




RESEARCH ARTICLE

Uniformly distributed floor sources of buoyancy can give rise to significant spatial inhomogeneities within rooms

Carolanne V.M. Vouriot^{1,2}, Thomas D. Higton¹, P.F. Linden² , Graham O. Hughes¹, Maarten van Reeuwijk¹  and Henry C. Burridge^{1,*} 

¹Department of Civil and Environmental Engineering, Skempton Building, South Kensington Campus, Imperial College London, London SW7 2BX, UK

²Department of Applied Mathematics and Theoretical Physics, Centre for Mathematical Sciences, University of Cambridge, Wilberforce Rd, Cambridge CB3 0WA, UK

*Corresponding author. E-mail: h.burridge@imperial.ac.uk

Received: 7 December 2022; **Revised:** 26 April 2023; **Accepted:** 7 May 2023

Keywords: Natural ventilation of buildings; Buoyancy driven flows; Horizontal convection

Abstract

Displacement ventilation, where cool external air enters a room through low-level vents and warmer air leaves through high-level vents, is characterised by vertical gradients in pressure arising from the warmer indoor temperatures. Models usually assume that horizontal variations of temperature difference are small in comparison and are, therefore, unimportant. Small-scale laboratory experiments and computational fluid dynamics were used to examine these flows, driven by a uniformly heated floor. These experiments and simulations show that the horizontal variations of temperature difference can be neglected for predictions of the bulk ventilation rate; however, they also evidence that these horizontal variations can be significant and play a critical role in establishing the pattern of flow within the room – this renders the horizontal position of the low- and high-level vents (relative to one another) important. We consider two cases: single-ended (where inlet and outlet are at the same end of the room) and opposite-ended. In both cases the ventilation flow rate is the same. However, in the opposite-ended case a dead zone is established in the upper part of the room which results in significant horizontal variations. We consider the formation of this dead zone by examining the streamline patterns and the age of air within the room. We discuss the implications for occupant exposure to pollutants and airborne disease.

Impact Statement

Exposure to indoor air pollution and airborne diseases are major factors in human health and well being. Guidance on appropriate ventilation rates is typically based on bulk ventilation rates, either in terms of the amount supplied per individual or as air exchange rates for a space. Such bulk measures assume homogeneous conditions within a space while, in practice, there are often significant spatial variations in properties. This paper shows that in displacement ventilation, where it is commonly assumed that horizontal variations are negligible, in fact simply altering the horizontal position of an outlet vent can lead to large variations in indoor air quality. Consequently, exposures based on average values can be misleading. This finding also has important implications for the location of sensors to measure the conditions of indoor environments, something which is becoming increasingly commonplace.

1. Introduction

Where feasible, displacement ventilation strategies, i.e. the low-level introduction and high-level extraction of air within building spaces, offer the potential for reduced energy usage (Linden, 1999; Wachenfeldt, Mysen, & Schild, 2007) and improved indoor air quality (Bhagat & Linden, 2020; Sandberg, 1981). The physics of displacement ventilation strategies have long been studied (e.g. Gladstone & Woods, 2001; Linden, Lane-Serff, & Smeed, 1990) with the standard assumptions that spaces are horizontally uniform; this assumption implies that the horizontal location of ventilation openings is irrelevant. We seek to highlight that this is not always the case, even in the limit where the bulk vertical ventilation flow within a space is naturally driven by buoyancy forces (arising from ubiquitous temperature differences) without any enhancement by external winds. In such cases, predictions of the ventilation flow, and the indoor conditions, are typically made by assuming the space to be either well-mixed or stratified in layers (two layers in the simplest case). Which arises is broadly taken to depend on whether the low-level heat (buoyancy) sources can be considered to be (horizontally) distributed or localised (Linden, 1999), with localised heat sources taken to generate coherent vertical flows, i.e. plumes. In this latter case, the bulk vertical ventilation flow is carried across interface(s) by these plumes; knowledge of the vertical evolution of the volume flux within the plumes can be combined with the location of the interface(s) to provide estimates of the ventilation flow rate. Our focus is the former case, i.e. displacement ventilation driven by distributed heat sources, and we challenge the traditional view (see figure 1) of the room conditions that arise.

It is well known that horizontal fluid motion can play a significant role in determining the transport in confined regions of heated fluid (Hughes & Griffiths, 2008), yet such knowledge is rarely considered when designing indoor spaces. The displacement ventilation flow and layered stratification generated in the presence of the heat input from a single person within a room (of albeit limited size) have been investigated via high resolution direct numerical simulations in the case that the inflow is mechanically forced (Yang, Ng, Chong, Verzicco, & Lohse, 2022). With this freedom to force the horizontal inflow, Yang et al. (2022) showed that, by altering the horizontal position of the upper outlet from one side of the room to the other with respect to the inlet, the height of the interface was altered, and that this height is critical to predicting where exposures within a space might be high. Studies of naturally driven displacement ventilation have also shown the ability of external wind to force the flow to transition, from one which maintains a layered stratification, to one in which horizontal motions mix the space when the wind is sufficiently strong. For example, Hunt and Linden (1999) define a Froude number condition to describe the relative forcing of wind to buoyancy; above a threshold Froude number, one can expect layered stratifications to become mixed. We assert that in cases when layered stratifications are disrupted, while well-mixed models might still be capable of predicting bulk flow rates with suitable accuracy, predicting exposures based on horizontal uniformity may be risky. This has already been evidenced by Yang et al. (2022) for ventilation flows that can be forced arbitrarily hard. However, here we show our assertion to have wider validity by investigating displacement flows ‘naturally’ driven only by distributed floor sources of buoyancy. Although such sources are widely expected to result in conditions reasonably approximated by the well-mixed assumption, e.g. figure 1, we show that this is not the case.

In the current study, we deploy both experiments and simulations of ventilation flows through low- and high-level vents driven by uniformly distributed floor sources of buoyancy (equivalent to the heat loads in operational rooms during typical heating season), to examine the distribution of buoyancy (temperature), and associated flow fields within idealised empty rooms. We examine the impact of changing the room aspect ratio and changing the relative horizontal positions of the low-level inlet and high-level outlet vents, so that in some configurations the flow exits the room in the opposite corner to the inlet and, in other configurations, the flow exits the room at the same end of the room as the inlet. We choose to describe these as (i) opposite-ended configurations, and (ii) single-ended configurations, respectively. We note that within our terminology we attempt to draw a distinction from the terms ‘cross-flow ventilation’ (sometimes simply ‘cross-ventilation’) and ‘single-sided ventilation’, which are




Figure 1. Schematic of the traditional well-mixed model for a room ventilated, via low- and high-level vents, by a flow driven by distributed floor source of heat (buoyancy). Our research shows this view to only be appropriate for predicting bulk ventilation rates.

taken to refer (implicitly in the case of ‘cross-flow’, and explicitly for ‘single-sided’) to ventilation flows through openings in walls (see CIBSE, 2005), and are sometimes taken to refer to wind dominated ventilation. However, we note that our terminology for the configuration is deliberately not explicit in distinguishing between vertically (i.e. in walls) or horizontally (i.e. in ceilings or floors) aligned openings as, with some considered thought, our findings can be applied to both.

The remainder of the paper is structured as follows. In § 2 we describe the fundamental parameters which determine the scale of the bulk ventilating flow. Details of the laboratory (§ 2.1) and numerical (§ 2.2) experiments are presented before moving on to our results, § 3. We present results for the bulk flow rates attained, and then focus on the buoyancy field within the room from our experiments in § 3.1. These are then compared and contrasted with those of our simulations, first examining the effect of relative horizontal position of the inlet and outlet vents in § 3.2, before investigating the fluid mechanics which underpins such findings (§ 3.3), and then considering the implications for occupant experience within § 3.4. Finally, we draw conclusions (§ 4).

2. Methodology

This study was inspired by observations made during an experimental campaign, which were then supplemented by a set of numerical simulations, of the natural ventilation flow in a room with low- and high-level vents separated by a vertical distance, H . Within both our experiments and simulations the flow is driven by a buoyancy source, uniformly distributed over the floor, within a room connected to the external environment by a low-level inlet and a high-level outlet – such set-ups are typically expected to generate relatively well-mixed conditions within the room (Linden, 1999).

Bulk ventilation flow rates Q were created within the room due to the stack pressure, arising due to the buoyancy, over the height H . In the case that vents are vertically aligned (e.g. a door-like ventilation opening) we measure H from the half-height of the opening. These flows were generated by an integral source buoyancy flux F (or equivalently an integral heat flux W ; with $W = \rho c_p F / (\beta g)$, where g is the gravitational acceleration, and ρ , c_p and $\beta = 1/T_{ref} \approx 1/300$ K are the fluid density, specific heat capacity and thermal expansion coefficient, respectively). In all cases the buoyancy flux was input evenly over the whole floor. The flow is geometrically constrained by the low- and high-level vents of areas A_l and A_h , respectively, providing resistance to the bulk flow via the effective area

$$A^* = \frac{\sqrt{2}C_l A_l C_h A_h}{\sqrt{(C_l A_l)^2 + (C_h A_h)^2}} \approx \frac{\sqrt{2}C_d A_l A_h}{\sqrt{A_l^2 + A_h^2}}, \quad (2.1)$$

where C_l and C_h are the loss coefficients at the low- and high-level vents respectively, and the right-hand side becomes an equality when $C_d \equiv C_l = C_h$.

For the case of a well-mixed room the uniform buoyancy is F/Q , and in the presence of low- and high-level vents, Linden (1999) provides the expected volume flux as

$$Q_M = (A^2 FH)^{1/3}. \quad (2.2)$$

To aid comparison with the well-mixed state, we present our results via the dimensionless volume flux

$$\hat{Q} = \frac{Q}{Q_M}, \quad (2.3)$$

and dimensionless time, scaled by the ventilation air-change rate, namely

$$\hat{t} = \frac{t}{V/Q}, \quad (2.4)$$

where V is the volume of the room.

Since our interest is to investigate in-room quantities, we present measurements of the buoyancy $b = b(x, y, z, t) = [\rho_a - \rho(x, y, z, t)]g/\rho_a = \beta[T(x, y, z, t) - T_a]g$, where ρ_a and T_a are the (uniform) density and temperature of the fluid within the ambient environment, respectively. We note that temperature differences within the room are typically at least two orders of magnitude smaller than absolute temperatures, and hence the Boussinesq approximation is valid. We use an overbar to denote quantities averaged over a time interval t_i , which is suitable to obtain good statistical estimates of the steady mean, e.g. \bar{b} . In order to present variations, within the room, about the ‘well-mixed’ state, and because temperature is the active scalar of interest, we present the scaled buoyancy

$$\overline{\Delta\mathcal{T}}(x, y, z) = \frac{\bar{b}(x, y, z)}{F/Q} = \frac{\rho c_p [\bar{T}(x, y, z) - T_a]}{W/Q}, \quad (2.5)$$

so that $\Delta\mathcal{T} = 1$ for a fully mixed room. Finally, we utilise angled brackets, $\langle \cdot \rangle$, for spatial averages. For example, we describe data averaged over the (spanwise) width of the box as ‘width-averaged’, denoting these data $\langle \overline{\Delta\mathcal{T}} \rangle_y = \overline{\Delta\mathcal{T}}(x, z)$. Similarly, averages over horizontal cross-sections are denoted $\langle \overline{\Delta\mathcal{T}} \rangle_{x,y} = \overline{\Delta\mathcal{T}}(z)$, and averages over the room are denoted $\langle \overline{\Delta\mathcal{T}} \rangle_V$, with $\langle \overline{\Delta\mathcal{T}} \rangle_V = 1$ for a fully mixed room. We note that the buoyancy data gathered via experiments (see § 2.1) are, effectively, integrated as light rays traverse the width of the box (i.e. in the direction of the coordinate we denote ‘y’) – a process which we describe as producing ‘width-averaged’ data. In the numerical simulations, the width-averaged results are obtained by taking an average of 54 vertical planes across the room.

2.1. Laboratory experiments

Experiments were conducted using the salt-bath technique, where room-scale ventilation can be studied using small-scale models submerged in water (Linden, 1999). The use of water as the fluid medium ensures that dynamic similarity, in both Reynolds and Rayleigh numbers, is broadly achieved at a room : model scale of around 10 : 1, and using salt (rather than heat) as the buoyancy scalar, as was realised in our experiments. Modelling the buoyancy, arising in buildings due to temperature differences in air, via saline differences in water also simplifies flow visualisation and enabled the observations included in this study. We choose to describe the experiments in the orientation of an actual room, such that we describe the fluid as being affected by a floor-level buoyancy source, rising up and exiting via the high-level vent.




Figure 2. The width-averaged scaled buoyancy field $\langle \overline{\Delta T} \rangle_y$ for the experiment (§ 2.1). The observations described based on the video images (§ 3.1), that the flow field is strongly heterogeneous is evident in this image of the time-averaged buoyancy field.

A cuboid, as a small-scale model of a room, was connected to the ambient environment by two openings, a low-level inlet vent (aligned in the vertical plane much like a small doorway) and a high-level outlet vent; the positions of the vents are illustrated in figure 2 – see table 1 for details. The high-level outlet was positioned in the ceiling approximately two thirds of the way along the room from the inlet, and hence the bulk flow must be both upwards and across the room. A saline solution was supplied to the scale model at floor level across a porous plastic panel. The source solution was dyed, and the model backlit, to enable the width-averaged two-dimensional buoyancy field $\langle b \rangle_y = b(x, z, t)$ to be measured using the dye-attenuation technique (Allgayer & Hunt, 2012; Cenedese & Dalziel, 1998). Note that, experimental data directly adjacent to the walls are compromised due to parallax errors, camera viewing restrictions and reflections from the walls and hence these regions are excluded from the results presented – these regions are highlighted in figure 2. Appendix A provides further details of the experimental set-up and procedures.

Our experiments, investigating the flow established by a buoyancy source (uniformly) distributed over the entire area, included the low-level vent being vertically aligned, akin to a doorway, but we only examined cases for which a unidirectional inflow was achieved at the low-level doorway. We note that when presenting data from, and describing the flow field of, our experiments we have chosen to invert the set-up. This decision was taken to ensure a direction of the flow that is consistent with our numerical simulations and our primary application of interest, i.e. building ventilation; akin to this we choose to describe the confining experimental box as a ‘room’.

2.2. Numerical simulations

Numerical simulations were conducted with OpenFOAMv2106 using the transient buoyantPimpleFoam solver. Simulations were run using the Reynolds-averaged Navier-Stokes (RANS) $k-\omega$ shear stress transport (SST) turbulence model (kOmegaSST). The RANS models only solve for the mean flow and model the remaining turbulent scales, as such, they are not able to resolve all the flow features which

Table 1. Relevant parameters, and resulting flow rate, ranges spanned by our experiments and simulations. Note that in all of the above we take the discharge coefficient $C_d = 0.65$, Appendix C presents the sensitivity to this choice.

Parameter (range)	Experiments	Simulation
Room volume, V (m ³)	0.041	149
Room length, L (m)	0.45	10
Room width, W (m)	0.3	5.5
Vertical distance, H (m)	0.27	2.7
Effective area, A^* (m ²)	0.8×10^{-3} – 1.8×10^{-3}	0.08–0.33
Buoyancy flux, F (m ⁴ /s ³)	5.8×10^{-6}	0.17
Vent Reynolds number, Re (–)	3.36×10^3 – 4.17×10^3	3.95×10^4 – 4.63×10^4
Bulk Rayleigh number, Ra (–)	8.41×10^{11} – 4.25×10^{11}	8.09×10^{10} – 3.40×10^{10}
Dimensionless volume flux, \hat{Q} (–)	0.95–1.03	0.99–1.07

might otherwise be resolved, for example, with large-eddy simulations. In the present work, however, the focus is on the overall flow pattern in a room set by large-scale buoyancy variations, which RANS models are expected to capture. The relatively low cost of RANS was therefore exploited to consider different ventilation configuration and a range of vent areas. The k - ω SST model was specifically chosen for its ability to model flows near walls, shear flows and potential relaminarisation, all of which are needed to simulate ventilating flows.

A room domain, of size 10 m \times 5.5 m \times 2.7 m, was meshed and connected to larger exterior boxes through low- and high-level vents which both lay in horizontal planes and with effective areas in the range $0.012 \leq A^*/H^2 \leq 0.049$, see table 1. For all simulations discussed herein, the area of the high-level vent, A_h , was smaller than that of the low-level vent, A_l , falling in the range $0.25 \leq A_h/A_l \leq 0.5$ – this avoided excessive momentum flux at the low-level vent (associated with unwanted cold drafts in rooms). For the example illustrated in § 3, $A_h/A_l = 0.5$. The larger exterior allowed the flow to be entirely driven by the buoyancy source by ensuring sufficient spatial separation between any prescribed boundary conditions and the flow within the room. The horizontal position of the high-level vent was varied from being at the far end of the room away from the low-level vent (much like the experimental configuration), to being at the same end of the room as the low-level vent; as mentioned above we describe these configurations as ‘opposite-ended’ and ‘single-ended’, respectively. A wintertime scenario was investigated where the external ambient temperature was set to $T_a = 278$ K (5 °C) and the flow was driven by imposing a 6200 W heat flux on the floor. An additional scalar transport equation was solved for a passive tracer representing the age of air t_{AoA} . Defined to be equal to 0 at the inlet, the age of air represents the time required to reach a certain point in the space and is often used to assess the efficiency of a ventilation system (Sandberg, 1981). Herein, we present the age of air scaled by the ventilation air change rate such that $\hat{t}_{AoA} = t_{AoA}V/Q$. Appendix B provides further details of the numerical set-up.

3. Results

Before examining the buoyancy field, and ultimately the flow patterns that influence it, we examine the bulk ventilation flow rates attained. Table 1 shows details of the experimental and simulated configurations considered. The breadth of the underlying physical parameter ranges explored is significant, and results in the vent Reynolds and bulk Rayleigh numbers each varying by at least an order of magnitude. For completeness, we define the vent Reynolds number as $Re = Q/(A^{*1/2}\nu)$, where ν is the kinematic viscosity of the fluid, and the bulk Rayleigh number $Ra = FH^3/(Q\kappa\nu)$, where κ is the thermal or molecular diffusivity, as appropriate for the simulations or experiments, respectively. Despite the breadth of our investigation, the range of dimensionless volume fluxes attained, \hat{Q} , varies only slightly from unity

(by always less than 7%) – this demonstrates that the realised ventilation flow rate is in good agreement with the classical theory based on the well-mixed assumption, and hence adheres to the scaling, namely (2.2). Moreover, the relatively narrow range of \hat{Q} observed highlights that if one is only interested in predicting the bulk ventilation flow then one can conclude that the well-mixed model provides appropriate estimates – at least to within uncertainty, much of which centres around the parameterisation of the losses at the vent via the coefficient C_d ; see Appendix C for the sensitivity of \hat{Q} on C_d .

In subsequent sections, we report detailed findings for a subset of the configurations reported in table 1; for convenience of communication only, we use one representative example experiment and one simulation for consistent illustration. Crucially, examination of the scalar field shows that a relatively well-mixed state, expected for indoor flows driven by distributed floor sources of buoyancy, was not observed, and this hitherto unreported observation forms the motivation of this study. We proceed to show that, if one is required to predict the scalar field, for example, to predict occupant thermal comfort or exposure to pollutants, then the well-mixed model may be inadequate.

3.1. Experimental observations and measurements of the width-averaged buoyancy field

Figure 2 shows the width-averaged (scaled) buoyancy field $\langle \overline{\Delta T} \rangle_y$, from the statistically steady-state, flow analysed. It is evident that buoyancy within the room is not uniformly mixed, e.g. there is a cooler region close to the floor, which rises at the far wall, and a warmer region in the upper part of the room. This results in spatial variation of buoyancy within the room spanning a range of around 20% of the mean buoyancy. Moreover, the horizontal variation in buoyancy within the lower two thirds ($0 \leq z/H \leq 2/3$) of the room is greatest. In many modern rooms, ceilings are relatively low, and this region is therefore likely to be that experienced by occupants; thus, predicting these variations is a worthy research challenge.

We note that more significant inhomogeneities were observed in near wall regions which are compromised due to camera viewing restrictions, parallax and reflections from the walls. The observations apparent in these images are more clearly evident in the video included as supplementary material and movies, available at <https://doi.org/10.1017/flo.2023.11>. In this video, as with our other experiments, a statistically steady state had been attained before recording began (the experiment had already been running for over nine air-change periods, i.e. a duration of $\hat{t} \approx 9.5$). However, unlike our other experiments – i.e. unlike those which provided our quantitative experimental results – the buoyant source fluid ejected was left undyed, and instead a sudden (low momentum) release of dyed neutrally buoyant fluid was made in the ambient environment just outside the inlet vent (by bursting a water filled balloon). This introduced a transient passive scalar (i.e. dye) concentration field within the statistically steady flow. The video illustrates the complexity of the flow field as relatively cool dyed fluid enters the room and travels along the floor as a gravity current that is warmed by the convection established above the floor buoyancy source. Fluid in the gravity current impacts the far wall, against which it rises (see the supplementary material and movies and figure 2). A portion of this rising flow travels relatively directly to the outlet vent, while a portion of it remains in the upper part of the room and a further portion slumps to form an intruding gravity current above, and oppositely directed to, the incoming current. At the same time large-scale convective motion is evident throughout the entire room.

The similarities of this flow to those driven by ‘horizontal convection’ (Hughes & Griffiths, 2008) are apparent (as discussed in § 3.3 below). However, the extent to which the vertical alignment of the low-level doorway, and the momentum associated with the horizontal inflow, might be critical in establishing the flow field that results in the horizontal convection observed is unclear. It had been established (Yang et al., 2022) that when the inflow is forced sufficiently hard horizontally, relative to the stratification generated by a localised buoyancy source, then the canonical flow patterns expected are disrupted. To enable efficient examination of a wider range of conditions, including: a different room aspect ratio, eliminating the horizontal momentum of the incoming flow and altering the relative horizontal positions of the low- and high-level vents, we now turn to the results from our investigation via numerical simulations.

3.2. Numerical results for different relative inlet to outlet vent positions

The horizontal momentum of the inflow was eliminated within the simulations by choosing to align the low-level vent horizontally within the ‘floor’ of the room domain so that the incoming ventilation flow is purely vertical. Moreover, to test the sensitivity of our findings to the room aspect ratio we simulated a room of $L/H \approx 4.35$, cf. $L/H \approx 1.50$ within the experiments.

Figure 3 shows the width-averaged buoyancy field obtained from our numerical simulations. The left-hand column plots results from the opposite-ended configuration (in which the low- and high-level vents are positioned in the region of opposite corners of the room) and the single-ended configuration results are shown in the right-hand column (in which the vents are both positioned near the left-hand end of the room, i.e. closer to $x/L = 0$, but with the low-level inflow vent near the back of the room (being $y/W = 1$) and high-level outflow near the front of the room (being $y/W = 0$)). The flow patterns of the simulations (figure 3) are comparable to those of the experiments (figure 2), with significant spatial variations in the buoyancy field, especially in the bottom third of both rooms, and are robust to the significant variation in aspect ratio between rooms (depicted to scale within these figures) and to the fact that in our simulations the incoming flow was vertical. This agreement suggests that the horizontal momentum of the inflow in our experiments does not play a significant role in establishing the dynamics within the room, and this dynamics is robust to changes in aspect ratio (at least over the range 1.50–4.35, which spans many geometries observed in rooms). Moreover, differences between the experimental data and those of the numerical simulations are of a similar magnitude to the differences that arise when the position of the upper vent in the simulations is moved from being nearer the right-hand wall furthest from the inlet, i.e. the opposite-ended configuration (as was the case in the experiments), to being nearer the left-hand wall, i.e. single-ended. For example, in the opposite-ended simulations, there is an area in the upper part of the room in which the buoyancy (temperature) is greater than any observed in the single-ended configuration – this occurs predominantly at the inlet (left-hand) side of the room.

Figure 3 shows the buoyancy field on the central vertical plane ($y/W = 0.5$), which exhibits broadly the same features as the width average, but with slightly greater variation between the two flow configurations and more evidence of convection across the plane. We return to the variations between different regions of the room in § 3.4.

The variation between the simulated opposite-ended and single-ended configurations are quite marked given that the two simulations are expected to give rise to a well-mixed room (Linden, 1999), and the configurations are notionally identical except for the horizontal location of the upper vent. We now consider the physics which give rise to these observations.

3.3. Untangling the flow

Before examining the details of the flows established, we note the following similarities to the flows typically described as horizontal convection (Hughes & Griffiths, 2008). Firstly, the overall sense (direction) of the circulation in the room is set by the location of the inlet relative to the distributed heat/buoyancy forcing – i.e. the sense of the circulation of the large-scale convective flow is set by the largest horizontal buoyancy contrast. Secondly, in a steady state, the strength of this circulation adjusts such that the incoming ambient air acquires sufficient buoyancy to rise against the far endwall, thus coupling the bulk buoyancy with the large-scale convective flow in the room. While the coupling between the bulk buoyancy and the large-scale convective flow persists, the orientation of the inlet (e.g. lying horizontal in the floor or vertical in a wall) can be expected to result in only minor changes in the flow field within the room. Furthermore, we expect this coupling to continue until the momentum flux of the inlet is sufficiently forced; which it is not in any of our cases, since the flow is naturally forced by the buoyancy. Thirdly, as a consequence of the fact that the incoming ambient air acquires sufficient buoyancy to rise against the far endwall, and hence the bulk buoyancy is coupled with the large-scale convective flow within the room, this convective flow can be anticipated to run along the full length of the floor and full height of the room, regardless of the room aspect ratio. Hence, the similarities observed




Figure 3. Scaled buoyancy fields from the numerical simulations in which a uniformly distributed heat source at floor level drives the flow through the (horizontally aligned) low- and high-level vents, located on the floor and ceiling of the room, respectively. The results presented are independent of the choice of parameters within the ranges shown in table 1 but for the simulation shown, $A_l = 0.030 \text{ m}^2$ and $A_h = 0.015 \text{ m}^2$, with $H = 2.7 \text{ m}$, taking $C_d = 0.65$ gives $A^*/H^2 = 0.012$. The top row presents, (a) the width-averaged buoyancy field $\langle \overline{\Delta T} \rangle_y$, cf. figure 2, and the bottom row presents, (b) the scaled buoyancy field at the central plane $\overline{\Delta T}(y = W/2)$. In each row, the opposite-ended (OE) configuration is shown in the left column and the single-ended (SE) configuration in the right column. Note that thick black lines mark the confined edges of the room, gaps show the positioning of the low- and high-level, inlet and outlet, vents respectively. Note that, as with the other figures included herein, data are only shown from within the room, see figure 10 for an illustration of the room within the full computational mesh.

between all the flows we examined can be expected. Moreover, the coupling of the large-scale convective flow which rises up the far wall, with the buoyancy that drives the ventilation, also underlies many of the differences we observed. For example, in the opposite-ended case the outlet vent is relatively close to the location where the large-scale convective flow rises; for the single-ended case the large-scale convective flow must return back along the length of the ceiling before reaching the outlet vent – the latter results in a large-scale overturning motion that fills the room, typical of classical horizontal convection flows.

To begin examining the physics of the flows examined herein, we present the scaled buoyancy fields in the lower portion of the room $0 \leq z/H \leq 0.2$ (figure 4). Regardless of the locations of the outlet vents the similarities between the two configurations are striking. To elucidate the buoyancy driven flow that arises, we overlay white arrows representing the velocity vectors of the width-averaged flow field. The magnitudes of, and flow patterns indicated by, these velocity vectors are very similar. In both cases, the cold air drawn in through the vent in the floor creates a rising fountain of cooler air near the inlet with relatively large vertical velocity, fluid from which slumps back and propagates along the floor in a cool gravity current. Fluid in this current is continually warmed as it propagates by the heat emitted from the floor. In a manner similar to evolution of the convective mixed layer observed in horizontal convection (Mullarney, Griffiths, & Hughes, 2004), this heat is mixed into, rather than penetrates through, the gravity current. In the region close to the endwall, the fluid within the current has warmed to a temperature of $\langle \overline{\Delta T} \rangle_y \approx 1$, i.e. close to that of the well-mixed state. Figure 5(a) illustrates the horizontal evolution of the temperature within the gravity current by presenting vertical profiles of the scaled buoyancy at four different horizontal locations.

The evolution of the (scaled) velocity profile, $\hat{u}(z)$, within this gravity current is further shown in figure 5(b). We present the (width-averaged) horizontal velocities, $\langle \hat{u} \rangle_y$, scaled by the average velocity in an idealised two-layer horizontal flow within the room, i.e. $U = 2Q/(WH)$. One can see from the




Figure 4. Scaled buoyancy fields $\langle \overline{\Delta T} \rangle_y$ from the numerical simulations presented only in the lower portion of the room, $0 \leq z/H \leq 0.2$. Velocity vectors are overlaid in white. The opposite-ended (OE) configuration is shown in the left column and the single-ended (SE) configuration in the right column. Note that in each figure thick black lines mark the confined edges of the room, gaps mark the positioning of the low-level inlet vents.




Figure 5. Vertical profiles within the lower portion of the room ($0 \leq z/H \leq 0.2$) from both the opposite-ended configuration (solid lines) and the single-sided configuration (dashed lines): (a) width-averaged scaled buoyancy $\langle \overline{\Delta T} \rangle_y$ and (b) scaled horizontal velocity $\langle \hat{u} \rangle_y$. In both panels profiles are shown at four horizontal locations: $x/L = \{0.3, 0.5, 0.6, 0.8\}$ along the length of the room.

height at which the horizontal velocities change sign, that the gravity current occupies only around the bottom 15% of the room. Moreover, the magnitudes of the scaled velocities integrated over this bottom region indicate that the full ventilation flow Q is accommodated within the gravity current. We note that if one simply looks at peak horizontal velocities, the gravity current accelerates slightly from $x/L = 0.3$ towards the middle of the centre ($x/L = 0.5$ and $x/L = 0.6$) before then decelerating slightly towards $x/L = 0.8$.

To understand how flows in the two different room configurations can be so similar in the lower portion of the room and yet be quite different higher in the room, we investigate the pressure field driving the flow. Figure 6(a) shows the width-averaged pressure within the room $\langle \bar{P} \rangle_y = P(x, z)$ relative to the ambient pressure outside, $P_a(z)$. We present all pressures scaled by the natural pressure scale $\Delta P = \rho_a(FH/A^*)^{2/3} = \rho_a b_m H$, where b_m is the buoyancy attained in the well-mixed state, so that ΔP represents the full scale of the available driving pressure from the buoyancy source within the room. Results for $\Delta \hat{P}_a = (\langle \bar{P} \rangle_y - P_a(z))/\Delta P$ (figure 6(a)) show that the pressure differences are broadly horizontally homogeneous, span a range of magnitude equal to unity, and the height of the neutral pressure level z_{npl} (at which $\Delta \hat{P}_a = 0$, see e.g. Connick & Hunt, 2020) lies below the middle of the room,




Figure 6. Width-averaged pressure fields within the room scaled by $\Delta P = \rho_a(FH/A^*)^{2/3}$. Top row, (a) scaled pressure difference relative to the ambient pressure $\Delta \hat{P}_a = [P(x, z) - P_a(z)]/\Delta P$. Middle row, (b) scaled pressure anomaly relative to the pressure within a well-mixed room $\Delta \hat{P}_m = [P(x, z) - P_m(z)]/\Delta P$, where $P_m(z) = P_a(z) - \rho_a z F/Q$. Bottom row, (c) scaled pressure anomaly $\Delta \hat{P}_m$ in the lower portion of the room, $0 \leq z/H \leq 0.2$. The left-hand panels show the opposite-ended (OE) configuration, the right-hand panels show the single-ended (SE). Thick black lines mark the confined edges of the room, gaps marks the positioning of the low- and high-level, inlet and outlet vents, respectively.

as expected when the greatest restriction to flow is provided by the high-level vent, note $A_l/A_h = 0.5$, see § 2.2. This form of the pressure field indicates why the bulk flow rates can be reasonably predicted by assuming a well-mixed state since this is driven by the integrated buoyancy.

To understand the internal flows within the room, it proves useful to examine the width-averaged pressure relative to the pressure that would be present if the room were truly well-mixed, which we denote $P_m(z)$. To do so, we exploit knowledge of the neutral pressure level which we attain from horizontally averaging $\Delta \hat{P}_a$ (figure 6a), giving $z_{npl}/H \approx 0.3$ for both the opposite-ended and the single-ended cases. We define the pressure in the well-mixed state to be $P_m(z) = P_a(z) + \rho_a(z - z_{npl})F/Q$, so that the scaled pressure anomalies are

$$\Delta \hat{P}_m = \frac{\langle \bar{P} \rangle_y - P_m(z)}{\Delta P}, \tag{3.1}$$

and at the height $z = z_{npl}$, $\Delta \hat{P}_m = 0$. Figure 6(b) shows this scaled pressure anomaly and shows that throughout the bulk of the room $\Delta \hat{P}_m \approx 0$ with the strongest deviations occurring close to the lower inlet vent. It is notable that these data show there to be an adverse horizontal pressure gradient within the upper portion of the room, such that any fluid rising at the far wall which overshoots the vent, might be expected to remain within the region above the lower inlet vent. Focusing on the lower portion of the room figure 6(c) presents this data for the range $0 \leq z/H \leq 0.2$. While figure 6(a) showed that,

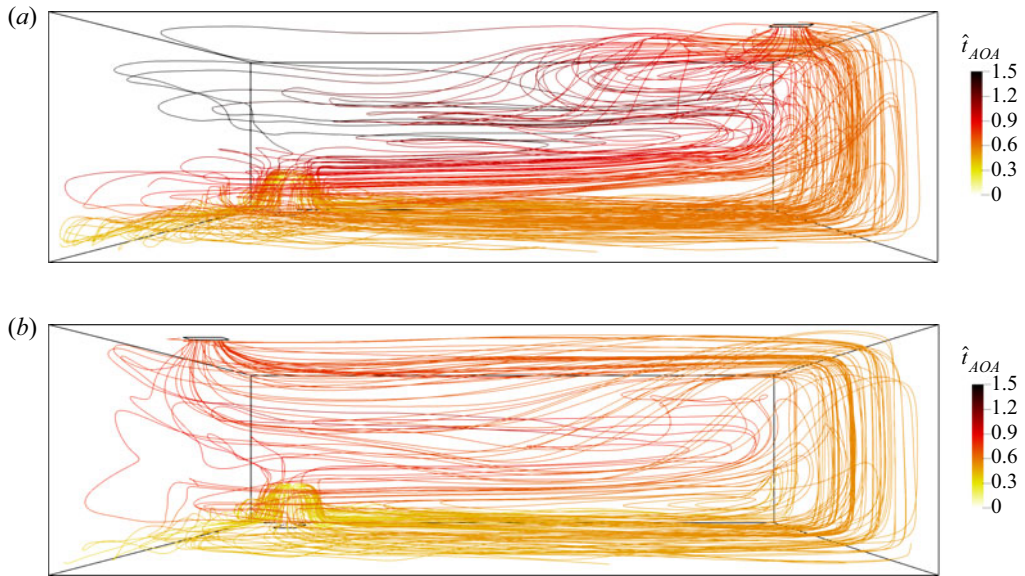


Figure 7. Three-dimensional illustration of the streamlines in (a) the opposite-ended configuration, and (b) the single-ended configuration. In both cases, the room is viewed from the front, and 100 streamlines which originate from (i.e. seeded at) the low-level inlet vent are shown. The streamlines are coloured by the scaled age of air. Note that, as with the other figures included herein, data are only shown from within the room, see figure 10 for an illustration of the room within the full computational mesh.

on average the scaled vertical pressure gradients remain close to unity in this bottom region, the data presented in figure 6(c) show that the vertical gradients in the pressure anomalies from the well-mixed state are relatively strong in this region. This is especially true in the region above the vent where $\Delta\hat{P}_m$ changes by approximately -0.07 over the height $0 \leq z/H \leq 0.2$, giving a vertical gradient of $d\Delta\hat{P}_m/d(z/H) \approx -0.3$ for the pressure anomalies from the well-mixed state, cf. $d\Delta\hat{P}_a/d(z/H) \approx 1$ for the pressure differences relative to ambient. However, within the bottom region the flow is driven by horizontal pressure gradients which are far smaller, $d\Delta\hat{P}_m/d(x/H) \approx -0.02$. Despite these horizontal pressure gradients being nearly two orders of magnitude smaller than the vertical pressure gradients that drive the bulk flow, these horizontal pressure gradients are sufficient to establish the large-scale horizontal convection within the rooms which have significant implications for the occupant experience within the room.

Although the flow patterns and temperatures in the bottom portion of the room are very similar between the two configurations, the streamlines (figure 7) show significant differences in the flow fields between the two configurations. In the case of the single-ended configuration, the relatively cool air rises at the inlet but falls back due to its buoyancy and is then carried over the full floor area, being warmed as it travels towards the far endwall. After rising near the right-hand endwall, the air (broadly speaking) travels back along the length of the room in the upper part of the room and out through the high-level vent. In contrast, in the opposite-ended configuration, the warmed fluid rises in the region near the right-hand endwall and predominantly travels up and out the high-level vent at the far end of the room. As a consequence, a relatively low volume flux of fluid travels through the region in the upper portion of the room above the region of the inlet vent, creating a re-circulation or ‘dead’ zone. We now consider these findings in the context of the potential impact on occupants. Examination of the vorticity field highlighted similarities between the opposite-ended and single-ended configurations within the lower layer; similarities which were reflected, when viewed relative to the outlet, in the upper portion of the room too.




Figure 8. The (scaled) age of air \hat{t}_{AoA} at a plane through the outlets ($y = 0.8 W$); within the opposite-ended (left OE) and single-ended (right SE) configurations.




Figure 9. Histogram showing the probability density function of the scaled age of air \hat{t}_{AoA} within the room. The colours highlight the portion associated with three vertical regions: green – top, $T(2H/3 \leq z \leq H)$; red – middle, $M(H/3 \leq z \leq 2H/3)$; and blue – bottom, $B(0 \leq z \leq H/3)$. In the opposite-ended (OE) configuration the mean in each vertical region is $\mu_T = 1.24$, $\mu_M = 1.11$ and $\mu_B = 0.74$ with respective standard deviation $\sigma_T = 0.34$, $\sigma_M = 0.26$ and $\sigma_B = 0.16$. In the single-ended (SE) configuration it is $\mu_T = 0.77$, $\mu_M = 0.79$ and $\mu_B = 0.63$ with $\sigma_T = 0.11$, $\sigma_M = 0.11$ and $\sigma_B = 0.12$.

3.4. Implications for thermal comfort and exposure

Although the buoyancy field indicates variation in occupant thermal comfort, a more pertinent measure of indoor air quality and occupant exposure to pollutants is the scaled age of air \hat{t}_{AoA} which determines how long fluid has remained within the room relative to the bulk ventilation air-change rate Q/V (see § 2 and Sandberg, 1981). Figure 8 plots the scale age of air on a vertical slice through the central plane of the outlet vent, data are presented for both configurations – differences are notable. In the opposite-ended configuration, the age of air in the upper part of the room above the lower inlet vent is ≈ 1.5 , clear evidence of the impact of the dead zone shown in the top left corner of figure 7(a), with no such counterpart in the single-ended configuration. Despite these differences we find that the age of air at the outlet is almost identical for both configurations, having values 0.775 and 0.774 for the opposite-ended and the single-ended cases, respectively, (see figure 8). It is unclear whether this agreement is a coincidence or a feature of the flows having the same ventilation rate.

Histograms of the scaled age of air are plotted in figure 9. The mean (scaled) age of air within the room is 1.01 ± 0.35 for the opposite-ended configuration and 0.71 ± 0.14 for the single-ended (with the bound indicating one standard deviation) – these differences in mean and variance are both statistically significant. Figure 9 also illustrates that, for the single-ended case, only approximately 6% of the air within the room exhibits $\hat{t}_{AoA} \geq 1$; i.e. the fluid within the room that has had time to fully circulate the room once. All other fluid within the room must be, on average, in the process of being efficiently circulated around and out of the room. In contrast, not only does the age of air distribution for the opposite-ended case show a much wider variance, with around 20% of the air within the room satisfying

$\hat{t}_{A_{0A}} \geq 1.5$. This suggests the potential for ‘dead zones’ (regions in which fluid is largely recirculated) is more prevalent in the opposite-ended configuration than in the single-ended configuration. Note also that all of the fluid with $\hat{t}_{A_{0A}} \geq 1$ lies in the upper two thirds of the room but, perhaps more pertinent for exposures, much of this more stale air lies in the middle two thirds of the room, $0.9 \text{ m} \leq z \leq 1.8 \text{ m}$, which includes the breathing zone.

4. Conclusions

This paper reports the results of laboratory experiments and corresponding computations of the flow produced by a uniformly heated floor in a ventilated room. The room is connected to the exterior by upper and lower vents, and the heated floor produces natural displacement ventilation where cool air enters through the low-level vent and leaves via the high-level vent. We focused on the effects of the relative location of the upper vent in relation to the lower vent by considering two cases: the single-ended case where the outlet is at the same end of the room as the inlet, and the opposite-ended case where the outlet vent is at the other end of the room. Two rooms were studied, a relatively short room with length 1.7 times the height and a long room with length 3.7 times the height.

In all cases, large-scale convective flows travelled along the full length of floor and up the far wall, and all our findings were, within the bounds tested, insensitive to whether the momentum of the incoming cool air is vertical or horizontal, and insensitive to the aspect ratio of the room – these facts can be understood as follows. Firstly, we observe that the overall sense of the circulation in the room is set by the inlet location relative to the distributed heat/buoyancy forcing (due to this setting the vertical level of the greatest horizontal buoyancy contrast). The steady-state strength of circulation adjusts such that the incoming ambient air acquires sufficient buoyancy to rise against the endwall, thus coupling the bulk buoyancy with the large-scale convective flow in the room; a coupling which, for these buoyancy-driven flows, results in the momentum flux associated with the inflow being of little consequence irrespective of the orientation of the vent. Furthermore, since the incoming ambient air acquires sufficient buoyancy to rise against the endwall, this large-scale convective flow can be expected to run along the full length and height of the room, regardless of the room aspect ratio.

A key conclusion of this research, which follows from the physics of this large-scale convective flow, is that, even though, as expected, the ventilation rate of the room does not depend on the location of the outlet vents, the flow within the upper part of the room is quite different in the single-ended and opposite-ended cases. In the opposite-ended case, the outlet vent is relatively close to the location where the large-scale convective flow rises; for the single-ended case the large-scale convective flow must return along the length of the ceiling before reaching the outlet vent – the latter results in a large-scale overturning motion that fills the room, typical of classical horizontal convection flows. In the single-ended case the upper regions of the room are relatively uniform in temperature, while in the opposite-ended case there are significant differences in temperature from one end to the other. In this latter case the results exhibit spatial inhomogeneities of the buoyancy scalar of the order of 20 % of the mean buoyancy.

These inhomogeneities have implications for the exposures that occupants might experience if positioned at different locations within the room. These findings are supported by calculations of the age of air and the streamline patterns, which show that a relatively stagnant ‘dead zone’ is formed above the inlet in the opposite-ended case. While most of the air rising up the far endwall exits at the outlet vent a small amount travels back towards the inlet in the upper two thirds of the room and accumulates there. Over time this leads to an accumulation of buoyancy in this upper region. The horizontal variation would imply that pollutant entering from outside would also accumulate and have concentrations above the mean for the room as a whole. Similarly, these variations have implications for indoor sources of pollution and the spread of airborne biological material released inside the room.

Although the cases considered in this paper are somewhat idealised, one should expect the pattern of the large-scale convection observed to be relatively robust and hence one could expect flow patterns similar to those reported herein to be realised in some operational rooms. For example, based on the

understanding of horizontal convection (e.g. Hughes & Griffiths, 2008), the flows observed are expected to be robust to the inclusion of additional localised heat sources. This can be expected because: firstly, the time scale for the fluid above the horizontal buoyancy source to be swept to the endwall of the room is comparable to the time scale of rise through the lower stratified region of buoyant parcels so that, with the associated vertical and horizontal length scales, the momentum of the horizontal flow remains dominant in organising the overall circulation; secondly, the lower stratified region acts to confine vertically parcels that are initially buoyant at the horizontal boundary; and finally, in the presence of such statistically steady convective flows, the buoyancy of the warmest fluid in the room interior is approximately equal to (i.e. is set by) that of the fluid above the bottom horizontal boundary adjacent to the far endwall. These effects will act to render the convective circulation relatively robust to perturbations that arise due to the presence of localised heat sources. As such, our findings have three further implications that are worth noting. First, point measurements of temperature, carbon dioxide and pollutant concentration, commonplace in the assessment of building spaces, should be treated with caution as they may not be representative of the room as a whole. Second, the results show in the cases investigated, that a ‘single-ended’ displacement design of positioning all low- and high-level vents at one end of the room provides a more uniformly mixed indoor environment. This might deserve consideration when revising any design guidance which typically promotes opposite-ended strategies for natural ventilation, as a result of expected contributions by the wind enhancing the bulk ventilation flow rate (e.g. CIBSE, 2005). Third, we acknowledge that our findings increase the complexity of making useful predictions from simple design rules. However, they do highlight the rich variety of dynamics that can arise in indoor air flows and the value of deepening our understanding as we spend ever more time indoors. Ultimately, our results set a challenge to understand and classify indoor air flows into a broader set of classes and, for each, to quantify the typical magnitudes of the variations in scalar concentrations that can be expected to arise.

Supplementary material and movies. Supplementary material and movies are available at <https://doi.org/10.1017/flo.2023.11>.

Acknowledgements. C.V.M.V., H.C.B. and M.v.R. gratefully acknowledge F. Mendonça and his team at OpenCFD for providing support and advice on setting up the numerical simulations. C.V.M.V., H.C.B. and P.F.L. acknowledge the support of the SAMHE Project consortium.

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

Funding statement. The PhD works of C.V.M.V. and T.D.H., data from which inspired and underpinned this work, were funded by the Engineering and Physical Sciences Research Council (EPSRC) via the Imperial College London Centres for Doctoral Training in: Fluid Dynamics Across Scales (grant EP/L016230/1), and Sustainable Civil Engineering (grant EP/R512540/1), respectively. The contributions from H.C.B. and P.F.L. were funded by the EPSRC COVID-19 Transmission Risk Assessment Case Studies – education Establishments (CO-TRACE) project, and the extension, School Air quality Monitoring for Health and Education, SAMHE (both under grant EP/W001411/1). H.C.B. was further supported in his role within WP2.2.2: Ventilation Effects, within the PROTECT COVID-19 National Core Study on transmission and environment, managed by the Health and Safety Executive on behalf of HM Government. Computational resources were provided by the UK Turbulence Consortium (grant EP/R029326/1).

Author contributions. Through their PhD research, C.V.M.V. performed all simulations, and T.D.H. performed all experiments – including carrying out their respective analyses through to the creation of the plots included. H.C.B. supervised the PhD work of C.V.V. and T.D.H., with co-supervision provided by M.v.R. and G.O.H., respectively. H.C.B. supervised the data presentation and wrote the manuscript. P.F.L. identified the value of the finding reported herein, thus inspiring the creation of this paper. All authors edited the text, providing invaluable insights and suggestions for the analysis.

Data availability statement. Raw data are available from the corresponding author (H.C.B.).

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

Appendix A. Details of the laboratory experiments

A cuboid, as a small-scale model of a room, was connected to the ambient environment (a visualisation tank of cross-section 1.3 m × 1.3 m filled with fresh water to a depth of 1.25 m) by two openings, a

low-level inlet vent (aligned in the vertical plane much like a small doorway) and a high-level outlet vent, the positions of the vents are illustrated in [figure 2](#) – see [table 1](#) for details. The high-level outlet was positioned in the ceiling approximately two thirds of the way along the room from the inlet, and hence the bulk flow must be both upwards and across the room. A saline solution of buoyancy $b_0 = 0.35 \pm 0.01 \text{ m s}^{-2}$ (the density of which was measured using an Anton Paar DMA 4500 density meter with precision $1 \times 10^{-5} \text{ g cm}^{-3}$) was supplied to the scale model at floor level across a porous plastic (high density polyethylene) panel (Vyon F 6 mm thickness) with a flow rate of $Q_0 = 1.67 \times 10^{-5} \text{ m}^3 \text{ s}^{-1}$ (controlled by a peristaltic pump (Cole Parmer Masterflex L/S) and monitored using an in-line ultrasonic flow meter (Atrato 740-V10-D)). The source volume flux was always less than 1% of the predicted bulk ventilation rate ([Gladstone & Woods, 2000](#)) and was therefore assumed to be largely representative of the flow pattern in the full-scale application. A two-chamber system was used to provide a source buoyancy flux that was evenly distributed over the floor area. The principle of this system was to ensure that any horizontal pressure gradients across the source were negligible compared with the pressure gradients imposed by the pump to drive the flow across the porous plastic panel of the source (see [Higton \(2022\)](#), for further details).

The source solution was dyed with a red food colouring, Allura red (AC, E129), to enable the width-averaged two-dimensional buoyancy field $\langle b \rangle_y = b(x, z, t)$ to be measured using the dye-attenuation technique ([Allgayer & Hunt, 2012](#); [Cenedese & Dalziel, 1998](#)). The model ‘room’ was illuminated from behind using a $1 \text{ m} \times 1 \text{ m}$ light panel (Applelec LED light sheet), and the light attenuated by the dyed fluid in the model was measured by a Sony A7Rii digital camera (14-bit depth RGB images with a spatial resolution of approximately $0.14 \text{ mm pixel}^{-1}$). Note that experimental data directly adjacent to the walls are compromised due to parallax errors, camera viewing restrictions and reflections from the walls and hence these regions are excluded from the results presented – these regions are highlighted in [figure 2](#).

Appendix B. Details of the numerical simulations

Numerical simulations were conducted with OpenFOAMv2106 using the transient buoyantPimpleFoam solver and the $k-\omega$ SST turbulence model. The chosen OpenFOAM version also incorporates the effects of buoyancy on turbulence production by using the buoyancyTurbSource finite volume option ([OpenCFD, 2020](#)). Full details of the governing equations and methods of solution are available in the OpenFOAM documentation ([OpenCFD, 2021](#)).

The computational domain includes the room, a cube of dimensions $10 \text{ m} \times 5.5 \text{ m} \times 2.7 \text{ m}$, along with two exterior boxes linked to the room through inlets and outlets (see [figure 10](#)). These two boxes were included in order to properly model the effect of flow at the two openings and the resulting flow in the room. The bottom exterior box is centred around the inlet and has dimensions $2.8 \text{ m} \times 2.5 \text{ m} \times 1 \text{ m}$. The top exterior covers the same surface as the room with dimensions $10 \text{ m} \times 5.5 \text{ m} \times 3.7 \text{ m}$. The inlet and outlet connecting the exterior boxes to the room have the same size as the vents with respective areas A_l and A_h , and height 0.2 m . The sensitivity of the simulations to the size of the external boxes was tested and these dimensions were chosen as they did not impact the results. The mesh is defined as a perfect orthogonal mesh with $\Delta x = \Delta y = \Delta z = 0.05 \text{ m}$ in the room, inlet and outlet and $\Delta x = \Delta y = \Delta z = 0.1 \text{ m}$ in the exterior, requiring computations to be executed on 1.4×10^6 hexahedral grid cells. A grid convergence study was also performed: the mesh used in this study accurately represents the bulk flow as well as the buoyancy and tracer distribution when compared with a finer mesh (with $\Delta x = \Delta y = \Delta z = 0.025 \text{ m}$), it allows us to capture the variations in the flow pattern observed in the single-ended and opposite-ended configurations. Full details are presented by [Vouriot \(2022\)](#).

At all walls, including the room and the exterior boxes (apart from the domain inflow and outflow), a no-slip velocity boundary condition is used. All walls are also assumed to be adiabatic, with the exception of the floor where a constant 6200 W heat input is imposed. At the domain inflow the temperature is set to ambient (defined as 278 K in this simulation) and at the domain outflow a Neumann boundary condition is used setting the gradient of the temperature. Velocity boundary conditions at the inflow and outflow are calculated from the pressure field where the pressure difference $\Delta p_0 = -\rho_a g H_{\text{domain}}$ is




Figure 10. Computational grid used in the numerical simulations shown for the single-ended configuration, including the modelled exterior, inlet, outlet and room.

imposed across the domain inflow and outflow. The boundary layer region is not fully resolved, instead wall functions are used with standard parameters for: the turbulent thermal diffusivity, the turbulent kinetic energy, the turbulent viscosity and specific dissipation rate. The stratification is first established using a steady run, the simulations are then run with a transient solver for 8600 s with statistics averaged only over the last 3600 s. A transient solver had to be used due the presence of long-time period fluctuations (potentially caused by internal wave modes) observed in the results generated by steady solvers. Second-order backward finite difference time marching and second-order central differencing schemes are used. Full details and the specific OpenFOAM set-up used are described by [Vouriot \(2022\)](#).

Appendix C. Sensitivity of the dimensionless volume flux on the choice of discharge coefficient

For completeness, the values of A^* , and thereby \hat{Q} , depend on C_d , and in [table 1](#) we took $C_d = 0.65$. Should we have taken $C_d = 0.6$ then $1.00 \leq \hat{Q} \leq 1.08$ for the experiments and $1.05 \leq \hat{Q} \leq 1.12$ for the simulations. Conversely, taking $C_d = 0.7$ would then give $0.91 \leq \hat{Q} \leq 0.98$ for the experiments and $0.95 \leq \hat{Q} \leq 1.02$ for the simulations.

References

- Allgayer, D.M., & Hunt, G.R. (2012). On the application of the light-attenuation technique as a tool for non-intrusive buoyancy measurements. *Experimental Thermal and Fluid Science*, 38, 257–261.
- Bhagat, R.K., & Linden, P.F. (2020). Displacement ventilation: A viable ventilation strategy for makeshift hospitals and public buildings to contain COVID-19 and other airborne diseases. *Royal Society Open Science*, 7(9), 200680.
- Cenedese, C., & Dalziel, S.B. (1998). Concentration and depth field determined by the light transmitted through a dyed solution. In *Proceedings of the 8th international symposium on flow visualization, Sorrento, Italy, vol. 8*, pp. 1–37.
- CIBSE (2005). *Chartered institution of building services engineers, applications manual 10: Natural ventilation in non-domestic buildings*. Chartered Institution of Building Services Engineers.
- Connick, O., & Hunt, G.R. (2020). Hybrid ventilation of a room: A theoretical model for the combined effects of mechanically-imposed and buoyancy-induced driving pressures. *Building and Environment*, 169, 106546.
- Gladstone, C., & Woods, A.W. (2000). On the application of box models to particle-driven gravity currents. *Journal of Fluid Mechanics*, 416, 187–195.

- Gladstone, C., & Woods, A.W. (2001). On buoyancy-driven natural ventilation of a room with a heated floor. *Journal of Fluid Mechanics*, 441, 293–314.
- Higton, T.D. (2022). *The fluid mechanics of horizontal exchange flows in naturally ventilated buildings* (PhD thesis). Imperial College London, London, UK.
- Hughes, G.O., & Griffiths, R.W. (2008). Horizontal convection. *Annual Review of Fluid Mechanics*, 40, 185–208.
- Hunt, G., & Linden, P.F. (1999). The fluid mechanics of natural ventilation—displacement ventilation by buoyancy-driven flows assisted by wind. *Building and Environment*, 34(6), 707–720.
- Linden, P.F. (1999). The fluid mechanics of natural ventilation. *Annual Review of Fluid Mechanics*, 31, 201–238.
- Linden, P.F., Lane-Serff, G.F., & Smeed, D.A. (1990). Emptying filling boxes, the fluid mechanics of natural ventilation. *Journal of Fluid Mechanics*, 212, 309–335.
- Mullarney, J.C., Griffiths, R.W., & Hughes, G.O. (2004). Convection driven by differential heating at a horizontal boundary. *Journal of Fluid Mechanics*, 516, 181–209.
- OpenCFD (2020). buoyancyTurbSource finite-volume option. Retrieved from <https://www.openfoam.com/news/main-news/openfoam-v20-12/solver-and-physics-solver-and-physics-buoyancy-fvoption>.
- OpenCFD (2021). ESI OpenCFD Release OpenFOAM® v2106. Retrieved from <https://www.openfoam.com/news/main-news/openfoam-v2106>.
- Sandberg, M. (1981). What is ventilation efficiency? *Building and Environment*, 16(2), 123–135.
- Vouriot, C.V.M. (2022). *The potential for spatial variations in exposure and infection risk within a typical UK classroom: COVID-19 as a case study* (PhD thesis). Imperial College London, London, UK.
- Wachenfeldt, B.J., Mysen, M., & Schild, P.G. (2007). Air flow rates and energy saving potential in schools with demand-controlled displacement ventilation. *Energy and Buildings*, 39(10), 1073–1079.
- Yang, R., Ng, C.S., Chong, K.L., Verzicco, R., & Lohse, D. (2022). Do increased flow rates in displacement ventilation always lead to better results? *Journal of Fluid Mechanics*, 932, A3.