fan-power coefficients c_i , and the presence of in-zone disinfection devices, which would add elements j included in the clean-air delivery and disinfection energy calculations of Eqs. (1) and (2). Of course, which values are included in the optimization space will need to account for the physical constraints of the HVAC system (e.g., certain filter types may not be compatible with the filter rack present in the duct).

3.2.2. Objective function and constraints

Once the set of decision variables has been chosen, the next step is to choose the objective functions and constraints to include in the optimization problem. For our purposes, we are interested in multiobjective optimization to evaluate the tradeoff between energy consumption and infection risk in the space. The corresponding performance metrics are the expected number of transmissions N_T (from Eq. (2)) and the total energy consumption $E = E_{HVAC} + E_{aux}$ (from Eqs. (2) and (6)). The goal of the optimization is not only to minimize these metrics separately, but also to identify the Pareto set in the two-dimensional objective space to identify the efficient set of decision variables.

Although the decision space is already bounded based on the chosen optimization ranges for the optimization variables, it is also important to ensure that selected values do not have unintended consequences for the space. For example, although a given AHU may physically be able to supply air at a temperature $T_{AHU}^{sp} = 65$ °C, this value may not provide adequate dehumidification of the supply air, which could thus lead to unacceptably high humidity within the zone. To remove such variables from the decision space, we thus include the thermal and moisture comfort metrics ϵ_T and ϵ_ω from Eq. (6) as explicit constraints in the optimization problem. Specifically, thresholds are set on both met-

rics, and if a particular set of variables leads to either threshold being violated, then that set of variables is deemed infeasible and not part of the Pareto set. Note that in extreme circumstances (e.g., if the comfort bounds T_z^{min} , T_z^{max} , ω_z^{min} , and ω_z^{max} have been chosen too tightly), then there may not be any feasible points left in the optimization space. Thus, to obtain a solution in such cases, it may be necessary to iteratively relax the constraint-violation thresholds until at least one feasible point is identified. However, for well-designed systems, such infeasibility is not expected to be an issue.

Both objective functions and constraints are calculated from the outputs of the simulation models as illustrated in Fig. 1. Each of these metrics aggregates over the simulation interval to produce a single scalar value from the time-varying model outputs. These values are then processed via the Pareto optimization algorithm as described next.

3.2.3. Pareto optimization algorithm

Although there is a wealth of literature on multiobjective optimization (and optimization in general), we propose in the interest of usability a fairly simple-minded optimization algorithm for our problem. Specifically, we choose a finite set of points $\tilde{\Theta} \subset \Theta$ that approximates the full space and then identify the Pareto set $\tilde{\Theta}^*$ from just that finite set by simulating the performance metrics for each point within that set. This choice is appropriate for the following reasons:

- The number of independent decision variables is generally between 3 and 5, which means a simple grid over the continuous variables will not give a prohibitive number of points.
- Some decision variables (e.g., filter type or on/off choice for equipment) are implicitly discrete, and thus continuous optimization methods cannot be applied in a straightforward manner.
- Enumerating the full set of points avoids the need to calculate gradient information, as optimization can be performed simply by looking at the simulated metrics for each of the finite set of points.

For other problems where these properties do not hold, it may be necessary to use more sophisticated optimization algorithms for the problem to remain tractable. However, as we will show in the examples, this strategy is suitable for the current problem.

To build the set $\tilde{\Theta}$, we simply choose a finite set of cases for each scalar decision variables and take $\tilde{\Theta}$ to be their Cartesian product. For variables that are discrete, the finite set is already defined,

while for continuous variables, we choose a minimum and maximum value along with a handful of values in between. For example, suppose our decision variables are the supply temperature setpoint \bar{T}_{AHU}^{sp} , the minimum ventilation rate \bar{f}_{vent} , and whether to activate UV disinfection \bar{a}_{UV} . We would thus define the set

$$\tilde{\Theta} = \{13, 13.5, 14, \dots, 18\} \times \{0.5, 0.6, 0.7, \dots, 1.5\} \times \{0, 1\}$$

i.e., choosing 11 values between 13 °C and 18 °C for \bar{T}_{AHU}^{sp} , 11 values between 0.5 m³/s and 1.5 m³/s for \bar{f}_{vent} , and two values for \bar{a}_{UV} corresponding to “on” and “off” respectively. Thus, to optimize over this set, we need perform $11 \times 11 \times 2 = 242$ simulations.

Once the set $\tilde{\Theta}$ has been chosen, the overall optimization algorithm is straightforward:

1. Initialize the Pareto-optimal set $\tilde{\Theta}^*$, the suboptimal set, $\tilde{\Theta}^\circ$, and the infeasible set $\tilde{\Theta}^\times$ to the empty set \emptyset .
2. For each $\theta_i \in \tilde{\Theta}$:
 - 2.1 Simulate the model using parameters θ_i to determine each $z_k(\theta_i)$.
 - 2.2 If all bounds $z_k^{min} \leq z_k(\theta_i) \leq z_k^{max}$ are satisfied for $k \in \{K_1 + 1, \dots, K_2\}$, add θ_i to $\tilde{\Theta}^*$. Otherwise, add θ_i to $\tilde{\Theta}^\times$.
3. For each $\theta_i \in \tilde{\Theta}^*$:
 - 3.1 For each $\theta_r \in \tilde{\Theta}^* \setminus \{\theta_i\}$:
 - 3.1.1 If all $z_k(\theta_r) \leq z_k(\theta_i)$ for $k \in \{1, \dots, K_1\}$, remove θ_i from $\tilde{\Theta}^*$ and add to $\tilde{\Theta}^\circ$.
4. Return the Pareto-optimal set $\tilde{\Theta}^*$, the suboptimal set $\tilde{\Theta}^\circ$, and the infeasible set $\tilde{\Theta}^\times$.

Thus, we require one model simulation for each case in $\tilde{\Theta}$ and K_1 comparisons for each pair of cases in $\tilde{\Theta}$. Note that more efficient algorithms are known [62], but this naive implementation is sufficient for the problem sizes we consider. In terms of Fig. 1, Step 2.1 is the model simulation, Step 2.2 is the feasibility check, and Step 3.1.1 is the feasibility check. The returned sets $\tilde{\Theta}^*$, $\tilde{\Theta}^\circ$, $\tilde{\Theta}^\times$ (representing the Pareto-optimal, suboptimal, and infeasible subsets) are thus disjoint and characterize each case in the original set $\tilde{\Theta}$.

Once the approximate Pareto-optimal set $\tilde{\Theta}^*$ has been determined, the only decision left is to choose a specific point to implement within the system. This selection could be made via thresholds on the variables, e.g., choosing the solution with the lowest energy consumption E such that the number of transmissions $N_T \leq 1$ if any such points exist, or otherwise choosing the solution with the lowest N_T . The exact selection criteria will depend on a building manager’s current priorities, but by providing the full Pareto set, the fundamental tradeoff can be fully evaluated. Suboptimal and infeasible points can also be explored as desired, but they generally should not be implemented.

4. Illustrative examples

In this section, we show how the proposed models and optimization framework can be used to generate insights and guide operational decisions in buildings. We split the examples into two main categories, the first using data obtained from real buildings (during pre-pandemic operation) and the second using data obtained from building simulations. For both categories, we note that the associated observations and discussion are intended as specific for the exact buildings and spaces being illustrated. In particular, calculated infection rates for a particular example space should not be interpreted as typical values for all spaces of that type in other buildings, as specific characteristics can vary significantly. Indeed, a key point of the examples is that energy and infec-

tion metrics depend on many individual factors, and thus the quantitative analysis proposed in this paper is necessary to achieve consistent positive outcomes.

4.1. Real data for a commercial building

To illustrate the potential uses of our proposed modeling and optimization framework, we start by applying the models to real building data. For these examples, we will use datasets from two zones in a real building: (1) an “Office” zone consisting of 25,000 ft² of open-plan office floorspace (with some conference rooms and small offices) served by the same AHU; and (2) a “Gym” zone consisting of 5,000 ft² of floorspace containing a weight room, an aerobics studio, and locker rooms. Both datasets come from the same commercial building in the US Midwest, and they cover the time range from August 2019 through March 2020 (after which point the building was closed due to the COVID-19 pandemic). Data was collected at a 1-min frequency and then re-sampled to 15-min intervals for our purposes. The Office zone is nominally occupied from 6am to 7 pm on weekdays and closed on weekends, while the Gym zone is occupied from 5 am to 9 pm on weekdays, 7am to 8 pm on Saturdays, and 7 am to 2 pm on Sundays.

As discussed in Section 2.1, calculation of infection risk requires values of key parameters relating to occupant activity. For the Office zone, we assume a breathing rate of 0.7 m³/h, while in the Gym zone, we use a higher breathing rate of 1.6 m³/h. These values correspond to “low” and “moderate” activity respectively for a typical adult occupant [63]. We note that there is likely to be high variance in the breathing rate for Gym due to the different activities (aerobic exercise, weightlifting, calisthenics, resting, etc.), so our chosen value is best interpreted as a time-averaged mean. For the exhaled quanta concentrations in the zones, it reasonable to suspect that Gym occupants have a higher overall exhaled particle concentration, as they are more likely to be breathing deeply and through their mouth compared to Office occupants. However, it is also plausible that actively infectious individuals would be less likely to exercise if they are feeling symptomatic, and thus we would expect Gym infectors to have below-average viral load (and also exhaled quanta concentration). Given this uncertainty and the absence of definitive data, we use the same mean exhaled quanta concentration of 20 q/m³ in both spaces, corresponding to light vocalization [4].

Finally, the models require a time-varying occupancy profile, which is unfortunately not provided by the dataset. Thus, for both zones, we use a surrogate profile obtained by rescaling default schedules from ASHRAE [55] using design occupancy, which is taken as 5 and 10 occupants per 1000 ft² respectively [54]. Zone areas and volumes are taken from floorplan diagrams, and we assume MERV8 filters with no occupants wearing masks, as is consistent with pre-pandemic operation.

4.1.1. Data trends and metrics

To begin analysis, we start by examining the daily trends inherent in the datasets. The key BMS measurements used in the infection-risk model are the outdoor air ventilation (which contributes directly to clean-air delivery) and the total supply air flow (which provides clean-air via filtration). To illustrate daily and weekly trends, Fig. 2 shows values of these variables along with the assumed time-varying occupancy profiles. For the Office zone, we see that total supply air flow is reasonably constant throughout the occupied period on a given day, but that the amount of flow provided does change from week to week, likely corresponding to seasonal changes. The provided ventilation varies much more significantly, perhaps due to the action of the economizer, but it still covers a fairly narrow range. By contrast, we see that there is much greater variation in the Gym zone, both within a given

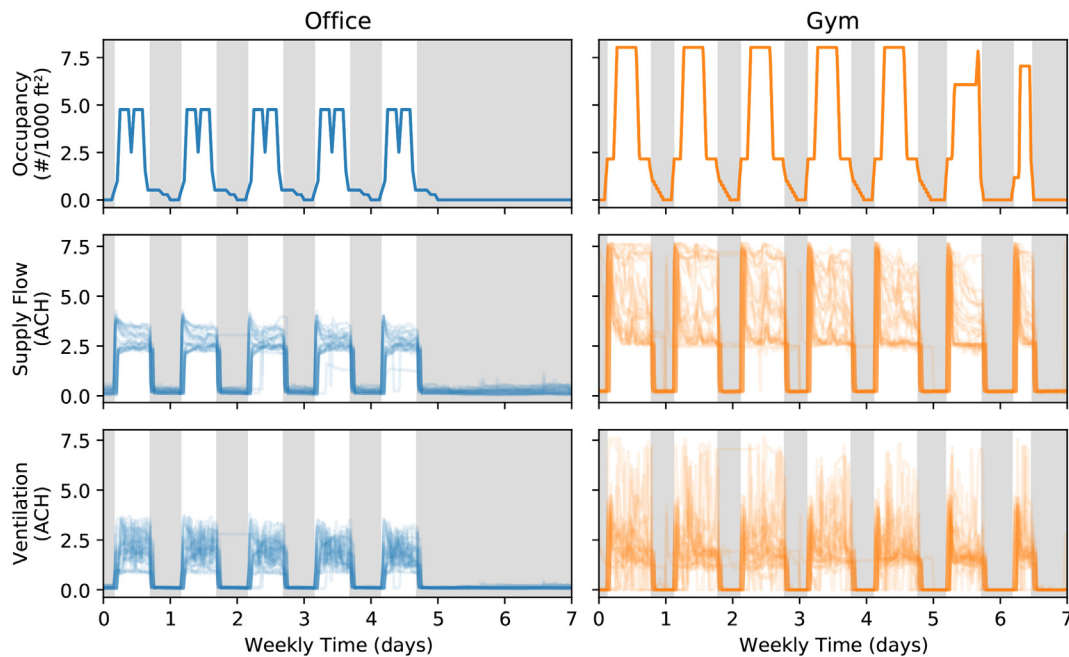


Fig. 2. Weekly flow data and occupancy profiles for the real building zones. Plots show 26 overlaid weekly profiles (with time zero corresponding to midnight on Monday). Shaded regions correspond to unoccupied hours when the HVAC system is deactivated.

day and from week to week. We almost always see a large spike in airflow at the beginning of the occupied period, which indicates that the zone heats up overnight and thus the temperature controllers take aggressive action as soon as they are activated in the morning. There are also some other flow spikes likely associated with extra occupancy around lunchtime and at the end of the workday (but unfortunately the occupancy profile does not reflect those features). We do see that the ventilation varies much more significantly, which indicates the presence of measurement noise in the sensor but could also be caused by operation of the economizer. Consistent with ASHRAE standards, we see higher (normalized) airflows in the Gym zone compared to Office. We note also that in accordance with the operation of the building, there is zero occupancy in the Office zone on the weekends, but the Gym zone remains open. Thus, the HVAC system is inactive in the Office zone but active in the Gym zone during these days.

To assess the infection risk present in these zones, we take the data from the previous plot and simulate the dynamic infectious-particle models from Section 2.1 to calculate average clean-air delivery rate and expected transmissions for each day of operation. Given the uncertainty in the exhaled quanta concentration, we also show a shaded confidence interval corresponding to exhaled quanta concentrations between 5 q/m^3 and 80 q/m^3 (with the points assuming the mean value of 20 q/m^3). These values are plotted in Fig. 3. From these results, we can make three key observations. First, we note that infection risk in both zones is lower on weekends than on weekdays despite similar or significantly lower clean-air delivery rates. This behavior simply mirrors the occupancy cycles of the respective spaces, with lower weekend occupancy for the Gym and zero weekend occupancy for the Office. Second, we note that although the clean-air delivery to the Gym zone is almost always higher than that of the Office zone, the average infection risk for the zone is nearly 10 times higher than that of the Office zone. This discrepancy is explained by the occupancy differences in the two spaces: in the Gym, there is a higher density of people who are all breathing at elevated rates. (Indeed, the Gym has a 2.3 times higher breathing rate and roughly a 1.5 times

higher occupant density, which suggests a $2.3^2 \times 1.5 \approx 7.9$ times higher infection rate, all other effects equal.) Thus, although the HVAC system removes more of the particles, there is an inherent higher infection risk compared to the Office zone where occupants are quietly seated at their desks. Third, we note that clean-air delivery rate is generally higher in warmer months and lower in colder months. This behavior is consistent with the fact that supply air flows are elevated in the cooling season to mitigate additional heat loads, whereas in the heating season, flows are lower and may sit at their configured minimum bounds. Under steady-state conditions, transmission rate is inversely proportional to clean-air delivery rate, and thus the observed 40% decrease in clean-air delivery for the Gym zone roughly corresponds to the observed 65% increase in transmissions when comparing the winter months to the summer months.

Before moving on, we briefly discuss the uncertainty associated with the two variables. For the clean-air delivery rate, the only uncertain terms are the passive decay rate for the infectious particles (due to both deactivation and deposition) and the filtration efficiency for the in-duct filter. Because both effects exhibit particle-size dependence, their actual effect will depend slightly on the specific particle distribution produced by the infector. However, these values are reasonably well characterized, and their range of variation is generally small. The remaining terms in the clean-air calculation are flow measurements taken directly from the BMS. Although flow sensors can be very noisy, we expect only minuscule error when averaged over the course of a full operating day. Thus, we plot the clean-air delivery without an uncertainty region. By contrast, the infection rate depends linearly on the exhaled quanta concentration, which is known to vary by up to three orders of magnitude depending on vocalization level and mode of respiration, from nasal breathing to speaking and singing [4]. Accordingly, we have shown an uncertainty range corresponding to a factor of 4 above or below the chosen mean value. A key takeaway is that, although the baseline infection risk in the Gym is inherently higher than in the Office, this analysis does *not* suggest that Gym zones are always highly dangerous or that Office

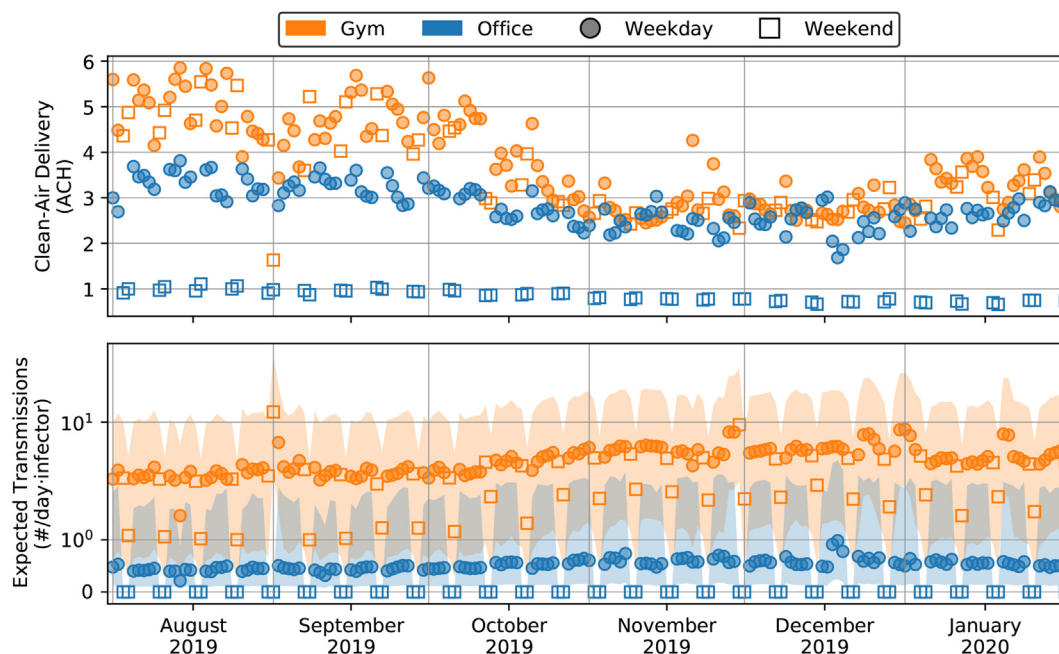


Fig. 3. Daily average clean-air delivery and expected transmissions for the real building zones. Each point reflects the average value of the metric over a 24-h simulation. For transmissions, points are calculated assuming a mean exhaled quanta concentration of 20 q/m^3 , while the shaded region shows the resulting range for exhaled quanta concentrations between 5 q/m^3 and 80 q/m^3 .

zones are always a safe haven. Instead, it is necessary to examine both types of spaces on a case-by-case basis to determine the correct course of action.

4.1.2. Prediction and optimization

We turn now to showing how these models can be used for the purposes of prediction and optimization. In this use case, we wish to predict the operation of the HVAC system one day in advance, which means we do not have access to flow measurements from the BMS and must predict them using the controller and zone-temperature models. Thus, we start by showing example day-ahead predictions of clean-air delivery using the models from Section 2.

Important inputs to the thermal models include the internal heat and moisture disturbances. To predict these values in advance, we use the following procedure: back-calculate disturbances on the prior 7 days of data as discussed in Section 3.1.2, filter that trajectory by retaining only the most dominant Fourier modes in the frequency domain, use the first day of that trajectory as the disturbance for the simulation. This procedure implicitly assumes that successive weeks follow approximately the same disturbance trajectories. Certainly more sophisticated disturbance-prediction algorithms are possible, but even this extremely simple strategy produces adequate results on the test data. In Fig. 4 we show day-ahead clean-air delivery predictions for each zone under summer and winter conditions. Even though these predictions are far from perfect, they nevertheless capture the general daily trends in airflow and thus produce sufficiently accurate infection-risk predictions. As a result, the models are suitable for use in optimization.

To close this example, we present some sample optimization results using the data just presented. In each of the four cases, the optimization period covers 24 h of operation (corresponding to the first day from each case in Fig. 4). The decision variables are the supply temperature setpoint (sampled at 25 points between the bounds of $55 \text{ }^\circ\text{F}$ and $65 \text{ }^\circ\text{F}$) and whether or not to activate 2 ACH of supplemental in-zone filtration (provided by stan-

dalone filtration devices) during occupied hours. The two primary objective functions are energy cost and expected transmissions, and the constraints are maintenance of zone temperature and humidity within comfortable ranges during occupied hours. Following the optimization strategy presented in Section 3.2, each case thus runs 50 simulations (corresponding to the different values of the decision variables) of the dynamic models from Section 2, and then the Pareto search described in Section 3.2.3 is applied to categorize each solution. These results are shown in Fig. 5. Note that in both cases, the number of transmissions is determined using the mean exhaled quanta concentration of 20 q/m^3 . The effect of uncertainty in this parameter would be to uniformly scale the calculated value up or down, which does not affect the optimality characterization of the points.

From the sample optimization results, we see that both decision variables decrease the airborne transmission rate and generally increase energy consumption as their values are increased. For supply temperature setpoint, there are some regions where energy cost is reduced with increasing supply temperature (which corresponds to conditions where the reduced latent load is larger than the increased fan power), but they usually cannot be exploited. A key observation is that in the summer conditions, it is not possible to increase the supply temperature setpoint all the way to the established upper bound, as the reduction in de-humidification at the AHU cooling coil eventually leads to violation of the humidity comfort constraint. By contrast, during winter, the full range of supply temperature setpoints is available, as the outdoor air does not need to be de-humidified. In all cases, when choosing a course of action, only Pareto optimal solutions should be considered, since feasible but suboptimal solutions consume unnecessary energy to achieve a given infection rate (or vice versa). Thus, as priority shifts from energy to infection risk, the general strategy is to increase supply temperature up to a certain point (which depends on weather conditions) and then activate the supplemental in-zone filtration units (while dropping the supply temperature back to its minimum value) before increasing the supply temperature setpoint once again to achieve maximum clean-air delivery. Overall,

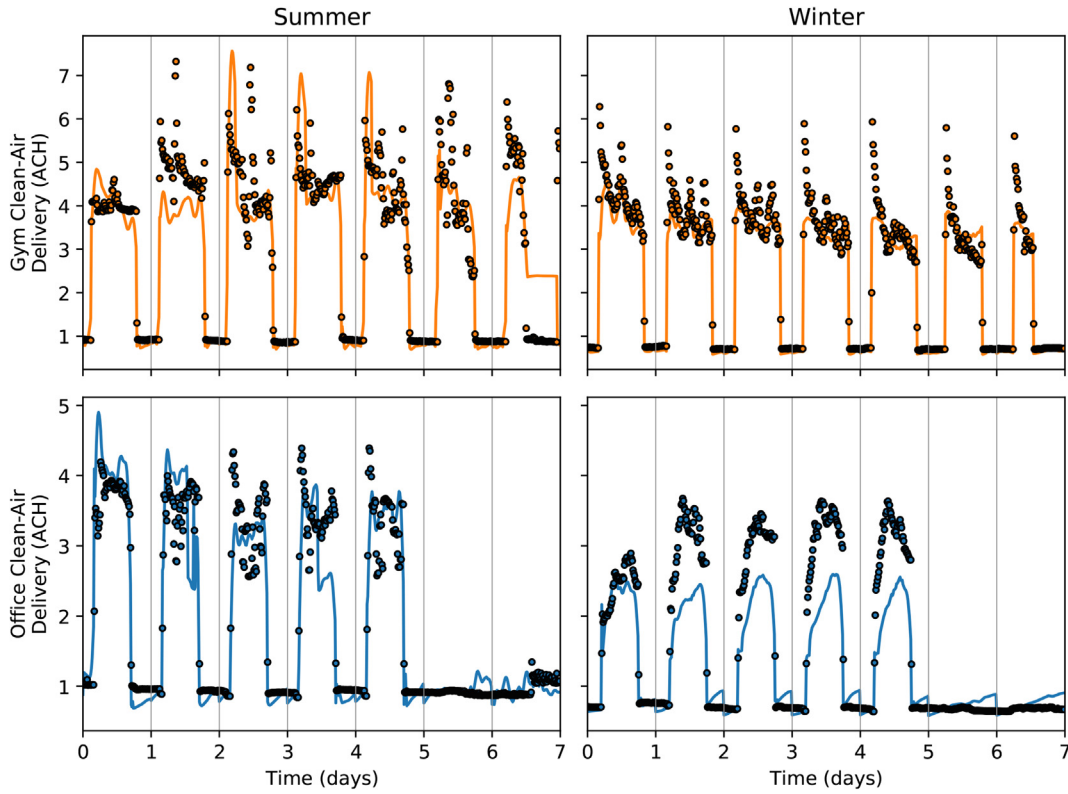


Fig. 4. Day-ahead clean-air delivery predictions using the full dynamic models. Points show values computed from the actual BMS data, while lines show values predicted by the model. Data for Summer are taken from August 2019, while data for Winter are taken from February 2020.

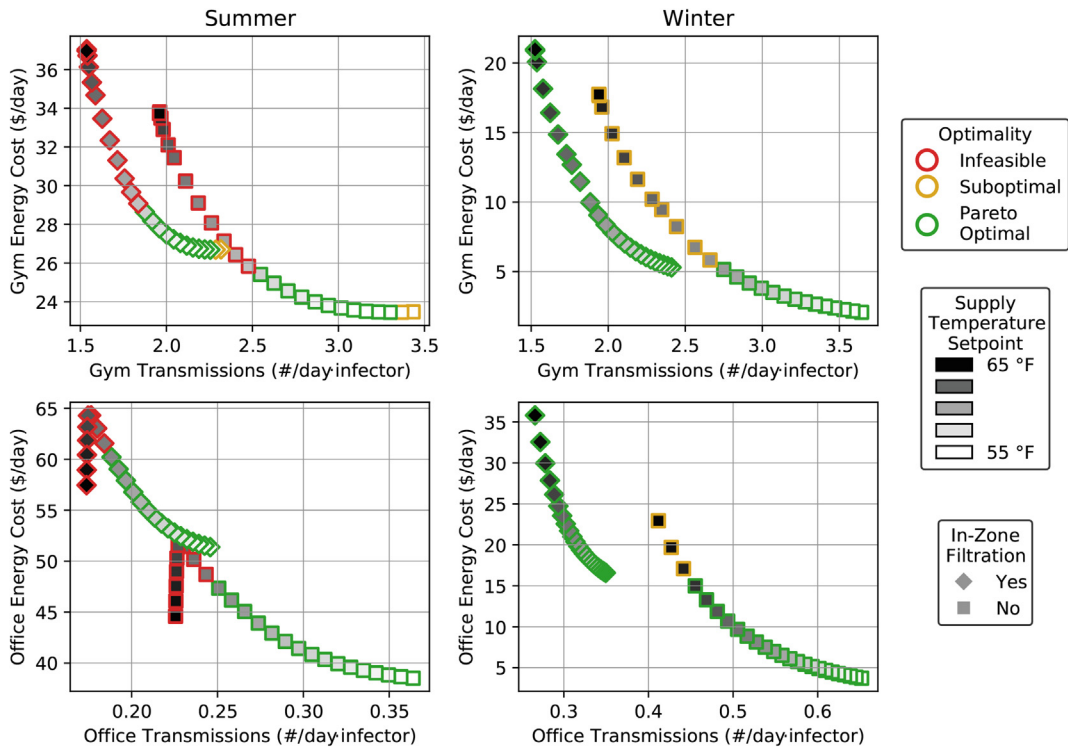


Fig. 5. Optimization results for the real building zones under two different weather conditions. Disturbance and weather data for each case are taken from the first day of the corresponding case in Fig. 4.

we see that even with just these two decision variables being considered, the range of achievable infection rates (and the most energy-efficient strategy to achieve them) depends on both space characteristics and outdoor conditions. Therefore, it is important to update operational decisions regularly to account for changes to these variables.

4.2. EnergyPlus simulations

In our second set of examples, we use simulated building data produced from EnergyPlus [49] using typical meteorological year (TMY) weather data. Although the EnergyPlus simulated data is not as realistic as the real data in the previous section (in particular it is much less noisy and is subject to much more regular disturbance profiles), the key benefit is that we can directly compare the effects of different control actions. Specifically, for a given day's weather, we can simulate *all* candidate sets operational variables rather than just the default operating policy. We will use this feature to more thoroughly assess the available control space and the accuracy of our models.

For this purpose, we consider the first floor of the "Large Office" reference building provided by EnergyPlus, which consists of 38,000 ft² of floorspace served by the same AHU with a design occupancy of 192 occupants. HVAC equipment in the building is auto-sized by EnergyPlus using the chosen TMY weather data. We assume the same occupant breathing rate of 0.7 m³/h and exhaled quanta concentration of 20 q/m³ as in the previous example. Similarly, we assume MERV8 in-duct filters and that occupants are not wearing masks. For optimization variables, we allow the minimum ventilation rate to be varied from its default level (chosen by EnergyPlus in accordance with ASHRAE standards) to double that value, while we allow the supply temperature setpoint to be adjusted from its default of 55 °F up to 65 °F.

4.2.1. Metric trends

To begin the simulation analysis, we examine the values of our two main objective functions and how they vary as the operational variables are changed. For each of the four extreme points in the optimization space, we simulate a full year of the building operating with those variables using EnergyPlus to obtain the required BMS data. We then use that BMS data with the proposed models to calculate daily values for energy cost and infection rate. We perform this process for both a cold climate (TMY data from Chicago, IL) and a hot climate (TMY data from Houston, TX). The results of these simulations are shown in Fig. 6. For the purposes of energy accounting, we assume cooling is provided via electricity at a COP of 3, while heating is provided by gas at 90% efficiency. We use an electricity price of 0.12 \$/kW.h and a gas price of 0.80 \$/therm. For brevity, we do not show thermal discomfort, although all temperature comfort violations are below 4 °F·h/day, and humidity comfort violations are generally confined to the winter with indoor conditions being too dry. To illustrate the relative efficiency of the two operational variables for reducing infection risk, we also plot the cost of extra clean air. This value is calculated as

$$\text{Extra Clean Air Cost} = \frac{\text{Energy Cost} - \text{Baseline Energy Cost}}{\text{Clean Air Delivery} - \text{Baseline Clean Air Delivery}}$$

which gives the per-unit cost of the extra clean air provided relative to baseline operation. Lower values of this metric indicate a more efficient source of clean air.

From these results, we see similar behavior as in the previous real-data examples, but with new trends emerging from the new optimization variable. First, we note once again that infection rate is generally higher in the winter months and lower in the summer months, simply due to variation in the average airflow provided by the HVAC system. We note that this effect is much more pro-

nounced in the Cold climate compared to Hot due to the weather-dependent cooling loads (and resulting airflow differences). Second, we see that increased ventilation almost always requires additional energy consumption, and the magnitude is much higher in particularly hot or cold weather. The explanation for this effect is that because the building is equipped with an economizer using an aggressive suitability threshold, the system will automatically increase ventilation above its minimum value when energetically favorable. Thus, any additional ventilation that results from increasing the minimum ventilation setpoint necessarily requires an energy penalty. By contrast, increasing the supply temperature setpoint can potentially *reduce* energy consumption, and so it tends to be much less costly. Finally, we note that the HVAC system by itself has a somewhat narrow range to reduce the infection rate. Comparing the maximum and minimum infection rates on each day, we see less than a factor of two reduction. Thus for spaces with unacceptably high infection rate, it may be necessary to install standalone-filtration or other disinfection devices to achieve infection-risk goals.

For the purpose of decreasing infection risk, we see that increasing the supply temperature setpoint is less effective than increasing the ventilation rate in the heating season but more effective in other seasons. When choosing between the two, we generally prefer the variable with the lower cost of extra clean air. The bottom plot in Fig. 6 shows that the extra clean air provided by increasing the supply temperature setpoint costs next to nothing throughout much of the year, but during the heating season, its cost can jump to more than \$200/ACH. This observation is consistent with the fact that the HVAC systems often operate near minimum-flow constraints during the heating season due to the small cooling load, and thus adjusting the supply temperature setpoint has little to no effect on clean-air delivery but a large effect on energy cost (since extra heating is required to meet the higher supply temperature). However, adjusting the supply temperature setpoint in the cooling season can have a significant impact on clean-air delivery, with only a minor change to energy consumption. By contrast, Fig. 6 shows that doubling the ventilation rate yields a more consistent cost between \$50/ACH and \$100/ACH for extra clean air throughout the year, which tracks with the changing weather conditions throughout the year. Overall, this analysis illustrates that the two operational variables can have very different effects on energy cost and infection risk throughout the year, which motivates the proposed optimization approach to ensure that the most efficient operational variables are selected based on current weather conditions.

4.2.2. Tradeoff curves

Given the trends observed in the previous section, we now look at the optimization tradeoffs in more detail. For each of the simulation cases, we choose a representative day for each season and simulate our dynamic models using the appropriate parameters and weather data. Specifically, we grid over the optimization space (supply temperature between 55 °F and 65 °F; minimum ventilation rate between 1 and 2 times the default value), simulate one day of HVAC operation (using the proposed models from Section 2) for each grid point, and plot the resulting energy consumptions and infection rates. These results are shown in Figs. 7. A key observation is that the shapes of the mapped optimization sets (and thus the Pareto-optimal points within those sets) vary significantly based on the weather, and thus optimal operation changes as well.

Before moving on, we highlight some of the key features of these mappings. First, for the cases that come to a sharp point in the upper left corner, it is evident that the value of the minimum ventilation rate becomes nearly degenerate at high values of the supply temperature setpoint. The explanation for this behavior is that the higher supply temperatures result in increased *total* sup-

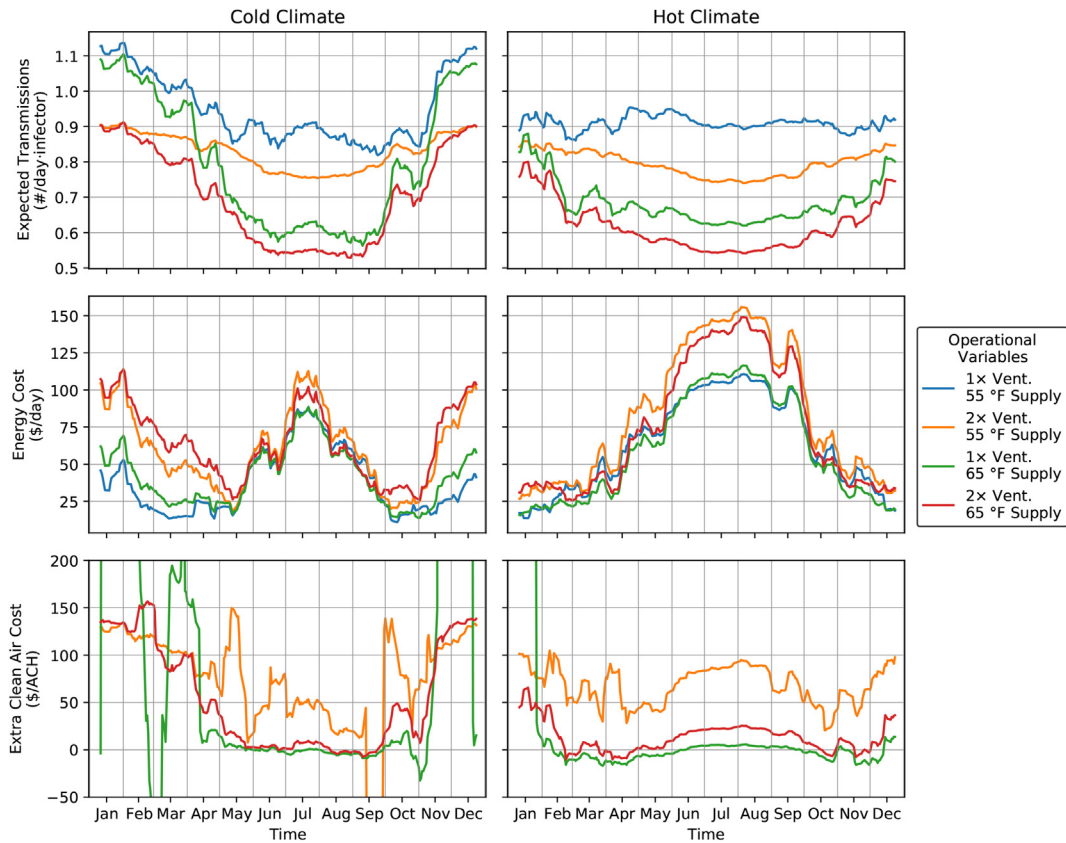


Fig. 6. Daily metrics calculated from EnergyPlus simulation data. Lines show ten-day rolling averages for non-holiday weekdays. Values of operational variables are the four extreme points in the optimization set. Extra clean air cost is calculated relative to the baseline operation with both operational variables at their minimum value.

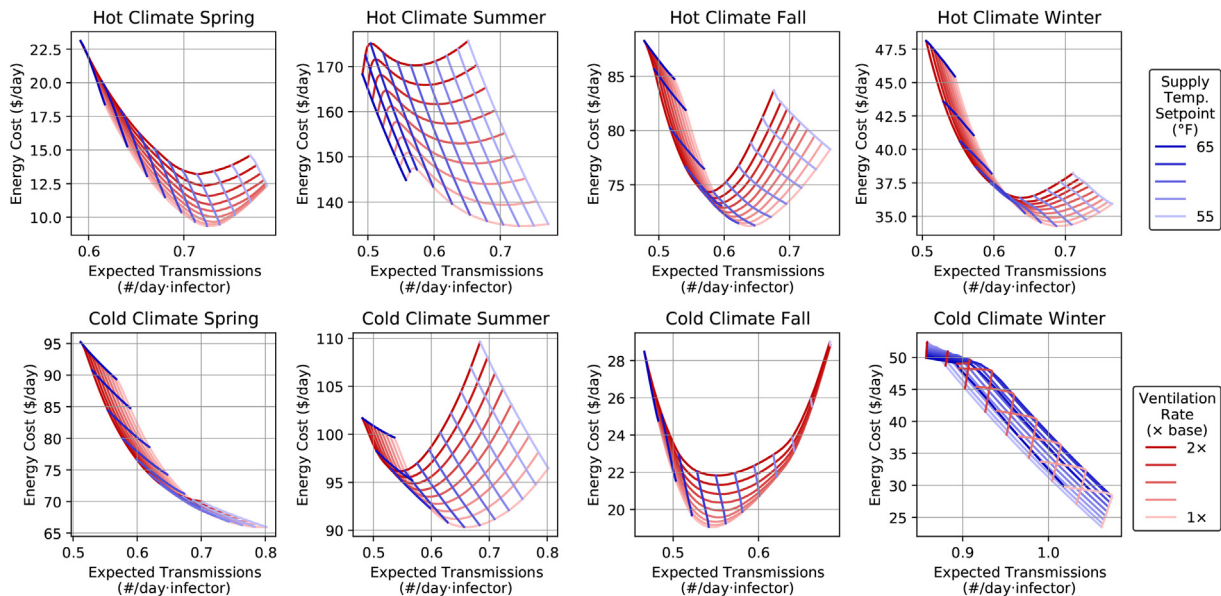


Fig. 7. Example optimization tradeoffs for different seasons and cities. Plots show how the two-dimension optimization range gets mapped onto the to primary objective functions. Red lines hold ventilation rate constant and vary supply temperature setpoint, while blue lines hold supply temperature constant and vary ventilation rate.

ply airflow, and so the economizer will naturally increase outdoor-air flow (to maintain roughly a constant outdoor-air fraction). Thus, increasing the *minimum* ventilation rate essentially has no effect because the economizer is already providing more than that amount of ventilation. Second, we note that increasing the ventilation rate uniformly leads to an increase in energy use. As observed in the previous examples, this behavior is due to the aggressive

economizer on this building: any time extra ventilation is energetically favorable, it will be supplied automatically by the economizer; thus, any increased ventilation that results from changing the minimum setpoint is necessarily at an energy penalty. We note that the use of energy-recovery ventilation systems could significantly reduce this cost [45], but installation of such devices would be costly and perhaps not worthwhile solely for reducing infection

risk. Third, we see that increasing the supply temperature setpoint can decrease energy cost at first but then eventually begins to increase it. This is generally due to the tradeoff between latent cooling load (which decreases at higher supply temperatures) and fan power (which increases at higher flows caused by higher supply temperatures). Because both of these effects are nonlinear, their resulting sum can take on the wide variety of shapes observed in the figure depending on ambient and internal conditions. Finally, we note that the unique behavior observed in the Cold Climate Winter case is due to the fact that the system is operating at minimum-flow conditions. Thus, increasing the supply temperature setpoint above its minimum bound initially does not change the flow rate (since it is still at minimum) but requires additional energy consumption to heat the air to that temperature. Only after a sufficiently high supply setpoint does airflow rise above its minimum bound and we see more typical behavior.

Overall, we see in these example cases that the Pareto-optimal set generally but not always contains points with the lowest ventilation rate and sometimes points with the highest supply temperature setpoint. Thus, as priorities shift from energy minimization (e.g., during non-pandemic periods) to infection reduction (e.g., during a pandemic), the optimal policy is to first increase supply temperature setpoint and maybe increase the minimum ventilation rate after that. We note from the relative flatness of the sets in Fig. 8, that many ambient conditions make it possible to achieve significant reduction in expected transmissions with only minor increase in energy cost, and thus for a properly optimized system, energy consumption can remain nearly constant across pandemic and non-pandemic conditions. However, for systems with different HVAC configurations, different behavior might be optimal, and so it is important to consider the specifics of each application.

4.2.3. Optimization accuracy

To close the examples, we briefly discuss the accuracy of the proposed framework with respect to whether the optimized actions deliver the benefits predicted by the models. For this purpose, we choose representative weeks for each climate and season (corresponding to the eight cases in Fig. 7) and solve the Pareto optimization for each weekday. Note that this Pareto search accounts for thermal comfort constraints, and so on some days, there may be very few feasible points. From the set of Pareto-optimal solutions, we choose three representative points: the point

that minimizes energy cost, the point that minimizes infection rate, and a point midway between them. For each of those points, we then run a new EnergyPlus simulation using the corresponding setpoints to see what the BMS data would be for those setpoints. We then calculate metrics from that data and compare to the metrics predicted by the models in the optimization problem. We also include values for the default operation of the building with both decision variables at their lower bounds.

A key challenge with this approach is that due to the simplistic disturbance forecast used (which essentially just reuses the disturbances from one week ago), it is likely that the model predictions may be systematically off due to week-to-week changes in disturbances (primarily due to solar radiation and heat exchange with the ambient). However, because the model predictions are being used for optimization, any linear re-scaling does not affect the location of the optimal solutions. Thus, to avoid the poor disturbance forecast influencing accuracy metrics, we re-scale the model's predictions to match the actual data using one single scale factor for the entire week. This procedure will correct for the effect of the disturbance schedule, but any other model deficiencies will still appear in the plots. These results are shown in Figs. 9 and 10, for the Cold and Hot climates respectively.

From these results, we see that after correcting for the poor disturbance forecast, the models accurately predict the resulting metrics for each representative solution. Although there are some cases where accuracy is not satisfactory (in particular, the winter energy cost in Fig. 9 and the spring clean-air delivery in Fig. 10), the dynamic models do capture the available operating space for the chosen representative solutions. We note in particular that the Balanced solution often achieves noticeable reduction in infection rate with minimal (if any) impact on energy consumption. Furthermore, because the scale correction does not affect which solutions are Pareto-optimal, even the low-quality disturbance forecast used in this example would be suitable for online optimization. Thus, we see that despite the relative simplicity of the proposed models, they can nevertheless capture the key behavior of real buildings and ultimately lead to actionable insights for building operators.

5. Conclusions and future directions

In this paper, we have presented a modeling and multiobjective optimization framework to assess the tradeoff between airborne

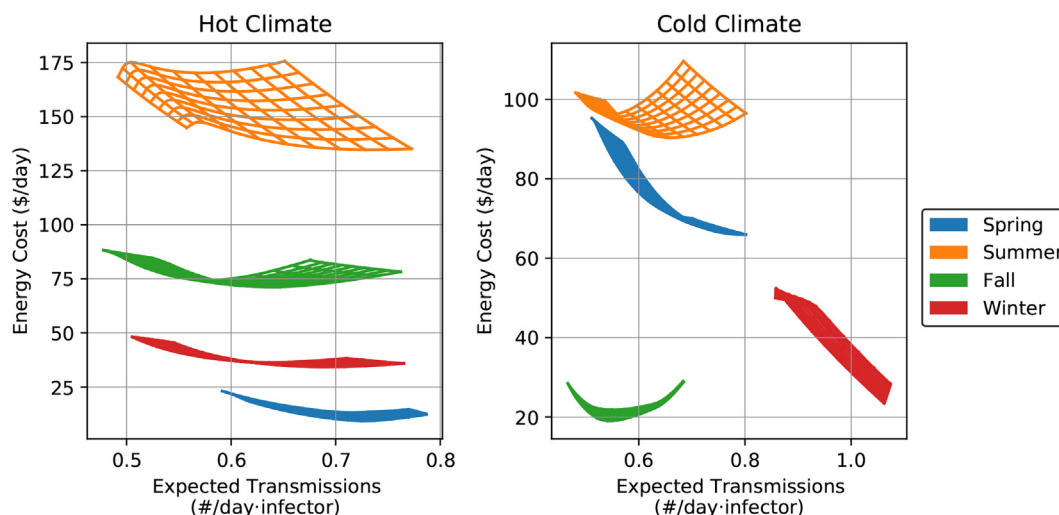


Fig. 8. Example optimization tradeoffs for different seasons and cities. Plots show the same values as in Fig. 7 but with all seasons plotted together for each climate.

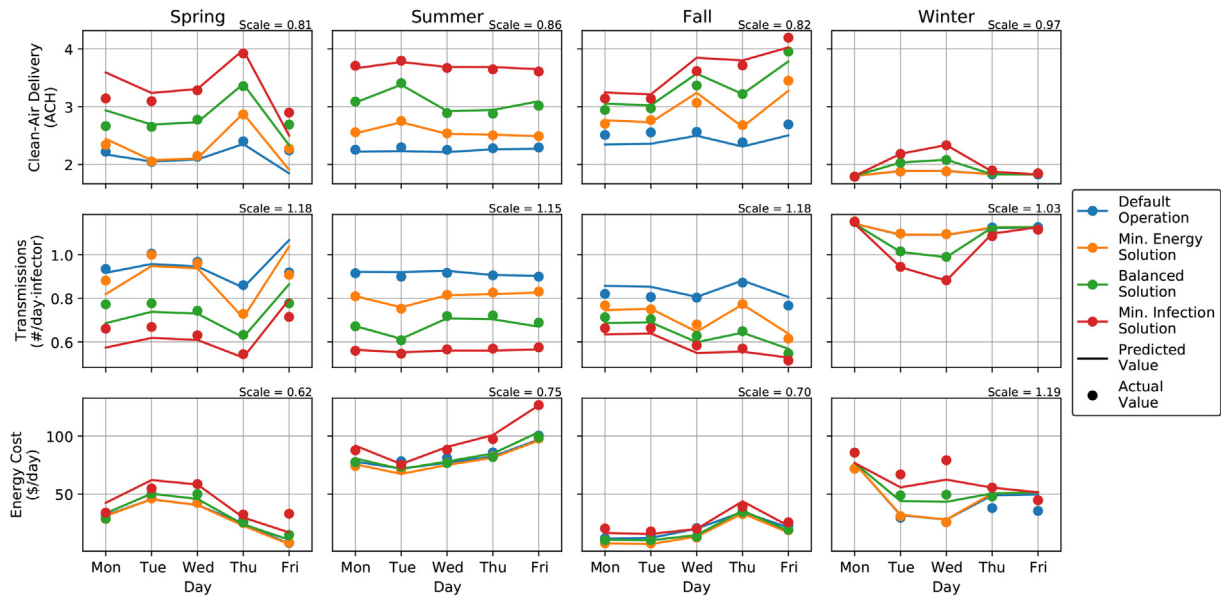


Fig. 9. Predicted and actual optimization metrics for the Cold climate. Colors correspond to different representative points on the Pareto surface. Actual values are taken from EnergyPlus simulations using the optimized decision variables, while predicted values are metrics predicted by the dynamic models (and scaled by the indicated scale factors).




Fig. 10. Predicted and actual optimization metrics for the Hot climate. Plots are as in Fig. 9.

disease transmission risk and energy consumption associated with the operation of building HVAC systems. Using a physics-based dynamic model for the hypothetical concentration of airborne infectious particles that would be produced by an infector in the space, the risk of airborne transmission can be estimated via the expected exposure of susceptible occupants to those particles. By accounting for filtration, ventilation, and other infectious-particle removal sources, this model can be used with real building data and space parameters to assess the infection risk in buildings as operated. Furthermore, by coupling to additional models for space temperature and humidity as well as the action of the HVAC regulatory control system, it is possible to predict the airflow that will be delivered to the space, from which infection risk and energy consumption can be estimated. This overall model can then be

embedded within a model-based optimization routine to determine Pareto-optimal values for operational setpoints and design variables. These optimization results thus provide quantitative guidance to building managers who can then make decisions consistent with their current health and sustainability goals.

To illustrate the insights provided by the proposed framework, we have shown a series of examples that apply the proposed methodology to real and simulated building data. Using BMS measurements from a real office space and gym, we show how clean-air delivery and infection rates can be estimated for each space. We then show that the proposed models can predict the future behavior of the HVAC system with sufficient accuracy to be used for optimization purposes. Finally, we have shown computed Pareto sets for the supply temperature setpoint and activation of sup-

plemental in-zone filtration devices, illustrating that both variables can reduce infection risk but unfortunately require additional energy consumption. Using EnergyPlus simulations, we have shown how infection risk and energy consumption vary throughout the year for a typical office building for different values of the supply temperature setpoint and minimum ventilation rate. We have then presented example Pareto tradeoff curves showing that increasing the supply temperature setpoint is often the most energy-efficient way to reduce infection risk, but also that the associated energy penalties vary significantly with weather conditions. We have closed by validating the proposed models' predictions for both baseline and optimized operating setpoints. Together, these examples illustrate that although the proposed models are relatively simple, they nevertheless capture the key building trends and can thus be a valuable source of insight for building operators.

In the future, we hope to extend the proposed framework to incorporate new indoor-air quality data sources (e.g., CO₂ and PM sensors) to generalize optimization goals. For example, the excess CO₂ concentration from occupants' respiration can be correlated with airborne disease transmission risk [38,42,64] and thus provides a rich additional datastream to calibrate the model and improve the optimization accuracy. Even when airborne disease transmission is no longer a major concern, the average CO₂ concentration could be optimized instead, with the primary tradeoff being between the increased energy consumption associated with ventilation and the wellness benefits to occupants associated with better indoor-air quality [45]. In this context, new design and operational variables could be introduced (e.g., whether to install energy-recovery ventilation and whether to activate it on a certain day) and optimized to account for specific space and climate characteristics. Such functionality would help support a paradigm shift to combat indoor respiratory infection and improve occupant wellness in the aftermath of the COVID-19 pandemic [5]. Overall, we hope that model-based optimization will increasingly be used as a tool to guide design and operation of building HVAC systems to improve health and sustainability.

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 research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors. Real building datasets were provided by Johnson Controls International. The authors would like to thank Michael Wenzel, Mohammad EIB-sat, Anas Alanqar, and Chenlu Zhang for their helpful discussions.

Appendix A. Infection model discussion

A.1. Clean-air delivery calculation

To model the time variation in infection quanta concentrations, we note that infectious particles are released into the air by the infectious individual and then removed by a variety of mechanisms. To standardize modeling of these effects, we define the concept of "clean-air delivery" to the space as follows: for a piece of equipment or physical process j that deactivates or captures infectious particles at a total rate r_j (measured in units of q/s) from an air stream with concentration C (measured in units of q/m³), the corresponding clean-air delivery rate is defined as

$$f_{\text{clean},j} = \frac{r_j}{C}$$

which thus has units m³/s corresponding to a hypothetical volumetric flow (and can be converted to air-change rates by dividing by the space volume). We use this definition because for most disinfection processes, r_j is directly proportional to C , and thus $f_{\text{clean},j}$ can be expressed as a simple fraction of airflow through the device. In addition, this formulation provides a common basis for comparison of different disinfection devices: multiple devices providing the same clean-air delivery rate thus become fungible from an infection-risk perspective and can be assessed in terms of energy consumption or other considerations to choose the best among them.

One complication of this clean-air framework is that for devices and processes that operate at the particulate level (i.e., deposition and filtration), the removal rate r_j depends on the size distribution of the particles containing the infectious material [4]. When the inlet size distribution is known, accounting for size-dependent effects is straightforward for a single device. However, because their action leads to changes in the airborne size distribution, it is not longer possible to fully and exactly describe the clean-air characteristics of a device without knowing (and accounting for the inherent recirculation effects on) the size distribution of airborne particles. In the interest of simplicity, we assume a constant steady-state particle size distribution (using the size bins defined by ASHRAE standards) and calculate removal rates accordingly. From previous studies for influenza, Azimi and Stephens [37] states that 20% of viral material is found within E1 particles (diameters 0.3 to 1 μm), 29% within E2 particles (diameters between 1 and 3 μm), and the remaining 51% within E3 particles (diameters between 3 and 10 μm). Assuming these values hold for the pathogens of interest, an effective clean air delivery rate can be calculated via a weighted average of the removal rates for a representative particle diameter within each size bin. However, see discussion in Section A.5 for how this effect could be modeled more accurately without significant changes to the existing framework.

Aside from size-dependent effects, the calculation of the clean-air delivery rate for most processes and devices is straightforward. For example, airborne particles naturally deposit on the ground at a given velocity $v_{\text{deposition}}$ due to gravity. Thus, in a room with volume V and height h , the removal rate of such particles by deposition is given by

$$r_{\text{deposition}} = \frac{1}{h} v_{\text{deposition}} VC$$

Similarly, due to natural decay of the viruses (assumed to occur with a first-order rate constant k_{decay}), the removal rate due to natural decay is thus given by

$$r_{\text{decay}} = k_{\text{decay}} VC$$

In practice, this decay rate may vary with temperature and humidity within the space [4], but for simplicity here we assume it is truly constant. As modeled above, we combine these two effects as "natural" disinfection sources in Eq. (1) with

$$f_{\text{clean,natural}} = k_{\text{natural}} V, \quad k_{\text{natural}} = k_{\text{decay}} + \frac{1}{h} v_{\text{deposition}}$$

Note that they scale with the total volume of the space and occur independently of any HVAC operational decisions.

For standalone filtration devices, the clean-air calculation is also straightforward: assuming a total volumetric flow f_{filter} through the filter, which captures a fraction η_{filter} of the particles the total removal rate is given by

$$r_{\text{filter}} = \eta_{\text{filter}} f_{\text{filter}} C$$

and so the clean-air delivery is simply

$$f_{\text{clean,filter}} = \eta_{\text{filter}} f_{\text{filter}}$$

Effective filtration efficiencies for various filter grades are known [44], modulo uncertainty in size dependence. Similarly, for UV devices that kill a given fraction η_{UV} of the microbes passing through them, the clean air production is

$$f_{\text{clean,UV}} = \eta_{\text{UV}} f_{\text{filter}}$$

If properly designed, UV disinfection tunnels can achieve very high disinfection rates $\eta_{\text{uv}} \approx 1$ [28], but the performance of smaller standalone devices may vary.

For devices operating in parallel, their total clean-air delivery is a simple sum of their individual values. Since all auxiliary in-zone devices are effectively in parallel, we thus include the summation term in Eq. (1), noting that the *maximum* clean-air capacity $f_{\text{aux},j}^{\text{max}}$ of a given auxiliary device needs to be scaled by the extent $a_{\text{aux},j}$ to which it is currently active. Note that for some devices $a_{\text{aux},j}$ may vary continuously, while for others, it may take only a binary value indicating whether the device is on or off.

For devices operating directly in series, however, the total clean-air delivery must account for the fact that the clean-air delivery for each device must be calculated in terms of the original inlet concentration rather than that device's particular inlet concentration. The canonical example of this configuration is in the AHU: when air is returned to the AHU, a given fraction x_a is vented to the ambient (and replaced with a corresponding volume of outdoor air), after which it passes through a filter with disinfection efficiency η_{filter} and possibly also through a UV kill tunnel that kills a fraction $a_{\text{uv}}\eta_{\text{UV}}$ of microbes in the airstream. Taken individually, the clean-air rates of each device would be given as

$$\begin{aligned} f_{\text{clean,vent}} &= x_a f_{\text{AHU}} \\ f_{\text{clean,filter}} &= \eta_{\text{filter}} f_{\text{AHU}} \\ f_{\text{clean,UV}} &= a_{\text{uv}} \eta_{\text{UV}} f_{\text{AHU}} \end{aligned}$$

However, because these devices operate in series, we must account for the fact that each device processes the outlet of the previous device. Thus, the inlet concentration to the filter is only $1 - x_a$ times its initial value, and thus the filter provides only $(1 - x_a)\eta_{\text{filter}}$ fractional disinfection. The overall clean-air delivery for the AHU is thus given by

$$f_{\text{clean,AHU}} = \eta_{\text{AHU}} f_{\text{AHU}}, \quad \eta_{\text{AHU}} = 1 - (1 - x_a)(1 - \eta_{\text{filter}})(1 - a_{\text{uv}}\eta_{\text{UV}})$$

consistent with the definition in Eq. (1). We can see immediately from the functional form that $f_{\text{clean,AHU}} \leq f_{\text{AHU}}$ which makes intuitive sense: the AHU cannot clean more air than passes through it.

A.2. Infection probability and received dose

To estimate the risk of infection in a zone, we calculate a hypothetical number of transmissions during the course of the evaluation time period (commonly on a daily basis) as given by Eq. (2). Desired limits can be placed on this value to guide operational and HVAC decisions. The key question then becomes how the time-varying trajectory of C resulting from the hypothetical infectors leads to infection risk for susceptible occupants in the space.

In the ideal case, infection risk would be calculated on an individual basis and then aggregated across the building to determine overall infection risk. Using the dynamic model Eq. (1) for the airborne infectious quanta concentration C in each zone, we can thus estimate an individual's received dose throughout the day and ultimately compute that individual's infection probability assuming a constant hazard rate (expressed in units of q^{-1}) for new infection. For a susceptible individual s , we have

$$P_s = 1 - \exp\left(-\int_0^\tau \beta_q q_s(t) dt\right)$$

with P_s giving that individual's infection probability in terms of the individual's infectious dose rate q_s , which is itself given by

$$q_s = (1 - \mu_s) f_{\text{breathe},s} C$$

We note that at steady-state conditions (i.e., $dC/dt = 0$), the model in Eq. (1) gives $C = q_i / f_{\text{clean}}$. Substituting these values (and converting the time-varying integral to a simple multiplication) gives

$$P_s = 1 - \exp\left(-\frac{\beta_q (1 - \mu_s) q_i \tau f_{\text{breathe},s}}{f_{\text{clean}}}\right)$$

which is equivalent to the common steady-state presentation of the Wells-Riley equation [4,37,40]. However, in variable-volume HVAC systems, the steady-state assumption is rarely satisfied, as the regulatory control layer constantly adjusts supply airflow f_{AHU} to maintain temperature within the space in response to variation in weather and occupancy. Thus, the resulting quanta concentration fluctuates throughout the day and should be modeled as a time-varying quantity to more accurately represent infection risk.

After each individual's infection probability has been calculated, these values can then be aggregated over all building occupants to find the average number of transmissions given by

$$N_T = \sum_s P_s$$

with distributional information available if needed. This value can of course be converted to a *rate* by dividing by the total evaluation time τ . Unfortunately, this person-based approach is not practical for implementation in real buildings. First, it implicitly requires that the infection quanta concentration be modeled in every space within the building, which may not be worthwhile for low-occupancy or otherwise low-risk zones. Second, even if every space were modeled, it is generally not possible to track each individual's movement throughout the building. Third, since we are interested in the HVAC operation of the zones, it is more practical to adopt a zone-based approach in which infection risk is associated with spaces rather than individuals.

To address these limitations, we take an approach similar to Bazant and Bush [4] and make a linear approximation as follows:

$$\begin{aligned} N_T &= \sum_s 1 - \exp\left(-\int_0^\tau \beta_q q_s(t) dt\right) \\ &\approx \sum_s \int_0^\tau \beta_q q_s(t) dt \\ &\approx \int_0^\tau \beta_q \sum_s q_s(t) dt \end{aligned}$$

Noting that Eq. (1) has already defined $q_s = \sum_s q_s(t)$ in the model, substituting this value gives the equation for N_T presented in Eq. (1). The main benefit of this formulation is that we can now split the whole-building R on a zone-by-zone basis. Thus, separate actions can be taken in each zone based on that zone's contribution to overall infection risk.

In arriving at this formulation, we have made use of the fact that $1 - \exp(-x) \approx x$ for small values of x , which allows us to remove the nonlinearity and shift the summation inside the integral. We note also that the linear approximation is conservative, i.e., uniformly *over-estimating* N_T , due to concavity of the exponential function. Assuming that the infection risk is fairly well controlled, buildings should be operating in the range where this approximation is highly accurate, at least over the course of one day. Expressed differently, if we are relying on the nonlinearity of the infection probability curve as a source of infection risk reduction, then we have already failed at infection control. Thus, the linear

approximation is always reasonable when modeling the mitigation of airborne transmission risk.

A.3. Key parameters

The infection quanta model and infection rate calculation Eqs. (1) and (2) include a number of key parameters related to infector and occupant behavior. Specifically, these quantities are the breathing rates $f_{\text{breath},I}$ and $f_{\text{breath},S}$ of infectors and susceptibles, the mask effectivenesses μ_I and μ_S , and finally the infector exhaled quanta concentration N_I . We discuss these parameters in more detail as follows.

The breathing-rate parameters f_{breath} are relatively straightforward to estimate, although their values are somewhat uncertain. The EPA provides values [63, Table 6-2] for average breathing rates based on age and activity level. For a 40-year-old individual, these values range from 0.276 m³/hr for sedentary activity to 2.94 m/hr for high-intensity activity. For typical office or commercial environments, a default value of 0.67 m³/hr is reasonable, but this value should be adjusted for specific spaces with different activity levels (e.g., gyms) and possibly also for different age distributions (e.g., in an elementary school). Note that because we generally do not know which specific individuals in the space are infectious, we will almost always assume $f_{\text{breath},I} \approx f_{\text{breath},S}$. However, the values can be considered separately in edge cases where distinct separate values are known.

For the mask effectiveness parameters μ_I and μ_S , unfortunately there is very large uncertainty relating to both mask material and fit. For example, an N95 mask may have a theoretical filtration efficiency above 99%, but if worn incorrectly (so that a significant amount of air can escape around the edges of the mask), then its effective efficiency may be significantly lower. Given this variability, we use a default value of roughly 70% filtration efficiency. As with the breathing-rate parameters, we generally do not know which specific individuals are infectious, and so there may not seem to be clear motivation to use different values for μ_I and μ_S . However, when aerosols are exhaled by humans, they generally contain a large amount of water that will evaporate as the particles come to equilibrium with the humidity in the space. As this evaporation occurs, the particles shrink in diameter by roughly a factor of 2 at 50% relative humidity [4]. Thus, we generally expect that $\mu_I > \mu_S$ since the particles are bigger when exhaled through the infector's mask than when they are inhaled through the susceptibles' masks. Although this effect is challenging to quantify, we retain separate μ_I and μ_S in case accurate experimental values are obtained in the future.

Finally, for the infector's exhaled quanta concentration C_I , we note that this value is commonly characterized in conjunction with the breathing rate $f_{\text{breath},I}$ as a single parameter q (denoted q_I in our equations above) giving a quanta generation rate in quanta per time [37,40]. However, given that COVID-19 is transmitted primarily in respiratory aerosols, it is suggested by Bazant and Bush [4] that the value be disaggregated as separate C_I and $f_{\text{breath},I}$ as we have modeled it, with C_I ranging from 10 q/m³ for quiet breathing, to q/m³ while speaking, and as high as 1000 q/m³ while singing. This choice of parameterization makes it clear that the dependence of infection risk (as determined by the total received quanta dose) is roughly *quadratic* in the average breathing rate. Of course, the true generation rate of infectious quanta is complicated by additional factors (e.g., intermittent coughing and variation in viral load due to natural disease progression), but this parameterization at least captures the fundamental trends with respect to activity level and vocalization.

A.4. Susceptible and infectious occupants

Two of the most important inputs in Eq. (1) are the numbers of infectious and susceptible occupants N_I and N_S . Given that these

values cannot be directly measured, we must *choose* appropriate values based on typical occupancy profiles and other considerations.

For the infectors N_I , one option is to choose the value based on the incidence of disease in the general population. Such statistics are generally available from the CDC and other organizations and can be multiplied by the design occupancy of a zone. Unfortunately, these values can vary significantly with time and location, which makes them difficult to incorporate into quantitative analysis. Thus, we instead opt for a conditional approach as follows: rather than trying to estimate the *actual* number of infectors N_I in the space, we instead *assume* that $N_I = 1$ and thus calculate infection risk conditioned on the fact that one infectious individual has made it into the space. (Note that to account for the typical occupancy cycles, we do not assume a constant $N_I \equiv 1$, but rather set $N_I = 1$ only during occupied hours and leave $N_I = 0$ during unoccupied hours.) With this assumption, the infection rate N_T from Eq. (2) thus scales linearly in the number of infectors (i.e., 2 infectors in the space will on average cause $2N_T$ infections over the course of a day, 3 will cause $3N_T$, etc.). We know from the mathematics that this approach will *overestimate* the number of new infections when the number of infectors is larger (since the susceptible population is correspondingly reduced), but if the situation is already that dire, HVAC measures alone will not be sufficient to prevent disease transmission, and the prudent choice is to simply close the building. The main point is that the $N_I = 1$ assumption decouples our simulations from external conditions while still providing valuable insight into the inherent risk posed by each zone. In addition, dividing the number of transmissions N_T by the average number of infectors $N_I = 1$ thus gives an estimate of the *reproductive number* for the space, i.e., the number of additional transmissions caused by each infector. To avoid runaway transmission of disease, buildings should thus be operated so that $R \leq 1$ (although due to uncertainty in other parameters, a considerable safety margin is likely warranted).

Finally, for the number of susceptibles N_S , the main consideration is to account for possible immunity in some or all building occupants. The conservative approach is of course to assume no occupants are immune, and thus N_S is simply the time-varying number of occupants in the space (less the chosen number of infectors N_I). Where vaccination rates are known, N_S can be further reduced by the fraction of occupants who are vaccinated, perhaps attenuating slightly to account for the fact that immunity is not absolute. For buildings where full time-varying occupancy profiles are not known, representative profiles are available in ASHRAE Standard 90.1 [55] that can be scaled by the design occupancy for each zone. Alternatively, time-varying CO₂ measurements can also be used to estimate occupancy patterns [65–67].

A.5. Size-resolved infection model

As alluded to in Section A.1, the size dependence of many infectious particle removal mechanisms means that the size distribution of airborne particles will vary over time based on the size-dependent balance of particle generation (by infectious occupants) and removal (via clean air delivery). The approach proposed previously is essentially to assume that the size distribution of infectious particles within the zone is equal to the size distribution as exhaled by infectious occupants. Unfortunately, this assumption may not hold if clean-air sources with strongly size-dependent efficiency are used. To more accurately model this effect, the best modification is to simulate a separate instance of Eq. (1) for a representative set of particle diameters. For example, using the ASHRAE-defined size bins, one would simulate separate models for C^{E1} , C^{E2} , and C^{E3} giving the infectious particle concentration

(still measured in q/m^3) only for particles within the E1, E2, and E3 size bins. All model parameters that exhibit size dependence (including k_{natural} , $f_{\text{aux},j}^{\text{max}}$, η_{filter} , μ_l , μ_s , and C_l) can then take on their appropriate values for particles within that range.

To calculate the expected transmission rate in this size-resolved model, the only modification is that exposure should now aggregate over the various size bins. Specifically, Eq. (2) should be modified to calculate

$$N_T = \int_0^\tau \beta_q \sum_{e \in \{E1, E2, E3\}} q_s^e(t) dt$$

in which the $q_s^e(t)$ are the size-resolved exposure rates calculated from the separate simulations of Eq. (1) for each size bin e . In the interest of simplicity, we have chosen not to adopt this model for the results presented in this paper, but its application is relatively straightforward provided that the number of size bins is kept reasonably small.

Appendix B. Thermal and energy model discussion

B.1. Temperature and humidity models

In addition to the net advective transfer provided by the supply and return air streams, the temperature model in Eq. (3) considers conductive heat transfer with the ambient proportional to the temperature difference $k_{az}(T_a - T_z)$, and all other sources of heat gain (e.g., plug, lighting, and metabolic loads), are assumed to be lumped into the Q_z input. For simplicity, we do not explicitly model the temperature dynamics of solid-mass component (as operating temperatures are relatively constant, and thus temperature-dependence of this effect can be largely neglected), and we do not explicitly calculate the effect of solar radiation (as accurate solar forecasts are often not available). However, these effects can be added to the model where relevant data is available. Separate zones within a building are assumed to be physically thermally isolated, and so no exchange (mass or energy) is modeled. Direct infiltration from the ambient is not considered, while exfiltration is implicitly included in the return-air stream. For zones that require heating (e.g., exterior zones in cold climates), we include the Q_h term as a direct addition of heat to remain agnostic to the heating source. This supplemental heating may be supplied by reheat coils inside the VAV boxes or by in-zone electric heaters. In any case, since it operates at the room level, we can accurately model it as a direct addition of heat.

As with the temperature model, the humidity model in Eq. (3) considers advective transfer from the supply air and an internal generation term (due primarily to occupants). However, since infiltration from the ambient is not considered, there is no direct exchange term between ω_z and ω_a . For robustness, the value of ω_z could be checked (and clipped appropriately) to ensure that it does not exceed the saturation point $\omega_{\text{sat}}(T_z)$. However, since such conditions should only arise during unrealistic operating scenarios (e.g., if the cooling coil is disabled during hot and humid ambient conditions), this check can generally be skipped. In the interest of brevity, we do not consider humidification equipment in the current model. Where present and active, it would be added as an additional generation term in the moisture model for in-zone equipment or in the AHU coil model for AHU equipment.

B.2. AHU coil model

We assume that the AHU controls its supply air temperature to setpoint using a heating coil and a cooling coil in sequence. Air passing through the AHU is thus heated or cooled depending on

whether its temperature is below or above the setpoint temperature $T_{\text{AHU}}^{\text{sp}}$. The previous Eq. (4) accounts for the action of the appropriate coil. In the interest of simplicity, we present these equations assuming unbounded cooling capacities (which means the setpoint can always be met), as total loads are implicitly constrained by bounds on flow. However, extension to directly account for finite coil capacity is straightforward.

To start, we calculate the AHU inlet conditions by mixing the outdoor-air and recirculating air streams. A straightforward energy balance gives

$$\begin{aligned} \omega_{\text{in}} &= x_a \omega_a + (1 - x_a) \omega_z \\ T_{\text{in}} &= T(x_a \mathcal{H}(T_a, \omega_a) + (1 - x_a) \mathcal{H}(T_z, \omega_z), \omega_{\text{in}}) \end{aligned}$$

in which $\mathcal{H}(T, \omega)$ gives the specific enthalpy of an air stream at temperature T and humidity ω , while $\mathcal{T}(h, \omega)$ gives the temperature of an air stream at specific enthalpy h and humidity ω . From here, if $T_{\text{in}} \leq T_{\text{AHU}}^{\text{sp}}$, the air is heated directly to the setpoint with no change in humidity. Thus, the AHU outlet conditions are $T_{\text{out}} = T_{\text{AHU}}^{\text{sp}}$ and $\omega_{\text{out}} = \omega_{\text{in}}$. Alternatively, if $T_{\text{in}} > T_{\text{AHU}}^{\text{sp}}$, the air is cooled as discussed next.

Because of the possibility of condensation, the cooling coil requires some additional modeling to determine the outlet humidity. For this purpose, we follow the common ‘‘contact mixture analogy’’ (e.g., as in Seem and House [59]). With this formulation, we assume that a fixed fraction of air (x_{bypass}) bypasses the cooling coil completely, while the remaining fraction ($1 - x_{\text{bypass}}$) equilibrates with the coil. As before, we assume that the cooling coil meets the setpoint, which gives $T_{\text{out}} = T_{\text{AHU}}^{\text{sp}}$. From this value, we can back-calculate the coil-equilibrating conditions as

$$\begin{aligned} T_{\text{coil}} &= \frac{T_{\text{out}} - x_{\text{bypass}} T_{\text{in}}}{1 - x_{\text{bypass}}} \\ \omega_{\text{coil}} &= \min(\omega_{\text{in}}, \omega_{\text{sat}}(T_{\text{coil}})) \end{aligned}$$

in which $\omega_{\text{sat}}(T)$ gives the saturation humidity at temperature T . This term is responsible for modeling the dehumidification provided by the coil, which is often a significant fraction of energy consumption in humid climates. Note that the first equation assumes linear mixing for the recombination of coil and bypass streams. We calculate the outlet humidity as

$$\omega_{\text{out}} = x_{\text{bypass}} \omega_{\text{in}} + (1 - x_{\text{bypass}}) \omega_{\text{coil}}$$

which will always be less than or equal to the inlet humidity ω_{in} .

Summarizing the heating and cooling models, we thus have the following relationships:

$$\begin{aligned} T_{\text{out}} &= T_{\text{AHU}}^{\text{sp}} \\ \omega_{\text{out}} &= \begin{cases} x_{\text{bypass}} \omega_{\text{in}} + (1 - x_{\text{bypass}}) \omega_{\text{coil}} & \text{if } T_{\text{in}} > T_{\text{AHU}}^{\text{sp}} \\ \omega_{\text{in}} & \text{else} \end{cases} \end{aligned}$$

From these conditions, we can thus calculate the thermal loads on the coils as follows:

$$\Delta h_{\text{coil}} = \mathcal{H}(T_{\text{out}}, \omega_{\text{out}}) - \mathcal{H}(T_{\text{in}}, \omega_{\text{in}}) - (\omega_{\text{out}} - \omega_{\text{in}}) \mathcal{H}_w(T_{\text{out}})$$

in which $\mathcal{H}(T, \omega)$ is the specific enthalpy of wet air at temperature as before, and $\mathcal{H}_w(T)$ is the specific enthalpy of liquid water. Note that this expression is written so that it gives the heat gained by the air stream, with a positive value indicating the stream was heated and a negative value indicating the stream was cooled. We can thus split into individual heating and cooling components as

$$\begin{aligned} \Delta h_{\text{coil}}^h &= \max(\Delta h_{\text{coil}}, 0) \\ \Delta h_{\text{coil}}^c &= \min(\Delta h_{\text{coil}}, 0) \end{aligned}$$

The calculated T_{out} , ω_{out} , Δh_{coil}^h , and Δh_{coil}^c then become the AHU conditions T_{AHU} , ω_{AHU} , Δh_{AHU}^h , and Δh_{AHU}^c indicated in Eq. (4). As above, we do not consider humidification equipment, but where present, it could be accounted for by increasing ω_{AHU} up to a given setpoint.

B.3. Fan power model

The fan power model in Eq. (5) defines electricity consumption of the fan as a cubic function of airflow f . Physically, the energy consumed by the fan is given by

$$w_{fan} = \frac{1}{\eta_{fan}} f \Delta P_{fan}$$

where η_{fan} is the mechanical efficiency of the fan, and ΔP_{fan} is the pressure rise across the fan. The cubic dependence on f comes from the fact that the pressure rise of the fan is equal to the pressure drop through the ductwork, which is commonly modeled as proportional to the square of flow. These considerations suggest that $w_{fan} \propto f^3$. In practice, the mechanical efficiency of the fan will vary over its operating range, and thus the cubic term alone is not sufficient to match the actual performance of the fan. By adding the additional linear and quadratic terms, this variation can be captured. Note that some sources will suggest adding an additional constant term to better match experimental behavior. However, we require for our application that w_{fan} is zero when f is zero (e.g., during unoccupied hours when the HVAC system is shut off), and so we do not include this term to avoid artificial inflation of predicted electricity consumption.

The most accurate way to determine the coefficients c_i in the fan model is to obtain a dataset of (f, w_{fan}) pairs from the actual fan and use standard regression. Unfortunately, it is uncommon to directly measure w_{fan} in real buildings, and so this approach is not practical. Instead, we define one set of coefficients \tilde{c}_i to match a representative AHU fan with efficiency $\tilde{\eta}_{fan}$ and pressure rise $\Delta \tilde{P}_{fan}$. We then rescale the \tilde{c}_i linearly based on the design values for η_{fan} and ΔP_{fan} (which are generally readily available) for the each AHU we model. This approach gives reasonable accuracy for a variety of zones and buildings.

A key requirement for the fan model is that it be able to account for the different electricity consumption associated with different filter types. In general, higher-efficiency filters require higher pressures to deliver the same airflow. This additional pressure must be overcome by the fan, which in turn increases electricity consumption. To model this effect, we simply rescale the coefficients to match the expected fan power at maximum flow f^{max} . From manufacturer's data, we know the pressure drop ΔP_{filter} associated with each filter type. We then assume that the design fan pressure rise ΔP_{fan} corresponds to a baseline filter type with pressure drop $\Delta \tilde{P}_{filter}$. Combined with the scaling for design pressure drop, we thus have the model

$$w_{fan} = \frac{\tilde{\eta}_{fan}}{\eta_{fan}} \frac{\Delta P_{fan} + \Delta P_{filter} - \Delta \tilde{P}_{filter}}{\Delta \tilde{P}_{fan}} \left(\tilde{c}_1 \frac{f}{f^{max}} + \tilde{c}_2 \left(\frac{f}{f^{max}} \right)^2 + \tilde{c}_3 \left(\frac{f}{f^{max}} \right)^3 \right)$$

Merging the prefactors with the representative \tilde{c}_i coefficients thus gives the specific c_i coefficients consistent with Eq. (5).

Appendix C. Controller model discussion

C.1. Zone temperature control

As shown above in Eq. (7) the controller model chooses values of f_{AHU} and Q_h by clipping the hypothetical values \tilde{f}_{AHU} and \tilde{Q}_h that would move zone temperature to the respective cooling and heat-

ing bounds T_z^{max} and T_z^{min} . Based on system structure, it is natural to first calculate f_{AHU} (as it has a nonzero lower bound due to the minimum outdoor-air ventilation constraint) assuming $Q_h = 0$ and then calculate Q_h second to provide the necessary supplemental heating. The key question is how to calculate the hypothetical values \tilde{f}_{AHU} and \tilde{Q}_h .

The ideal strategy to calculate the hypothetical variables is to implicitly solve (the discretization) of the thermal ODE model from Eq. (3). Forward simulation of this model defines the successor value of T_z as a function of the current T_z , the chosen f_{AHU} or Q_h , and the values of the remaining model inputs. Thus, we can solve those equations for the \tilde{f}_{AHU} and \tilde{Q}_h that drive T_z to T_z^{max} and T_z^{min} respectively, i.e., solving

$$\begin{aligned} \tilde{f}_{AHU} & \text{ satisfies } T_z^{max} = T_z + \int_0^\Delta \frac{dT_z}{dt} (\tilde{f}_{AHU}, 0, \dots) dt \\ \tilde{Q}_h & \text{ satisfies } T_z^{min} = T_z + \int_0^\Delta \frac{dT_z}{dt} (f_{AHU}, \tilde{Q}_h, \dots) dt \end{aligned}$$

in which the “...” represent the remaining model inputs (which would not all be known in a real system but are known to the model). Note that for physical reasons, both variables must be non-negative, but since that restriction is enforced by clipping after the fact, we place no sign restrictions on \tilde{f}_{AHU} and \tilde{Q}_h . From the model structure, we know that a suitable value \tilde{Q}_h always exists and a suitable \tilde{f}_{AHU} exists under standard operating conditions (i.e., $T_{AHU} < T_z$).

Unfortunately, because f_{AHU} enters the ODE in a product with the state variable T_z , an exact solution to the first equation requires iterative methods. The use of such algorithms can lead to slow and inconsistent solution times, and so we would like to avoid them. Thus, we instead opt for an approximate solution that is non-iterative. To determine \tilde{f}_{AHU} we assume that the $\rho_a c_a f_{AHU} (T_{AHU} - T_z)$ term in Eq. (3) can be approximated by a constant value Q_c over the simulation interval. By removing the $f_{AHU} \times T_z$ nonlinearity, our unknown variable now enters the ODE linearly and thus can be solved for directly using the exact discretization of

$$\rho_a c_a V \frac{dT_z}{dt} = k_{az} (T_a - T_z) + Q_z + Q_c$$

with the condition that $T_z = T_z^{max}$ after one timestep. With the corresponding value \tilde{Q}_c , we thus back-calculate

$$\tilde{f}_{AHU} \approx \frac{\tilde{Q}_c}{\rho_a c_a (T_{AHU} - T_z)} \quad (10)$$

In the event that the denominator is zero, we set \tilde{f}_{AHU} to $\pm\infty$ depending on the calculated sign of \tilde{Q}_c . In cases where T_{AHU} and T_z are constant during the integration interval, this calculation is exact, and the accuracy deteriorates as T_{AHU} and/or T_z vary within the interval. By assumption, T_{AHU} will be equal to its setpoint T_{AHU}^{sp} , and so there is little to no approximation error there. Similarly, during occupied hours, T_z will generally be near the cooling setpoint T_z^{max} or heating setpoint T_z^{min} depending on the season. However, during transitions at the beginning of the occupied period, T_z can change very quickly as the controller tries to restore setpoint, and so it is generally a good idea to replace T_z with an average value $(T_z + T_z^{max})/2$ to avoid overshoot. In any case, as long as the zone temperature setpoint stays reasonably constant throughout the occupied period, this calculation is sufficient to accurately model airflow to the space.

Note that it is in Eq. (10) where the primary impact of the supply-temperature setpoint T_{AHU}^{sp} is considered. Specifically, when T_{AHU} is higher and thus closer in value to T_z , we see that the

denominator becomes smaller, and so the delivered airflow \tilde{f}_{AHU} increases. A similar effect would also be possible by decreasing T_z (via its setpoint T_z^{\max}), but due to the potential impact on occupant comfort, we do not directly consider that as a manipulated variable. In any case, the key observation is that the impact of T_{AHU} on airflow (and thus also on infectious-particle concentration) is implicit due to the action of temperature control in the space.

After clipping \tilde{f}_{AHU} to find the true f_{AHU} as in Eq. (7), we can thus determine \tilde{Q}_h via the discretization of

$$\rho_a c_a V \frac{dT_z}{dt} = \rho_a c_a f (T_{AHU} - T_z) + k_{az}(T_a - T_z) + Q_z + Q_h$$

to arrive at $T_z = T_z^{\min}$. Here, we can solve for \tilde{Q}_h using the exact linear formulas as for \tilde{Q}_c above. The resulting value is then clipped to Q_h as in Eq. (7).

C.2. Outdoor air and disinfection device control

The outdoor-air fraction x_a is chosen in accordance with standard economizer logic: if the outdoor air temperature is colder than the pre-configured threshold T_a^{econ} (which means that the outdoor air is a free source of cooling), then the controller should use the maximum outdoor-air fraction x_a^{\max} ; otherwise, if the air is hotter than T_a^{econ} (which means the air is *not* a free source of cooling), then the controller should use the minimum x_a^{\min} . Calculation of these bounds is described next. The remaining quantities a_{UV} and $a_{aux,j}$ associated with disinfection devices are simply cycled on and off in accordance with the occupancy flag σ .

The two outdoor-air fraction bounds x_a^{\min} and x_a^{\max} are calculated as shown in Eq. (8). For x_a^{\min} , its value is simply the minimum fraction of f_{AHU} that will deliver the required minimum outdoor-air ventilation rate f_{vent}^{\min} (which cycles between zero and the pre-configured bound f_{vent} based on the occupancy flag σ). As discussed in the previous section, we always have that $f_{AHU} \geq f_{\text{vent}}^{\min}$, and so we know that $x_a^{\min} = f_{\text{vent}}^{\min}/f_{AHU}$ will always be between 0 and 1. For x_a^{\max} , the bound is calculated as the maximum fraction of f_{AHU} that will keep the temperature of the resulting mixture below the supply temperature setpoint T_{AHU}^{sp} (so as to avoid an unnecessary heating load). Assuming linear mixing, this condition is exactly satisfied when

$$T_{AHU}^{\text{sp}} = x_a^{\max} T_a + (1 - x_a^{\max}) T_z \Rightarrow x_a^{\max} = \frac{T_z - T_{AHU}^{\text{sp}}}{T_z - T_a}$$

which gives a valid $0 \leq x_a^{\max} \leq 1$ whenever $T_a < T_{AHU}^{\text{sp}} < T_z$. However, because those constraints are not always satisfied (in particular due to seasonal variation in T_a), we modify the calculation to clip both temperature differences to zero, i.e.,

$$x_a^{\max} = \frac{\max(T_z - T_{AHU}^{\text{sp}}, 0)}{\max(T_z - T_a, 0)}$$

with the convention that $x_a^{\max} = 1$ in the singular case $T_z = T_a$. With these modifications, the calculation gives $x_a^{\max} \geq 1$ whenever $T_a \geq T_{AHU}^{\text{sp}}$ and $x_a^{\max} < 1$ whenever $T_a < T_{AHU}^{\text{sp}}$ (i.e., when outdoor conditions are too cold for 100% outdoor air). This value is then clipped to be between the lower bound x_a^{\min} and the upper bound 1 to give the calculation in Eq. (8).

Note that the models here assume that a simple (but standard) outdoor-air economizer is responsible for outdoor-air control. However, if different equipment is present, these models may have to be modified. For example, if demand-controlled ventilation systems are present (which modulate outdoor-air flow to control indoor CO₂ concentration to a given setpoint), then it will be nec-

essary to model the CO₂ concentration in the space (with an additional transient mass balance similar to the ω_z model) and use logic similar to the f_{AHU} controller model to determine x_a . However, in the interest of simplicity, we do not consider such cases here.

References

- [1] L. Morawska, D.K. Milton, It is time to address airborne transmission of COVID-19, *Clinical Infectious Diseases* 71 (2020) 2311–2313.
- [2] S. Asadi, A.S. Wexler, C.D. Cappa, S. Barreda, N.M. Bouvier, W.D. Ristenpart, Aerosol emission and superemission during human speech increase with voice loudness, *Scientific Reports* 9 (1) (2019) 1–10.
- [3] G. Buonanno, L. Morawska, L. Stabile, Quantitative assessment of the risk of airborne transmission of SARS-CoV-2 infection: Prospective and retrospective applications, *Environment International* 145 (2020) 106112.
- [4] M.Z. Bazant, J.W. Bush, A guideline to limit indoor airborne transmission of COVID-19, *Proceedings of the National Academy of Sciences* 118 (17) (2021).
- [5] L. Morawska, J. Allen, W. Bahnfleth, P.M. Bluyssen, A. Boerstra, G. Buonanno, J. Cao, S.J. Dancer, A. Floto, F. Franchimon, et al., A paradigm shift to combat indoor respiratory infection, *Science* 372 (6543) (2021) 689–691.
- [6] M. Abkarian, S. Mendez, N. Xue, F. Yang, H.A. Stone, Speech can produce jet-like transport relevant to asymptomatic spreading of virus, *Proceedings of the National Academy of Sciences* 117 (41) (2020) 25237.
- [7] D. Johnson, L. Morawska, The mechanism of breath aerosol formation, *Journal of Aerosol Medicine and Pulmonary Drug Delivery* 22 (3) (2009) 229–237.
- [8] W.F. Wells, M.W. Wells, Air-borne infection, *Journal of the American Medical Association* 107 (21) (1936) 1698–1703.
- [9] J. Duguid, The size and the duration of air-carriage of respiratory droplets and droplet-nuclei, *Epidemiology & Infection* 44 (6) (1946) 471–479.
- [10] S. Asadi, A.S. Wexler, C.D. Cappa, S. Barreda, N.M. Bouvier, W.D. Ristenpart, Effect of voicing and articulation manner on aerosol particle emission during human speech, *PLoS One* 15 (1) (2020) e0227699.
- [11] L. Morawska, G. Johnson, Z. Ristovski, M. Hargreaves, K. Mengersen, S. Corbett, C. Chao, Y. Li, D. Katoshevski, Size distribution and sites of origin of droplets expelled from the human respiratory tract during expiratory activities, *Journal of Aerosol Science* 40 (2009) 256–269.
- [12] G. Johnson, L. Morawska, Z. Ristovski, M. Hargreaves, K. Mengersen, C. Chao, M.-P. Wan, Y. Li, X. Xie, D. Katoshevski, S. Corbett, Modality of human expired aerosol size distributions, *Journal of Aerosol Science* 42 (12) (2011) 839–851.
- [13] W. Wells, M.W. Wells, S. Mudd, Infection of air: Bacteriologic and epidemiologic factors, *American Journal of Public Health and the Nations Health* 29 (8) (1939) 863–880.
- [14] L. Bourouiba, E. Dehandschoewercker, J.W.M. Bush, Violent expiratory events: on coughing and sneezing, *Journal of Fluid Mechanics* 745 (2014) 537–563.
- [15] M. Rosti, S. Olivieri, M. Cavaioia, A. Seminara, A. Mazzino, Fluid dynamics of COVID-19 airborne infection suggests urgent data for a scientific design of social distancing, *Scientific Reports* 10 (1) (2020) 1–9.
- [16] R.K. Bhagat, M.D. Wykes, S.B. Dalziel, P. Linden, Effects of ventilation on the indoor spread of COVID-19, *Journal of Fluid Mechanics* 903 (2020).
- [17] P.F. Linden, The fluid mechanics of natural ventilation, *Annual Review of Fluid Mechanics* 31 (1) (1999) 201–238. [Online]. Available: doi: 10.1146/annurev.fluid.31.1.201..
- [18] A. Foster, M. Kinzel, Estimating COVID-19 exposure in a classroom setting: A comparison between mathematical and numerical models, *Physics of Fluids* 33 (2) (2021) 021904.
- [19] M. Nicas, W.W. Nazaroff, A. Hubbard, Toward understanding the risk of secondary airborne infection: emission of respirable pathogens, *J. Occ. Env. Hygiene* 2 (3) (2005) 143–154.
- [20] N.I. Stilianakis, Y. Drossinos, Dynamics of infectious disease transmission by inhalable respiratory droplets, *Journal of the Royal Society Interface* 7 (50) (2010) 1355–1366.
- [21] D. Lewis, COVID-19 rarely spreads through surfaces. So why are we still deep cleaning?, *Nature* 590 (7844) (2021) 26–28.
- [22] W.F. Wells, *Airborne Contagion and Air Hygiene. An Ecological Study of Droplet Infections*, Harvard University Press, Cambridge, MA, 1955.
- [23] L. Gammaitoni, M.C. Nucci, Using a mathematical model to evaluate the efficacy of TB control measures, *Emerging Infectious Diseases* 3 (3) (1997) 335.
- [24] C.J. Noakes, P.A. Sleight, Mathematical models for assessing the role of airflow on the risk of airborne infection in hospital wards, *Journal of the Royal Society Interface* 6 (suppl_6) (2009) S791–S800.
- [25] K. Lin, L.C. Marr, Humidity-dependent decay of viruses, but not bacteria, in aerosols and droplets follows disinfection kinetics, *Environmental Science & Technology* 54 (2) (2019) 1024–1032.
- [26] F.J. García de Abajo, R.J. Hernandez, I. Kaminer, A. Meyerhans, J. Rosell-Llompart, T. Sanchez-Elsner, Back to normal: An old physics route to reduce SARS-CoV-2 transmission in indoor spaces, *ACS Nano* 14 (7) (2020) 7704–7713.
- [27] M. Heßling, K. Hönes, P. Vatter, C. Lingenfelder, Ultraviolet irradiation doses for coronavirus inactivation—review and analysis of coronavirus photoinactivation studies, *GMS Hygiene and Infection Control* 15 (2020).
- [28] W. Kowalski, *Ultraviolet Germicidal Irradiation Handbook: UVGI for Air and Surface Disinfection*, Springer, 2010.
- [29] A. Konda, A. Prakash, G.A. Moss, M. Schmoltd, G.D. Grant, S. Guha, Response to letters to the editor on aerosol filtration efficiency of common fabrics used in

- respiratory cloth masks: Revised and expanded results, *ACS Nano* 14 (9) (2020) 10764–10770.
- [30] M. Abkarian, S. Mendez, N. Xue, F. Yang, H.A. Stone, Puff trains in speaking produce long-range turbulent jet-like transport potentially relevant to asymptomatic spreading of viruses, arXiv preprint arXiv:2006.10671, 2020.
- [31] F. Yang, A.A. Pahlavan, S. Mendez, M. Abkarian, H.A. Stone, Towards improved social distancing guidelines: Space and time dependence of virus transmission from speech-driven aerosol transport between two individuals, *Phys. Rev. Fluids* 5 (2020) 122501(R).
- [32] M. Gandhi, D.S. Yokoe, D.V. Havlir, Asymptomatic transmission, the achilles heel of current strategies to control COVID-19 2020 (2020) 2158–2160.
- [33] P. van den Berg, E.M. Schechter-Perkins, R.S. Jack, I. Epshtein, R. Nelson, E. Oster, W. Branch-Elliman, Effectiveness of three versus six feet of physical distancing for controlling spread of COVID-19 among primary and secondary students and staff: A retrospective, state-wide cohort study, *Clinical Infectious Diseases* 03 (2021).
- [34] J. Gettings, M. Czarnik, E. Morris, E. Haller, A.M. Thompson-Paul, C. Rasberry, T. M. Lanzieri, J. Smith-Grant, T.M. Aholou, E. Thomas, C. Drenzek, D. MacKellar, Mask use and ventilation improvements to reduce COVID-19 incidence in elementary schools—Georgia, November 16–December 11, 2020, *MMWR. Morbidity and Mortality Weekly Report* 70 (2021) 779–784.
- [35] E. Riley, G. Murphy, R. Riley, Airborne spread of measles in a suburban elementary school, *American Journal of Epidemiology* 107 (5) (1978) 421–432.
- [36] C. Beggs, C. Noakes, P. Sleight, L. Fletcher, K. Siddiqi, The transmission of tuberculosis in confined spaces: an analytical review of alternative epidemiological models, *Int. J. Tuberculosis and Lung Disease* 7 (11) (2003) 1015–1026.
- [37] P. Azimi, B. Stephens, HVAC filtration for controlling infectious airborne disease transmission in indoor environments: predicting risk reductions and operational costs, *Building and Environment* 70 (2013) 150–160.
- [38] S. Rudnick, D. Milton, Risk of indoor airborne infection transmission estimated from carbon dioxide concentration, *Indoor Air* 13 (3) (2003) 237–245.
- [39] H. Dai, B. Zhao, Association of the infection probability of COVID-19 with ventilation rates in confined spaces, *Building Simulation* 13 (6) (2020) 1321–1327.
- [40] J. Shen, M. Kong, B. Dong, M.J. Birnkrant, J. Zhang, A systematic approach to estimating the effectiveness of multi-scale IAQ strategies for reducing the risk of airborne infection of SARS-CoV-2, *Building and Environment* 200 (2021) 107926.
- [41] C. Sun, Z. Zhai, The efficacy of social distance and ventilation effectiveness in preventing COVID-19 transmission, *Sustainable Cities and Society* 62 (2020) 102390.
- [42] Z. Peng, J.L. Jimenez, Exhaled CO₂ as COVID-19 infection risk proxy for different indoor environments and activities, *Environmental Science and Technology Letters* (2021).
- [43] L.F. Pease, N. Wang, T.I. Salsbury, R.M. Underhill, J.E. Flaherty, A. Vlachokostas, G. Kulkarni, D.P. James, Investigation of potential aerosol transmission and infectivity of SARS-CoV-2 through central ventilation systems, *Building and Environment* (2021) 107633.
- [44] ASHRAE Epidemic Task Force, Building readiness, URL: <https://www.ashrae.org/technical-resources/building-readiness>, 2021, accessed: 2021-05-07.
- [45] P. MacNaughton, J. Pegues, U. Satish, S. Santanam, J. Spengler, J. Allen, Economic, environmental and health implications of enhanced ventilation in office buildings, *International Journal of Environmental Research and Public Health* 12 (11) (2015) 14709–14722.
- [46] AHAM, Method for measuring performance of portable household electric room air cleaners, ANSI/AHAM Standard AC-1-2020 (2020).
- [47] ASHRAE Epidemic Task Force, Core recommendations for reducing airborne infectious aerosol exposure, 2021, accessed: 2021-05-07. URL: <https://www.ashrae.org/file%20library/technical%20resources/covid-19/core-recommendations-for-reducing-airborne-infectious-aerosol-exposure.pdf>.
- [48] J. Stein, M.M. Hydeman, Development and testing of the characteristic curve fan model, *ASHRAE Transactions* 110 (1) (2004).
- [49] US Department of Energy, EnergyPlus 9.4.0, 2020. URL: <https://energyplus.net/>
- [50] M. Zaatari, A. Novoselac, J. Siegel, The relationship between filter pressure drop, indoor air quality, and energy consumption in rooftop HVAC units, *Building and Environment* 73 (2014) 151–161.
- [51] M. Orme, Estimates of the energy impact of ventilation and associated financial expenditures, *Energy and Buildings* 33 (3) (2001) 199–205.
- [52] H.R. Santos, V.M. Leal, Energy vs. ventilation rate in buildings: A comprehensive scenario-based assessment in the European context, *Energy and Buildings* 54 (2012) 111–121.
- [53] S.T. Taylor, C.H. Cheng, Why enthalpy economizers don't work, *ASHRAE Journal* 52 (11) (2010) 12–28.
- [54] ASHRAE, Ventilation for acceptable indoor air quality, ANSI/ASHRAE/IES Standard 62.1-2019 (2019).
- [55] ASHRAE, Energy standard for buildings except low-rise residential buildings, ANSI/ASHRAE/IES Standard 90.1-2019 (2019).
- [56] X. Zhou, J.E. Braun, A simplified dynamic model for chilled-water cooling and dehumidifying coilpart 1: Development, *HVAC&R Research* 13 (5) (2007) 785–804.
- [57] X. Zhou, J.E. Braun, A simplified dynamic model for chilled-water cooling and dehumidifying coils—part 2: Experimental validation, *HVAC&R Research* 13 (5) (2007) 805–817.
- [58] X. Li, Y. Li, J.E. Seem, P. Li, Dynamic modeling and self-optimizing operation of chilled water systems using extremum seeking control, *Energy and Buildings* 58 (2013) 172–182.
- [59] J. Seem, J. House, Development and evaluation of optimization-based air economizer strategies, *Applied Energy* 87 (3) (2010) 910–924.
- [60] S.N. Rudnick, Optimizing the design of room air filters for the removal of submicrometer particles, *Aerosol Science and Technology* 38 (9) (2004) 861–869.
- [61] M. Kawashima, C.E. Dorgan, J.W. Mitchell, Hourly thermal load prediction for the next 24 hours by ARIMA, EWMA, LR and an artificial neural network, *ASHRAE Trans. Part 1* 101 (1995) 186–200.
- [62] H.N. Gabow, J.L. Bentley, R.E. Tarjan, Scaling and related techniques for geometry problems, in: *Proceedings of the Sixteenth Annual ACM Symposium on Theory of Computing*, 1984, pp. 135–143.
- [63] U.E.P. Agency, Exposure Factors Handbook, National Center for Environmental Assessment, Washington, DC, 2011.
- [64] M.Z. Bazant, O. Kodio, A.E. Cohen, K. Khan, Z. Gu, J.W.M. Bush, Monitoring carbon dioxide to quantify the risk of indoor airborne transmission of COVID-19, *Flow*, vol. in press, 2021, medRxiv preprint: doi: 10.1101/2021.04.04.21254903.
- [65] C. Jiang, M.K. Masood, Y.C. Soh, H. Li, Indoor occupancy estimation from carbon dioxide concentration, *Energy and Buildings* 131 (2016) 132–141.
- [66] M.S. Zuraimi, A. Pantazaras, K.A. Chaturvedi, J.J. Yang, K.W. Tham, S.E. Lee, Predicting occupancy counts using physical and statistical CO₂-based modeling methodologies, *Building and Environment* 123 (2017) 517–528.
- [67] C. Jiang, Z. Chen, R. Su, M.K. Masood, Y.C. Soh, Bayesian filtering for building occupancy estimation from carbon dioxide concentration, *Energy and Buildings* 206 (2020) 109566.