

RESEARCH ARTICLE

Mixing and ventilation in a living laboratory due to fast and slow response modes

Costanza Rodda , John Craske  and Graham O. Hughes 

Department of Civil and Environmental Engineering, Imperial College London, London SW7 2AZ, UK

Corresponding author: Costanza Rodda; Email: c.rodga@imperial.ac.uk

Received: 28 May 2025; **Revised:** 6 February 2026; **Accepted:** 2 March 2026

Keywords: indoor air quality; ventilation; dynamical systems; mixing

Abstract

We present and analyse observational data from a highly instrumented classroom computer laboratory and develop a multi-zone model to describe its mechanical ventilation and mixing regime. The laboratory houses 70 workstations that are used heterogeneously in time and space, in a manner similar to a generic office environment. Our model predicts CO₂ concentration in the laboratory, accounting for air exchange between the occupied classroom and its ceiling plenum, and by parametrising irreversible mixing in each zone. Applying the model to our measurements helps identify critical components in the ventilation network, as highlighted by a strong separation of the time scales characterising the flow response. On the one hand, this time scale separation leads to a simplified model describing the CO₂ transport. On the other hand, it suggests that the forced exchange of volume between the room and the plenum is ‘overdriven’ in that reduced energy operation could be achieved without compromising air quality. More generally, our modelling approach offers a systematic method to enhance energy efficient ventilation of multi-zone systems.

Impact statement

Contaminant transport and mixing in enclosed spaces are shaped by complex physical processes such as advection and diffusion, which are sensitive to uneven and unsteady airflow patterns. These flow characteristics make it challenging to predict and optimise ventilation. Our study addresses this by representing the most relevant physical processes in a multi-zone model, which provides a clearer picture of how contaminants spread in indoor spaces. This research has practical implications for industries and communities involved in building design and management. It supports the creation of energy-efficient ventilation systems that reduce operating costs while minimising the accumulation of pollutants. By contributing to healthier and more sustainable indoor environments, our findings offer valuable tools for building professionals, researchers and practitioners working to improve air quality and energy performance in the built environment.

1. Introduction

1.1. Motivation

The COVID-19 pandemic and the unprecedented energy crisis many countries have faced recently have highlighted the need for new measures that maximise energy efficiency, and simultaneously ensure a safe and comfortable environment for building occupants. Ventilation is recognised as an essential ongoing measure to reduce the risk of the indoor spreading of airborne infectious diseases, such as the SARS-CoV-2 virus (Bazant & Bush, 2021; Bhagat *et al.*, 2020; Burrridge *et al.*, 2022). During the COVID-19



pandemic, many countries opted to increase the running time of heating, ventilation and air-conditioning (HVAC) systems, following recommendations from several studies (Dai & Zhao, 2020; Guo *et al.*, 2021; Sun & Zhai, 2020). However, HVAC systems are among the main contributors to a building's life cycle carbon emissions and impact heavily upon its running costs because they are associated with approximately 75 % of the total energy consumption (Laustsen, 2008).

Better design and control of HVAC systems require the development of relatively simple models that can be used efficiently at the scale of an entire building. This is challenging from a fluid mechanics perspective because the transport and mixing of air in buildings is spatially heterogeneous and occurs on a wide range of temporal and spatial scales.

1.2. Existing models

Several different approaches have been developed in the past decades to optimise HVAC control and improve the energy efficiency of systems (Afram & Janabi-Sharifi 2014). One such approach is demand-controlled ventilation (DCV), which uses feedback control methods to optimise the ventilation rate based on building occupancy variations. DCV has led to a substantial reduction in energy consumption (up to 45 % (Pang *et al.* 2020)) because a significant amount of energy would be otherwise wasted during unoccupied and low-occupation hours (Masoso & Grobler 2010). Occupancy measurement relies on either direct counting or an indirect proxy – typically CO₂ concentration. Advances in direct counting have been enabled by technologies such as wireless technologies (Wi-Fi), Bluetooth, camera-based systems and environmental sensors (Zhao *et al.* 2022). Nevertheless, significant errors remain with most of these techniques.

CO₂ concentration-based occupancy measurement relies on the strong correlation between CO₂ levels and the number of occupants. This correlation can be calculated for a single-zone space under the following assumptions: the air in the zone is well-mixed at all times and the occupants are the only source of CO₂. If these assumptions hold true, CO₂ concentration can be related to occupancy using a simplified mass-balance equation,

$$V \frac{dC}{dt} = -QC + NF, \quad (1.1)$$

where V is the volume and C is the CO₂ concentration above the outdoor concentration (under the assumption that the outdoor concentration does not change with time), Q is the ventilation rate, N is the number of people and F is the CO₂ generation rate per person (Lu *et al.* 2022).

In the context of building ventilation models, a single-zone model represents the entire room (or building) as one well-mixed control volume, characterised by a single concentration $C(t)$ and volume V . All supply, exhaust and internal sources are assumed to act on this homogeneous zone. In contrast, multi-zone models subdivide the space into several interconnected control volumes (zones), each with its own concentration and volume, and explicitly represent the mass fluxes between zones. Examples include separate zones for occupied rooms, corridors, plenums or ceiling voids. Multi-zone models, therefore, allow for spatial heterogeneity (e.g. differences between room and plenum), at the cost of introducing additional state variables and parameters.

In spaces containing multiple zones, transient airflows and complex interactions between zones can significantly affect CO₂ distributions and hence occupancy estimates.

1.3. Overview

In this work, we propose a technique that combines modelling and data analysis to address the challenge posed by the coupled problem of energy usage reduction and air quality control.

First, we derive a mass-balance equation for an arbitrary zone (sub-volume) within a building. Irreversible mixing in the zone is parametrised and our model makes explicit use of the physics that couples CO₂ and occupancy.

Then, we apply a dynamical systems approach to analyse the complex behaviours of airflows in multi-zone spaces. By using the spectral properties of the system, we identify response time scales that may be considered ‘fast’ and ‘slow’. This time scale separation enables us to project the system onto a lower-dimensional space, allowing for more effective analysis and optimisation. The use of time scale separation to simplify the equations describing a system and understand its properties has been used at length in different contexts, such as weather modelling (Peña & Kalnay 2004), geophysical flows (Vanneste 2013), combustion chemistry (Al-Khateeb *et al.* 2009) and homogenisation for stochastic modelling more generally, where it has also been referred to as ‘adiabatic elimination’ (Haken 1983; Pavliotis & Stuart 2008). However, its application in the analysis of building ventilation has not been exploited.

Finally, we examine the model performance against the real data, providing insights into the impact of various parameters on CO₂ concentrations and potential energy savings.

Parker and Bowman (2011) have applied a state-space formulation to study the analytical solutions and characterise the dynamical behaviour of multi-zone buildings in the context of contaminant transport. We adopt a similar approach, investigating the importance of key parameters in relation to the characteristics of our system. In this study, we demonstrate how our method can be applied to ventilation systems. Geometrically, we represent the ventilation response as trajectories in phase space, which proves to be an insightful way of characterising the system’s behaviour.

To illustrate the application of this approach in understanding and modelling air exchanges among connected spaces and parametrising the mixing of air within sub-spaces, we analyse observational data from a ‘living’ laboratory – a highly instrumented classroom at Imperial College London. The laboratory is representative of a wide range of mixed ventilation spaces with characteristics encompassing many building types, including offices (with not less than 6 m² person^{−1}), which are ubiquitous worldwide. The data measured in the laboratory provide a way to understand how the system evolves in response to different configurations and forcings. The model of our living laboratory is represented by a space composed of a plenum containing fan coil units above the suspended ceiling of the occupied room – a typical configuration for mechanically ventilated rooms. In this space, we consider carbon dioxide (CO₂) as a proxy for ventilation and respiratory contaminants. CO₂ can be modelled relatively easily since its transport is limited to advection (or ventilation), while other quantities like heat are more complicated because transport occurs by convection, conduction and radiation, and the sources are comparatively difficult to quantify.

The laboratory, measurements and dataset are introduced in § 2, and the analytical model is described in § 3. Section 4 focuses on comparing the model with the observed data, while in § 5, we discuss model reduction and potential generalisation for n -dimensional spaces. Our discussions and conclusions are given in § 6 and § 7.

2. A ‘living’ laboratory and dataset

This study is based on data measured in a classroom equipped with a cluster of computer workstations – a ‘living’ laboratory – located at the South Kensington Campus of Imperial College London. Figure 1 shows a three-dimensional (3-D) sketch of the laboratory highlighting its main features. The laboratory consists of two spaces: a main room hosting the occupants and workstations (20.1 m in length, 8.0 m in width and 3.2 m in height for a total volume $V_0 \approx 515$ m³), and a plenum situated above the room hosting the air conditioning system (20.1 m in length, 8.0 m in width and 0.5 m in height for a total volume $V_1 \approx 80$ m³). The laboratory is adjacent to a lecture room on the west side and an open-plan student space on the north and east sides. The south-facing facade has 12 double-glazed windows, each with a surface area of 2.6 m² that are locked in a closed position and cannot be opened by the occupants. The laboratory is accessible through two doors located on the north wall. These doors are self-closing (and alarmed) and, therefore, the period for which the doors are open (while people pass into and out of the room) is assumed not to affect the room interior. The laboratory is subject to an extensive range of spatially and temporally heterogeneous thermal forcing; it hosts 70 workstations, a similar number




Figure 1. Schematic diagram of the computer laboratory. The coloured arrows and lines indicate the airflow directions occurring within the ducts in the plenum. The blue lines/arrows indicate the supply of air from the air handling units that are injected into the room by the four FCUs. The red lines/arrows indicate air that has been extracted from the room through the five ceiling grilles. The CO₂ sensors are marked S_{ph} , where ‘p’ corresponds to the (horizontal) position and the index ‘h’ corresponds to height, for which the integers 1, 2, 3 and 4 correspond to 3, 2, 1 and 0.5 m, respectively, above floor level.

of potential occupants (floor space per person therefore ranging from approximately 2 to 160 m²), and is mechanically heated and cooled (see § 2.1). The laboratory has been instrumented throughout with numerous CO₂, humidity and temperature sensors (see § 2.3).

2.1. Heating ventilation and air conditioning (HVAC)

The laboratory is mechanically ventilated via a typical variable air volume (VAV) system controlled by a building management system (BMS). Temperature-controlled fresh air is pumped into a plenum above the room (see the blue pipe in Figure 1). Four fan coil units (FCUs) (Quartz Sapphire model SPR9) draw air from the plenum and heat/cool the air using a hot/cold water system according to the temperature setpoint for the room. The conditioned air is ducted from the FCUs to supply the room via linear slot diffusers placed at the north and south sides of the room (blue pipes and arrows in Figure 1). The rate of supply to the room is usually different from the rate of extraction, which is constant and occurs through five square grilles positioned along the longitudinal axis of the ceiling (red pipes in Figure 1).

To account for the possible mismatch between the rate at which air is supplied by the FCUs and extracted according to the VAV system, the plenum and the room exchange air through two additional grilles on the ceiling (not shown). In addition, sections of the linear slot diffusers are not connected to the FCUs, and allow air exchange between the room and the plenum. These air exchanges can occur in either direction (from the room towards the plenum and *vice versa*) depending on the air handling unit (AHU) regime.

2.2. Occupancy

Occupancy estimates are inferred from Imperial’s deployment of the HubStar (formerly LoneRooftop) Building Insights Dashboard, which records the number of Wi-Fi connections (hourly average and maximum values are available). Wi-Fi estimated occupancy has several advantages, such as protecting

the privacy of occupants and achieving a high accuracy (96 % according to the study by Simma *et al.* (2019)). These advantages and the wide availability of Wi-Fi signals in buildings have made it a convenient means of occupancy estimation in many studies in the past ten years (Zhao *et al.* 2022). Nevertheless, estimates have uncertainties associated with the varying number of Wi-Fi devices each person carries. In our case, the time resolution and the anonymisation of the recording of Wi-Fi connections lead to the estimated uncertainty of the order of 30 % (see Appendix B.1 in the supplementary material available at <https://doi.org/10.1017/flo.2026.10045> for more details).

2.3. Measurements

The laboratory is monitored continuously with 24 temperature sensors (Trend thermistors model T/TFR 4) and 8 combined CO₂/temperature/humidity sensors (TREND space sensors model RS-WMB-THC) connected to the BMS (via 4 Trend IQ4 controllers). The temperature sensors are installed on 4 vertical risers (i.e. 6 per riser) at different positions in the room. The combined sensors are situated on vertical walls at different horizontal positions in the room (see Figure 1 for more details). The CO₂ and temperature measurements were further supplemented by sensors in the supply and extract ducting of the BMS system.

The dataset presented in this paper spans from June 2022 to November 2022, with measurements sampled and recorded every 5 minutes. An overview of the temperature data collected from the sensors on the risers is shown in Figure 2. Each panel shows the distribution of the temperature measured by the sensors at the given height, after subtracting the mean room temperature at each measurement time. The largest departures from the mean occur at the lowest two measurement levels. The sensor at 0.5 m is located beneath the desks of the room and the one at 1.0 m is close to the computers; both are therefore influenced by local heat sources and partial shielding, which explains their larger deviations from the rest of the vertical profiles. For the other heights, the distributions lie within 1 °C, with the exception of a small number of outliers. This narrow distribution for most of the data suggests that the room temperature is rather uniform along the vertical.

For the rest of the present study, we focus on the data provided by the 8 combined CO₂/temperature/humidity sensors. We assume that the occupants are the only source of CO₂ in the room.

We supplement the data acquired in the room with a range of standard measurements recorded by the BMS at 5-min intervals. These measurements include the supply and extraction volume flow rates, Q_{in} and Q_{out} , respectively, and the supply and return temperatures from the FCU. The flow rates Q_{in} and Q_{out} are measured by sensors placed in the air ducts. The data are presented in § 4.

3. Analytical model

3.1. CO₂ budget

Let C_0 and C_1 be the bulk excess CO₂ concentration (relative to outdoors) in the room and plenum, respectively, and let C'_0 and C'_1 be the quasi-steady-state excess concentration associated with the ceiling zone and FCU, respectively (see Figure 3 for a schematic representation of the spaces in a vertical cross-section of the room). We use q and q' to denote the rates at which fresh air is supplied to the plenum and at which air is drawn into the plenum from the ceiling zone, respectively, and define $Q := q + q'$ (see Figure 3). The variable Q , therefore, corresponds to the volume flux driven by the FCUs.

First, we establish an equation for the CO₂ concentration in the FCUs, the combined volume of which is labelled V'_1 . Unless stated otherwise, all CO₂ concentrations given refer to excess values.

The dimensionless parameters γ_0 and γ_1 encode how the FCU flow is partitioned between different pathways within the plenum–room system. Since the supply air duct is not connected directly to the FCUs, let $\gamma_1 \in [0, 1]$ represent the fraction of the supply rate q that first mixes with the air in the plenum before entering the FCU:




Figure 2. Monthly distributions of vertical heterogeneities (expressed as $T - \bar{T}$, with $\bar{T}$ the room averaged temperature at a given time) from the temperature sensors for June–November 2022.

$$V_1' \frac{dC_1'}{dt} = \underbrace{(\gamma_1 q + q') C_1}_{\text{from plenum}} - \underbrace{Q C_1'}_{\text{to room}}, \tag{3.1}$$

for which we have assumed that $\gamma_1 q + q' \geq 0$, i.e. the FCU always extracts air from the plenum. For the control volume just below the room’s ceiling, we adopt a similar approach to the FCU, using $\gamma_0 \in [0, 1]$ to denote the fraction of FCU outlet air that first mixes with the rest of the room (before entering the ceiling zone):




Figure 3. (a) Schematic diagram (not to scale) of a vertical section through the ceiling plenum and computer laboratory below, with excess concentrations C_1 and C_0 , respectively; the corresponding control volumes are highlighted with a red dashed line. C'_1 and C'_0 indicate the excess CO_2 concentration in the FCU and ceiling zone (highlighted with a darker shade of grey). As illustrated by the white arrows, conditioned air is supplied to the plenum with a volume flux q . A fraction $\gamma_1 \in [0, 1]$ of that supply air mixes with the air within the plenum; the rest is drawn directly into the fan coil unit (FCU), whose fan drives a total volume flux $Q = q + q'$. After being fed into the computer laboratory, a fraction $\gamma_0 \in [0, 1]$ of air from the FCU either mixes with the air in the room, while the rest comprises a 'short circuit' and remains in the ceiling zone. Due to the fact that $Q \neq q$, the FCU drives a secondary volume flux $q' = Q - q$ between the computer laboratory and plenum. (b) Graph corresponding to the governing equations described in (3.2). Each node corresponds to a control volume (labelled with CO_2 concentration) and each branch corresponds to a flux of CO_2 between control volumes (labelled with the corresponding volume flux).

$$V'_0 \frac{dC'_0}{dt} = \underbrace{\gamma_0 Q C_0}_{\text{from room}} + \underbrace{(1-\gamma_0) Q C'_1}_{\text{from FCUs}} - \underbrace{q' \phi}_{\text{to plenum}} - \underbrace{q C'_0}_{\text{to extract}}, \quad (3.2)$$

where

$$\phi = \begin{cases} C'_0 & q' \geq 0, \\ C_1 & q' < 0, \end{cases} \quad (3.3)$$

allows the flux of CO_2 to be represented as either from the ceiling zone to the plenum ($q' \geq 0$) or from the plenum to the ceiling zone ($q' < 0$).

Let us now consider the CO_2 concentration in the room and plenum. We assume that the rate of CO_2 generation per person is $F = 0.012 \text{ g s}^{-1} \text{ person}^{-1}$, consistent with the range $F = 0.009\text{--}0.012 \text{ g s}^{-1} \text{ person}^{-1}$ for occupants aged 21–30 years undertaking typical office work (Persily & de Jonge 2017). Indeed, our data (to be discussed in Figure 8) will indicate that $F = 0.012 \pm 0.001 \text{ g s}^{-1} \text{ person}^{-1}$ provides the best fit for the model, as evaluated by the highest R^2 value. Furthermore, we assume that the infiltration and exfiltration of air are negligible (the validity of such an assumption is discussed in Appendix B.2). The system of linear ordinary differential equations (ODEs) describing the time-variation of concentration in the room and the plenum respectively is therefore

$$V_0 \frac{dC_0}{dt} = \underbrace{\gamma_0 Q C'_1}_{\text{from FCU}} - \underbrace{\gamma_0 Q C_0}_{\text{to ceiling zone}} + NF, \quad (3.4)$$

$$V_1 \frac{dC_1}{dt} = \underbrace{q'\phi}_{\text{from ceiling zone}} - \underbrace{(\gamma_1 q + q')C_1}_{\text{to FCU}}, \quad (3.5)$$

where N is the number of occupants. The system above is similar to the two-zone transient model proposed by Lawrence and Braun (2006), with the difference that we have relaxed the condition of complete mixing in the plenum (V_1) and occupied zone ($V'_0 + V_0$) by adding the γ parameters. The governing equations can be graphically represented in the network diagram in Figure 3b, where the nodes represent control volumes labelled with concentrations. The volume fluxes between control volumes are represented by the network's branches.

Let us now re-write the governing equations (3.1), (3.2), (3.4) and (3.5) in non-dimensional form by introducing the following variables:

$$\tau := \frac{q}{V_1} t, \quad \varepsilon := \frac{V_1}{V_0}, \quad \zeta := \frac{q'}{q}, \quad f_0 := \frac{NFV_1}{qV_0}. \quad (3.6)$$

Physically, τ is a dimensionless time based on the filling time scale associated with the plenum (in terms of its volume V_1 and the AHU supply volume flux q), ε represents the ratio of the plenum and room volumes, ζ is the ratio of the rates of room-plenum re-circulation and ambient fresh air introduction, and f_0 is the dimensionless forcing.

We further assume that the FCUs and ceiling buffer zone volume are much smaller than the plenum volume, i.e. $V'_1/V_1 \ll 1$ and $V'_0/V_1 \ll 1$. In physical terms, storage of CO_2 in the FCU and ceiling buffer zone, represented by the time derivative in (3.1)–(3.2), is therefore assumed to be insignificant, such that (3.1) and (3.2) give explicit algebraic expressions relating C'_0 and C'_1 to γ_0 , γ_1 , q and q' .

Under these assumptions, the resulting non-dimensionalised equations are

$$\frac{dC_0}{d\tau} = -\varepsilon\gamma_0(\zeta + 1)C_0 + \varepsilon\gamma_0(\zeta + \gamma_1)C_1 + f_0, \quad (3.7)$$

$$\frac{dC_1}{d\tau} = \zeta\phi - (\zeta + \gamma_1)C_1. \quad (3.8)$$

We can express the system (3.7) and (3.8) in the form

$$\frac{d\mathbf{C}}{d\tau} = \mathbf{A}\mathbf{C} + \mathbf{f}, \quad (3.9)$$

where $\mathbf{A}$ is a coefficient matrix, $\mathbf{C} = (C_0, C_1)^\top$ is the state vector and $\mathbf{f} = (f_0, 0)^\top$ is the forcing vector.

Assuming that $\mathbf{A}$ is diagonalisable (see discussion in § 5.2 for the general case in which this assumption is not valid), we can derive the solution for (3.9),

$$\mathbf{C} = \mathbf{R} \exp(\mathbf{A}\tau) \mathbf{R}^{-1} (\mathbf{C}(0) + \mathbf{A}^{-1} \mathbf{f}) - \mathbf{A}^{-1} \mathbf{f}, \quad (3.10)$$

where $\mathbf{C}(0)$ are the initial conditions, $\mathbf{A}$ is the eigenvalue matrix and $\mathbf{R}$ is the corresponding (right) eigenvector matrix. The derivation of (3.10) and explicit expressions for the eigenvalues and eigenvectors are provided in Appendix A.

3.2. Phase space diagrams

Each possible state of the system is represented by a point in phase space, and a sequence of consecutive states forms a trajectory. Therefore, phase space diagrams are useful for representing systems geometrically, and investigating their equilibria, stability and evolution.

If the forcing $\mathbf{f}$ is regarded as constant, then (3.9) can be expressed as the autonomous system:

$$\frac{d\mathbf{y}}{d\tau} = \mathbf{A}\mathbf{y}, \quad (3.11)$$

where

$$\mathbf{y} := \mathbf{C} + \mathbf{A}^{-1} \mathbf{f}. \quad (3.12)$$




Figure 4. Phase space diagram corresponding to the variable $\mathbf{y}$ (3.12), which accounts for the stationary equilibrium $\mathbf{A}^{-1}\mathbf{f}$ for fixed $\varepsilon = 0.17$, $\zeta = 1.8$, and several combinations of γ_0 and γ_1 values. The grey arrows show the flow evolution trajectories for different initial conditions and the black dots show one particular solution of the dynamical system at uniform time intervals. The red and blue arrows show the two eigenvectors.

Figure 4 illustrates the phase space associated with (3.11). We choose $\varepsilon = 0.17$ and $\zeta = 1.8$ in (3.7) and (3.8) because they represent the measurements in the living laboratory and are characteristic of a ventilated office space. The values of γ_0 and γ_1 are varied to understand how the (parametrised) mixing affects the rate of convergence towards one manifold of the system. The grey arrows in Figure 4 mark flow evolution trajectories for different initial conditions. The red and blue arrows mark the directions of the two eigenvectors and correspond to the fast (large negative eigenvalue) and slow (small negative eigenvalue) dynamics of the system, respectively.

The black dots show one particular solution of the dynamical system at uniform time intervals. They show that the solutions converge rapidly towards the slow eigenmode, which dominates the adjustment to the equilibrium state $\mathbf{y} = \mathbf{0}$.

In terms of the parameters, we can easily see that γ_1 has far less influence on the dynamics than γ_0 . This behaviour comes from the condition $\varepsilon \ll 1$, which corresponds to the volume of the plenum being

Table 1. Input parameters for the analytical model (3.4), (3.5). The values of q are measured by the BMS system. The forcing NF is estimated by the measured Wi-Fi connections and $F = 0.012 \text{ g s}^{-1}$ per person. Q is estimated by air speed measurements done in the room (see Appendix B.2).

	Occupied space	Plenum
Volume	$V_0 = 512 \text{ m}^3$	$V_1 = 80 \text{ m}^3$
Mixing parameter	$0 < \gamma_0 < 1$	$0 < \gamma_1 < 1$
Volume flux (ON)	$Q = 0.8 \text{ m}^3\text{s}^{-1}$	$q = 0.25 \text{ m}^3\text{s}^{-1}$
Volume flux (OFF)	$Q = 0.01 \text{ m}^3\text{s}^{-1}$	$q = 0.02 \text{ m}^3\text{s}^{-1}$

much smaller than the volume of the room below ($V_1 \ll V_0$). When $\gamma_0 \rightarrow 1$ and $\gamma_1 \rightarrow 0$, the CO_2 concentrations in the room and the plenum become comparable, as can be seen by the blue vector approaching a slope of one in the bottom right panel in Figure 4. When $\gamma_0 = 1$ and $\gamma_1 = 0$, the slow manifold maps onto $y_0 = y_1$, while the fast manifold corresponds to the difference between the CO_2 concentration in the room and plenum. In all other cases shown in Figure 4, the concentration is higher in the room than in the plenum. More generally, for other combinations of γ_0 and γ_1 , the physical interpretation of the slow and fast manifolds is less straightforward, as they involve other linear combinations of the concentrations in the rooms, which depend on the mixing properties of the system.

4. Data analysis and comparison with the model

This section is dedicated to analysing CO_2 measurements in the laboratory using the phase space representation, followed by comparison with the predictions obtained from the model presented in § 3.

Table 1 summarises the input values for the analytical model in (3.4) and (3.5). The value for the volume flux q – the fresh air fed to the plenum – is directly measured by sensors placed in the supply and return ducts with a sampling interval of 5 min; Table 1 gives the average calculated over a month of measurements. Here, Q – the volume flux from the FCU to the occupied room – is not routinely measured, but is obtained by taking *ad hoc* measurements with an air speed transducer (see Appendix B.3). The forcing term NF is calculated using the occupancy estimation from Wi-Fi connection measurements (for N) and taking $F = 0.012 \text{ g s}^{-1} \text{ person}^{-1}$.

Figure 5 shows the measured excess of CO_2 in the room (C_0) and in the plenum (C_1) for the entire month of July (weekends excluded) in a phase space diagram. The dataset is divided into two sets corresponding to the ON regime (Figures 5a and 5b) and the OFF regime (Figures 5c and 5d). Panels (a) and (c) depict the data in the y -coordinate system (as per (3.12)). The model (in terms of the y coordinate) is also shown on these panels as a reference (with the grey arrows marking the flow evolution trajectories, and the red and blue arrows marking the eigenvectors representing the fast and slow dynamics of the system, respectively). Panels (b) and (d) show the measured data in terms of excess CO_2 concentrations. The volume fluxes and volumes are input into the model to form the matrix A according to the measured values (see Table 1), while the parameters γ_0 and γ_1 cannot be measured, and therefore need to be fitted with the data. For the case ON, the data are always well represented by $\gamma_0 = 1$ and $\gamma_1 = 0$, which correspond to strong mixing within the room and to supply air that enters the FCU having effectively bypassed the plenum, respectively. It is striking how the concentrations converge rapidly onto the subspace describing the slow dynamics of the system, marked by the blue arrow. This convergence points to a significant time separation between the slow and fast dynamics of the system. Another characteristic we can deduce from this plot is that the exchange driven by the FCU could be considered as ‘overdriving’ the response of the room in the sense that lowering the fan speed setting would still achieve a quick collapse onto the slow dynamics line.

The data are more scattered in the OFF regime (Figures 5c and 5d), for which q and therefore Q are significantly smaller than in the ON regime (but are not zero, see Table 1). Panel (d) shows how the data




Figure 5. Phase space diagram showing the excess of CO_2 data in the room (C_0) and in the plenum (C_1) for July (excluding weekends). The spaces in the y -coordinate system in (a) data when ON, (c) data when OFF. The blue arrow marks the eigenvector related to the slow manifold, and the red arrow in panel (c) the fast manifold, both calculated for $\gamma_0 = 1$ and $\gamma_1 = 0$. The grey arrow shows the flow evolution trajectories as in Figure 4. The same dataset is plotted in the C_0, C_1 space (b) for data when ON and (d) data when OFF. The blue arrow shows the eigenvector calculated for $\gamma_0 = 1$ and $\gamma_1 = 0$, and the green arrow marks the eigenvector for $\gamma_0 = 0.2$ and $\gamma_1 = 0.5$. The markers' colour represents the time of the day, as indicated in the legend.

might be represented by the model with a range of γ parameters, with a limiting case that corresponds to the same conditions as in the ON regime (marked by the blue arrow). Most of the other data lie 'below' this eigenvector (i.e. $C_1 < C_0$) and the model can be fitted using a range of combinations of γ_0 and γ_1 , suggesting that mixing in both the room and the plenum plays an important role in the OFF regime. In Figure 5d, the green arrow marks the eigenvector calculated by setting $\gamma_0 = 0.2$ and $\gamma_1 = 0.5$. Since there is little sensitivity to γ_1 for most values of γ_0 , similar eigenvectors can be obtained for different values of γ_1 .

To better understand what causes the spread in the OFF regime, we plotted the phase space diagrams in Figure 6 to compare the CO_2 concentrations over six months, from June to November. To quantify how closely the data follow the slow manifold in the C_0 – C_1 phase space, we decompose each data point into a component along the slow-manifold line and a component normal to it. Let d denote the distance from the origin measured along the slow-manifold line and s the signed normal distance to this line. We then define $\theta := \arcsin(s/d)$, which, for small deviations, is approximately equal to the angular




Figure 6. Phase space diagram showing the CO_2 data in the room (C_0) and the plenum (C_1) for months between June 2022 and November 2022 at the times when the FCU is off. The blue arrow indicates the slow manifold, same as in Figure 5. Variances of θ are: June 0.025, July 0.084, August 0.057, September 0.179, October 0.0025, November 0.0023.

deviation between the state vector (C_0, C_1) and the slow eigendirection. Points lying exactly on the slow manifold have $\theta = 0$, whereas a non-zero θ measures the contribution of modes that do not lie in the slow subspace. The variance $\text{Var}(\theta)$ given in the caption of Figure 6 therefore provides a dimensionless, scale-invariant measure of the spread of the data around the slow manifold: larger values correspond to a wider angular dispersion of points and hence weaker alignment with the slow mode. The monthly values of $\text{Var}(\theta)$ show that the data are most tightly clustered around the slow manifold in October and November, where the corresponding angular standard deviation is only a few degrees, and most widely scattered in June–September, where the angular spread is substantially larger.

The monthly variation of the angular spread around the slow manifold is consistent with the seasonal change in the plenum–room temperature difference: larger spreads occur in months when the thermal contrast, and hence the potential for stratification and departures from well-mixed behaviour, is greater. This suggests a link between external temperature, the resulting plenum–room temperature difference, and the internal mixing within the room and plenum.

Table 2. Best-fit values of the mixing parameters γ_0 and γ_1 obtained by least-squares fitting of the model to the ON and OFF data for each month

Month	$\gamma_{0,\text{on}}$	$\gamma_{0,\text{off}}$	$\gamma_{1,\text{on}}$	$\gamma_{1,\text{off}}$
June	0.997	0.376	0.002	0.725
July	1.000	0.529	0.000	0.705
August	0.996	0.658	0.129	0.709
September	1.000	0.646	0.072	0.559
October	0.909	0.419	0.214	0.831
November	1.000	0.392	0.000	0.562




Figure 7. Boxplot of the monthly averaged temperature difference between plenum and room, and the CO_2 residuals quantifying the spread as shown in Figure 6b. The horizontal line within each box marks the median, and the boxes extend from the first to the third quartile. The white dot marks the mean value. The circle-shaped markers indicate the outliers.

To examine this link, we consider the temperature difference recorded by the BMS between the plenum and the room (Figure 7). In July, August and September, the air supplied to the plenum is warmer than the room air, so that the air entering the room from the plenum is lighter and tends to spread by displacement, promoting vertical stratification. By contrast, when the incoming air is colder than the room, the denser air descends and enhances stirring and mixing. These buoyancy-driven effects, which are captured in our model through the mixing parameter, therefore provide a plausible explanation for the observed increase in angular spread around the slow manifold in months with larger plenum–room temperature differences.

Finally, we can use the theoretical model to study the time evolution of the CO_2 within the room and the plenum. Figure 8 illustrates such a comparison for a chosen day (the second Tuesday of each month from June to November). The dataset spans very different conditions, with the occupation in any given hour ranging from a few people to over 50. The outside air temperature ranges from a few degrees Celsius in the autumn months to above 30°C in the months of July and August.

The data (blue and yellow dots) correspond to the CO_2 concentrations in the laboratory and plenum. The solid blue and yellow lines represent the solution of (3.7) and (3.8) fed by the experimental values given in Table 1. The mixing parameters γ_0 and γ_1 are obtained by a least-squares fit for each month, and separately for the ON and OFF regimes. The resulting best-fit values are listed in Table 2. For the ON regime, the fits consistently give $\gamma_0 \approx 1$ and $\gamma_1 \approx 0$, confirming that, when the FCUs are operating, almost all of the air passes through the room and there is little short-circuiting through the ceiling zone. The room and ceiling layer, therefore, behave as a single, well-mixed reservoir, which explains the rapid collapse of the trajectories onto the slow manifold. In contrast, the OFF regime exhibits smaller γ_0 and




Figure 8. CO₂ concentration (relative to outdoors) time evolution and model prediction. The data (blue and yellow dots) correspond to the measured concentrations in the laboratory on a representative day of the dataset (the first Tuesday of each month). The ambient CO₂ (410 ppm) has been subtracted from the data. Note that the range of CO₂ concentrations (left y-axis) differs significantly from month to month. The solid yellow and blue lines represent the solution of (3.7) and (3.8); see text for more details. The blue-shadowed histograms indicate occupancy in terms of the number of people in the room. The grey-shadowed region marks the ON regime. The dashed curves under each plot represent the residuals and the root mean squared error (RMSE) and R² values for the ON phase are annotated in the caption.

larger γ_1 systematically. Values of $\gamma_{0,\text{off}} \approx 0.4\text{--}0.7$ indicate that a substantial fraction of the flow short-circuits within the ceiling zone and returns to the plenum without fully mixing with the occupied space, while $\gamma_{1,\text{off}} \approx 0.6\text{--}0.8$ shows that the fresh air mixes in the plenum before being ventilated by the FCUs into the room.

Considering the uncertainty related to the occupancy estimation (see § 2.2) and that occupancy is sampled only hourly whilst the data are sampled every 5 minutes, the model is in good agreement with the measurements. For the regime ON, the R² values are above 0.7, except when occupancy values are below 5 (i.e. the CO₂ forcing is relatively weak), as in the September and October panels in Figure 8 for which the R² values drop below 0.4. The RMSE plots show that the largest errors correspond to the

OFF regime and to periods with large relative variations in occupancy that again are coarsely sampled in time. When the FCUs are on (in the shaded grey region of Figure 8), the CO₂ concentration in the room and the plenum are nearly always equal, consistent with the fitted values $\gamma_{0,\text{on}} \approx 1$ and $\gamma_{1,\text{on}} \approx 0$. In contrast, when the FCUs are switched off, the CO₂ response to the forcing is more unpredictable, and the concentration difference between the plenum and the room varies. This behaviour is in line with the smaller $\gamma_{1,\text{on}}$ and larger $\gamma_{1,\text{off}}$ obtained from the OFF fits, which reflect increased short-circuiting and stratification in the ceiling–plenum region and help explain the higher modelling errors in this regime.

5. Model reduction and generalisation to higher dimensions

5.1. Model reduction

We revisit the model presented in § 3.1 in the light of the behaviour highlighted by the comparison with the measured data. Numerous situations have emerged in which there is a sufficiently large separation of time scales that, from the perspective of the relatively slow evolution of the overall system, the transient effects associated with the fast dynamics can be neglected. Exploiting a separation of time scales in this way is used widely to derive a simplified low-dimensional description of a physical system and, in broad terms, corresponds to the ‘adiabatic elimination’ of fast variables in stochastic systems (Haken 1983).

The equation governing the second (fast) eigenmode can be written (Appendix A; (35))

$$\frac{dz_2}{d\tau} = \lambda_2 z_2 + \frac{1}{\lambda_2 R_{21}} \frac{df_0}{d\tau}, \quad (5.1)$$

where $|1/\lambda_2|$ is the response time scale of the eigenmode, z_2 is the amplitude of the eigenmode and $\mathbf{R}$ is the eigenvector matrix. This reduction is valid only in parameter regimes where $\rho := \max(\lambda_1, \lambda_2) / \min(\lambda_1, \lambda_2) \ll 1$; in the low- ζ limit with $\gamma_0 \approx 1$ or small γ_1 , where $\lambda_1 \approx \lambda_2$, the fast–slow splitting becomes ill-conditioned and the full two-dimensional system must be used (see § 6).

Since $\lambda_2 < 0$, we can deduce that an upper bound for finite changes Δz_2 is obtained by ignoring the first term on the right-hand side of (5.1): $\Delta z_2 \leq \Delta f_0 / (\lambda_2 R_{21})$. The upper bound becomes tighter in the limit that the corresponding $\lambda_2 \Delta t \rightarrow 0$. For larger values of Δt , the upper bound becomes increasingly conservative. In other words, it changes in f_0 over time scales shorter than $1/\lambda_2$ that are likely to lead to deviations from $z_2 \approx 0$, which would remain a valid approximation, and therefore enable a reduction of the dimension of the system if $\Delta f_0 \ll \lambda_2 R_{12}$. In practice, this condition might not be fulfilled when there are rapid changes in occupancy, such as at the beginning and end of a class.

By assuming that the equation for z_2 plays no role, the (one-dimensional) reduced model can be written (Appendix A; (35)) as

$$\frac{dz_1}{d\tau} = \lambda_1 z_1 + \frac{1}{\lambda_1} R_{11}^{-1} \frac{df_0}{d\tau}, \quad (5.2)$$

and the solution (3.10) reduces to

$$\mathbf{C} = \mathbf{R}_1 \exp(\lambda_1 \tau) \mathbf{R}_1^{-1} (\mathbf{C}_0 + \mathbf{A}^{-1} \mathbf{f}) - \mathbf{A}^{-1} \mathbf{f}, \quad (5.3)$$

where $\mathbf{R}_1$ is the column vector of the eigenvector matrix (corresponding to the two components of one eigenvector) and $\mathbf{R}_1^{-1} := \mathbf{R}_1^T / (\mathbf{R}_1^T \mathbf{R}_1)$ is the left inverse. Note that (5.3) depends only on the eigenvalue and eigenvector associated with the slow dynamics.

The full and reduced models are compared in Figure 9a (blue and green lines, respectively), where the time evolution of C_0 is plotted for a chosen day in August (same data as in Figure 6). The reduced model prediction (given by (5.3)) is consistent with the full model prediction (given by the integration of (3.4)). We can notice that the rapid changes in occupancy are approximated in the reduced order system by an instantaneous adjustment of the fast mode and are responsible for the differences between the outputs from the full and reduced order models (panel a).

In our study, the pronounced separation between fast and slow time scales arises from the specific geometry and flow configuration of the living laboratory. In particular, the plenum volume is much




Figure 9. (a) Comparison of the reduced (green line) and full model (blue line) with the observational data (black dots) for the month of August. The parameters for the two models are the same as the one used in Figure 6. (b) Occupancy data.

smaller than the occupied-room volume ($\varepsilon \ll 1$) and when the FCUs are operating (ON regime), the data are well described by $\gamma_0 \approx 1$, indicating that most of the supply air is well mixed in the occupied room. Under these conditions and assuming that the effective ceiling layer is thin, the excess CO₂ concentration in that layer is close to the bulk room concentration, so that additional storage in the ceiling zone does not introduce a distinct dynamical time scale. Consequently, the dynamics collapse rapidly onto a one-dimensional slow manifold associated with the room–plenum exchange. Although these characteristics might arise in other spaces, we want to emphasise that this behaviour is not universal. In configurations where the ceiling zone has a larger effective volume or where persistent stratification leads to significant departures of the ceiling concentration from the bulk room concentration, the ceiling layer can support an additional slow mode. In such cases, the projection onto a single slow eigenmode may be inappropriate and a higher-dimensional reduced model (or even the full multi-zone description) may be required.

5.2. Generalisation to higher dimensional models

The model presented in § 3.1 can be generalised to analyse the dynamics of an arbitrary n -zone configuration. In that case, the dynamics of the system would still be expressed as (3.9), but $\mathbf{C}$ would be a vector of size n and C_i would represent the CO₂ concentration in the i th zone. Prior to any reduction of the system, an element A_{ij} for $i \neq j$ of the $n \times n$ matrix $\mathbf{A}$ corresponds to a normalised ventilation into zone i from zone j , whereas the diagonal elements $A_{ii} \leq 0$ would represent the total normalised volume flux directed out of zone i , as described by Parker and Bowman (2011). Although $\mathbf{A}$ cannot, in general, be diagonalised because it might not have n independent eigenvectors, it can be transformed into Jordan normal form. A trivial example of a configuration for which $\mathbf{A}$ cannot be diagonalised comes

from n rooms connected in series, such that fresh air is supplied to the first room in the series, while air from the last room discharges to a surrounding environment. In that case, the matrix A , when written in Jordan normal form, contains a single ‘Jordan block’, with non-zero entries along the diagonal (corresponding to the ventilation rate) and ones along the superdiagonal. The ventilation rate corresponds to the only eigenvalue of the system with geometric multiplicity equal to one (i.e. the n eigenvectors of the system point in the same direction). More generally, systems consisting of multiple interacting flow loops will contain several Jordan blocks. In such cases, it would, in principle, be possible to apply the approach described in this paper by examining the eigenvalue and, therefore, time scale associated with each block and its corresponding eigenvector. Model reduction would then correspond to projecting the system onto the product of eigenspaces associated with slow dynamics.

6. Discussion

An analytical multi-zone ventilation flow model and a dynamical systems analysis have been applied in a four-zone case study to demonstrate the potential of the approach in detail. We have further reasoned that the same analysis can be applied to other ventilated multi-zone spaces to assess whether a similar separation of time scales, and hence a low-dimensional description, can be used, as discussed in § 5.2. It is important to note that although our case study focuses on a mechanically ventilated space, the same methods can be applied to any HVAC control system and even to naturally ventilated buildings, provided the airflows are known.

The analysis of the time scales involved in the system helps identify a slow manifold to capture the bulk response behaviour, dimensionality reductions that can be made in modelling the system and the transient regimes that play a significant role during adjustment processes. For instance, the data analysis in § 4 highlights the rapid convergence in the ON regime of CO₂ concentration onto the slow dynamics of the system. A dimensionality reduction of the system response is possible, i.e. a collapse onto a single dimension in this instance, because the resulting time scale separation between the fast and slow dynamics is more than a factor of 10: the characteristic fast time scale is approximately 30 s, while the slow time scale exceeds 5 minutes (characteristic time scales are calculated as $1/\lambda$ and then dimensionalised as per (3.6)).

More generally, we can analyse the eigenvalues expressed in (43), which correspond to an inverse time scale associated with each eigenvector, to explore the sensitivity of the system in parameter space. For $\varepsilon \ll 1$ in our case study, the two eigenvalues are of a different order of magnitude throughout most of the parameter space, supporting characterisation by a one-dimensional dynamical system. Regions where this dimensionality reduction becomes invalid can also be found along with the dependence on the values of the other three variables: ζ , γ_0 and γ_1 . Figure 10 illustrates such dependencies for $\varepsilon = 0.15$, with the colours corresponding to the ratio $\rho := \max(\lambda_1, \lambda_2) / \min(\lambda_1, \lambda_2)$. This ratio governs the separation between fast and slow time scales; when $\rho \ll 1$, the eigenvalues are well separated and the corresponding eigenvectors form a well-conditioned basis, so that trajectories collapse rapidly onto a one-dimensional slow manifold. It can be seen from the figure that the reduction to a one-dimensional description is expected to hold throughout most of the parameter space ($\rho \approx 0$), consistent with the finding (by inspection) in previous sections. If fresh air is introduced in the system, even a surprisingly small recirculation $q' \gtrsim 0.3q$ is sufficient to promote collapse onto the one-dimensional system. However, for $\zeta \lesssim 0.3$, there are always combinations of γ_0 and γ_1 for which the system remains bi-dimensional, which corresponds to $\rho \rightarrow 1$. In this limit, the recirculation between room and plenum is weak (small ζ), while the FCU flow is either almost entirely routed through the room ($\gamma_0 \approx 1$) or predominantly extracted from the ceiling zone (small γ_1). Under such conditions, the room and plenum adjust on similar time scales, the two eigenvalues coalesce and the fast–slow decomposition becomes ill-conditioned; the dynamics remain intrinsically two-dimensional.

In addition to suggesting a leading order representation of complex ventilation systems, our approach helps to identify when and how a system is running inefficiently, and therefore how the running cost and energy requirements might be reduced. The total energy consumption of an HVAC system




Figure 10. Ratio of eigenvalues (indicated by colour) as a function of the dependent variables (a) ζ and γ_0 with $\gamma_1 = 0.18$, and (b) ζ and γ_1 with $\gamma_0 = 1.0$. $\varepsilon = 0.15$ for both panels.

depends on numerous variables, including the supply air flow rate, and is typically complex to estimate (Atthajariyakul & Leephakpreeda 2004). Nevertheless, in the case study presented here, our analysis indicates that for optimal CO_2 concentration distribution between the room and the plenum, the air recirculated by the FCUs needs to generate a recirculation that is only 30 % of the fresh air supplied into the plenum. From this perspective, the excess air recirculated to the room is overdriving the system without providing any real benefit in pollutant removal. The model suggests that the FCU air recirculation for the studied configuration could be reduced by more than 50 %, resulting in a significant reduction in costs without impacting the air quality experienced by occupants.

To estimate the potential energy saving, let us consider the annual energy usage of the FCUs in the room. According to the technical specifications, each FCU installed in the room has a fan operation electrical power consumption of 0.3 kW for the high-speed setting (Versatile 2024). Considering that the FCUs are used for an average of 8 hours per day and that there are 4 FCUs installed in the room, the estimated annual energy consumption is approximately 3500 kWh year⁻¹. Note that this estimate focuses on fan power only because the thermal condition of the room still needs to be maintained. If we halve the FCU flow rate as suggested by the model, this will result in a decrease in annual energy usage by 1750 kWh year⁻¹. In terms of energy consumption per unit area, this reduction corresponds to 10.9 kWhm⁻² year⁻¹ (for a room with a surface area of 160 m²).

7. Conclusions

We have studied the evolution of pollutant concentration in connected spaces by combining analytical modelling with analysis of observational data from a living laboratory. The use of the data to validate the modelling has offered a robust approach for predicting pollutant concentrations in a ventilated space.

The formulation of an analytical model depends upon first identifying a number of zones in the ventilated space along with key parameters describing the zone geometries, the flow between zones and the mixing within each zone. The methodology then uses an eigenmode decomposition to analyse the model properties in phase space. The aim is finally to identify an appropriate reduced-dimension description of the ventilation flow system.

In the living laboratory examined in this study, we proposed several zones (sub-volumes) to represent the potential heterogeneity in the pollutant (CO_2) distribution. The analysis highlighted the importance

of parameters describing the relative volumes of the selected zones; in particular, significantly different time scales associated with the ventilation of each zone proved to be a key step towards a significantly simplified model. Moreover, although the values of some other parameters in the model were not known, the approach allowed the sensitivity of the system to those parameters to be quantified. Throughout the parameter space considered in this study, we were thus able to show that detailed knowledge of these parameters was unimportant.

Ventilation flows and pollutant distribution in buildings are undoubtedly governed by more complex physics than we have attempted to represent, but the present study has highlighted some useful general principles. First, and somewhat ironically, a good operational balance between ventilation and energy consumption seems likely when an inadequate model would result from a reduced-order representation, i.e. the eigenvalues of a modal decomposition have comparable values. Second, it is apparent that knowledge of pollutant sources – in this case, spatio-temporal variability in occupancy – represents a major source of uncertainty for models, in general. These effects are typically overlooked in modelling approaches that assume spaces to be well mixed. However, the dataset available from the numerous sensors situated throughout our living laboratory has the potential to help answer outstanding questions about spatial pollutant distribution, but a key future step is to better account for spatio-temporal characteristics (including any associated thermal effects) of the corresponding sources. A better understanding of potential accumulation zones for pollutants would help improve the design and placement of air extractors and benefit air quality and safety.

Supplementary material. The supplementary material is available at <https://doi.org/10.1017/flo.2026.10045>.

Acknowledgements. We are grateful to Mark Reader for providing the BMS data and insightful comments about the HVAC system. We thank Neal Streamer and Trend Controls for supporting the project, and providing the sensors and controllers. We thank Chris Banks from Imperial College London for facilitating access to the occupancy data from Imperial's deployment of the HubStar (formerly LoneRooftop) Building Insights Dashboard, which infers occupancy from Wi-Fi connections.

Author contributions. C.R. carried out the investigation, data analysis and wrote the original manuscript. G.O.H. and J.C. acquired funds, instrumented the laboratory, and developed the original research idea. C.R., G.O.H. and J.C. edited the manuscript and contributed to developing the methodology.

Funding statement. This work was supported by the Engineering and Physical Sciences Research Council [grant number EP/V033883/1] as part of the [D*] stratify project.

Declaration of interests. The authors declare no conflict of interest.

Ethical standards. The research meets all ethical guidelines, including adherence to the legal requirements of the study country.

References

- Afram, A., & Janabi-Sharifi, F. (2014). Theory and applications of hvac control systems – A review of model predictive control (MPC). *Building and Environment*, 72, 343–355.
- Al-Khateeb, A. N., Powers, J. M., Paolucci, S., Sommesse, A. J., Diller, J. A., Hauenstein, J. D., & Mengers, J. D. (2009). One-dimensional slow invariant manifolds for spatially homogenous reactive systems. *The Journal of Chemical Physics*, 131(2), 024118.
- Attahajariyakul, S., & Leephakpreeda, T. (2004). Real-time determination of optimal indoor-air condition for thermal comfort, air quality and efficient energy usage. *Energy and Buildings*, 36(7), 720–733.
- Bazant, M. Z., & Bush, J. W. M. (2021). A guideline to limit indoor airborne transmission of COVID-19. *Proceedings of the National Academy of Sciences of the United States of America*, 118(17), e2018995118.
- Bhagat, R. K., Wykes, M. S. D., Dalziel, S. B., & Linden, P. F. (2020). Effects of ventilation on the indoor spread of COVID-19. *Journal of Fluid Mechanics*, 903, 2020.
- Burridge, H. C., Fan, S., Jones, R. L., Noakes, C. J., & Linden, P. F. (2022). Predictive and retrospective modelling of airborne infection risk using monitored carbon dioxide. *Indoor and Built Environment*, 31(5), 1363–1380.

- Dai, H., & Zhao, B. (2020). Association of the infection probability of COVID-19 with ventilation rates in confined spaces. *Building simulation*, 13, 1321–1327.
- Guo, M., Xu, P., Xiao, T., He, R., Dai, M., & Miller, S. L. (2021). Review and comparison of hvac operation guidelines in different countries during the COVID-19 pandemic. *Building and Environment*, 187, 107368.
- Haken, H. (1983). Synergetics: An introduction: Nonequilibrium phase transitions and self-organization in physics, chemistry, and biology. In *Recent results in cancer research number* (Vol. 3). Springer. ISBN 9780387088662. <https://books.google.co.uk/books?id=09PuAAAAMAAJ>.
- Jha, N. K., Frank, D., Darracq, L., & Linden, P. F. (2021). Contaminant transport by human passage through an air curtain separating two sections of a corridor: Part II—two zones at different temperatures. *Energy and Buildings*, 236, 110728.
- Laustsen, J. (2008). Energy efficiency requirements in building codes, energy efficiency policies for new buildings. *International Energy Agency*, 85(2008).
- Lawrence, T. M., & Braun, J. E. (2006). Evaluation of simplified models for predicting CO₂ concentrations in small commercial buildings. *Building and Environment*, 41(2), 184–194.
- Lu, X., Pang, Z., Fu, Y., & O'Neill, Z. (2022). The nexus of the indoor CO₂ concentration and ventilation demands underlying co₂-based demand-controlled ventilation in commercial buildings: A critical review. *Building and Environment*, 218, 109116.
- Masoso, O. T., & Grobler, L. J. (2010). The dark side of occupants' behaviour on building energy use. *Energy and Buildings*, 42(2), 173–177.
- Pang, Z., Chen, Y., Zhang, J., O'Neill, Z., Cheng, H., & Dong, B. (2020). Nationwide hvac energy-saving potential quantification for office buildings with occupant-centric controls in various climates. *Applied Energy*, 279, 115727.
- Parker, S. T., & Bowman, V. (2011). State-space methods for calculating concentration dynamics in multizone buildings. *Building and Environment*, 46(8), 1567–1577.
- Pavliotis, G. A., & Stuart, A. (2008). *Multiscale Methods: Averaging and Homogenization*. Texts in Applied Mathematics. Springer. ISBN 978038773828. <https://books.google.co.uk/books?id=qIaD19q7Tl8C>.
- Persily, A., & de Jonge, L. (2017). Carbon dioxide generation rates for building occupants. *Indoor Air*, 27(5), 868–879.
- Peña, M., & Kalnay, E. (2004). Separating fast and slow modes in coupled chaotic systems. *Nonlinear Processes in Geophysics*, 11(3), 319–327.
- Simma, K. C. J., Mammoli, A., & Bogus, S. M. (2019). Real-time occupancy estimation using Wi-Fi network to optimize HVAC operation. *Procedia Computer Science*, 155, 495–502.
- Sun, C., & Zhai, Z. (2020). The efficacy of social distance and ventilation effectiveness in preventing COVID-19 transmission. *Sustainable Cities and Society*, 62, 102390.
- Vanneste, J. (2013). Balance and spontaneous wave generation in geophysical flows. *Annual Review of Fluid Mechanics*, 45, 147–172.
- Versatile. (2024). Fan coils and HVAC technical brochure. Retrieved 09 September, from 2024, https://www.versatile.ie/wp-content/uploads/2021/10/Versatile_Quartz_-FanCoilUnits.pdf.
- Zhao, L., Li, Y., Liang, R., & Wang, P. (2022). A state of art review on methodologies of occupancy estimating in buildings from 2011 to 2021. *Electronics*, 11(19), 3173.