

Time-dependent ventilation flows driven by opposing wind and buoyancy

I. A. COOMARASWAMY¹ AND C. P. CAULFIELD^{2,1†}

¹Department of Applied Mathematics and Theoretical Physics, University of Cambridge,
Centre for Mathematical Sciences, Wilberforce Road, Cambridge CB3 0WA, UK

²BP Institute, University of Cambridge, Madingley Rise, Madingley Road, Cambridge CB3 0EZ, UK

(Received 7 September 2010; revised 21 October 2010; accepted 6 November 2010;
first published online 21 February 2011)

We consider transient flow in a box containing an isolated buoyancy source, ventilated by a windward high-level opening and a leeward low-level opening, so that prevailing wind acts to oppose buoyancy-driven flow. Hunt & Linden (*J. Fluid Mech.*, vol. 527, 2005, p. 27) demonstrated that two stable steady states can exist above a critical wind strength: buoyancy-driven displacement ventilation with a two-layer stratification and wind-driven mixing ventilation with the whole interior contaminated by buoyant fluid. We present two time-dependent models for this system: a nonlinear ordinary differential equation (ODE) model following Kaye & Hunt (*J. Fluid Mech.*, vol. 520, 2004, p. 135), assuming ‘perfect’ vertical mixing of fluid within each layer, and a partial differential equation model assuming zero vertical mixing, following Germeles (*J. Fluid Mech.*, vol. 71, 1975, p. 601).

We apply these models to an initial-value problem – the filling box with constant opposing wind. The interface between the upper hot plume fluid and the lower cool ambient air can dramatically overshoot its final level before relaxing to equilibrium; in some cases, a fully contaminated transient can occur before the buoyancy-driven two-layer steady state is reached. However, we find that for an initially completely uncontaminated box, the system converges to a stable wind-driven steady state whenever it exists. By analysing phase diagrams of the ODE model for the flow, we establish a general method of determining which final state is attained and also explain the hysteresis observed by Hunt & Linden (2005). We confirm these transient behaviours by conducting salt bath experiments in a recirculating flume tank and establish quantitative agreement between theory and experiment. Our ‘zero mixing’ model is more accurate than our ‘perfect mixing’ model for our experiments, as the upper layer remains stratified for a substantial time.

Key words: convection in cavities, plumes/thermals, turbulent mixing

1. Introduction

In many developed countries, a large proportion of final energy consumption can be attributed to maintaining thermal comfort and air quality within residential and commercial buildings. In particular, mechanical ventilation systems in large-scale buildings entail a significant energy cost. In the United Kingdom, as in other countries with temperate climates, there is considerable scope for relying instead on natural

† Email address for correspondence: c.p.caulfield@bpi.cam.ac.uk

ventilation schemes. With growing concern about climate change and consensus over the need to move towards a low carbon economy (Stern 2007), the energy savings inherent in the use of natural ventilation mean that it can play a vital part in this transition.

However, airflow in naturally ventilated buildings is currently poorly understood, with many such buildings exhibiting unusual and undesirable patterns of thermal behaviour. Reynolds and Péclet numbers are typically of the order of 10^4 or higher, flows are turbulent and confined by complex geometries. Direct numerical simulation of such flows requires accurate modelling of dynamics over a wide range of length and time scales, and is thus extremely computationally costly. Hence, many studies of natural ventilation have attempted to develop reduced models of interior flows that capture the key flow phenomena relevant to building design (Linden 1999). To date, most have focused on describing steady states reached in the presence of steady forcing. In reality, however, heat loads, wind speeds and opening sizes may all vary with time, and achieving an understanding of the complex time-dependent phenomena that can occur is vital for design purposes. An instructive example is that of a lecture theatre with both low-level and high-level air vents (Hunt & Linden 1998; Linden 1999). Were 500 people to enter such an auditorium and sit down, they would collectively provide a steady heat source that in time might result in a steady internal stratification. Typical engineering practice is to design according to known results for the steady-state conditions. However, simple modelling indicates that the time taken to reach this equilibrium may be as much as an hour. This provides a clear illustration of why theoretical models need to capture transient behaviour accurately in order to close the gap to design reality.

Linden *et al.* (1990) established a paradigm for the study of displacement ventilation flows. They addressed the case of a room with one floor-level and one ceiling-level opening, containing a single localised heat source at floor level. To simplify matters, they considered a cuboidal box containing a point source of pure buoyancy. They showed that a balance between a plume filling box flow and a draining box flow results in a two-layer steady state in which an upper layer of warm plume fluid sits above a lower layer of cool ambient fluid drawn in through the bottom opening (see figure 1*a*). Drawing on the classical plume theory of Morton, Taylor & Turner (1956, henceforth ‘MTT’) as well as the ‘filling box’ studies of Baines & Turner (1969) and Worster & Huppert (1983), they showed that the steady-state interface height h is independent of the source buoyancy flux B_0 . For a box of total height H , $\hat{h} = h/H$ satisfies

$$V = \left(\frac{A^*}{C^{3/2} H^2} \right)^2 = \frac{\hat{h}^5}{(1 - \hat{h})}, \quad C = \left(\frac{6\alpha}{5} \right) \left(\frac{9\alpha\pi^2}{10} \right)^{1/3}, \quad (1.1)$$

where α is the ‘top-hat’ entrainment constant for a plume, assumed ≈ 0.1 (see Linden 2000; Turner 2006), A^* is the ‘effective area’ of the openings and V thus may be thought of as an appropriate non-dimensional parameter describing the vent area. (Note that $V = 1/\mu^2$ in the notation of Kaye & Hunt 2004 and others.) If the top opening has area A_T , discharge coefficient c_T , and the bottom opening has area A_B , discharge coefficient c_B (where each discharge coefficient accounts for the dissipative losses of flow through the associated opening, dependent on its precise geometry and the flow pattern around it – see for example Hunt & Holford 2000), then

$$A^* = \frac{\sqrt{2} c_B A_B c_T A_T}{\sqrt{c_B^2 A_B^2 + c_T^2 A_T^2}}. \quad (1.2)$$

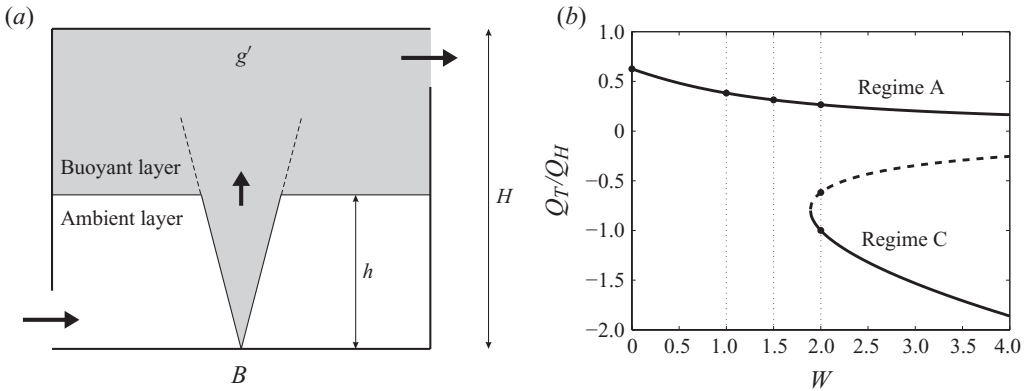


FIGURE 1. (a) A schematic diagram for displacement ventilation: an ‘emptying filling box’ with forward flow. (b) Hunt & Linden (2005)’s steady states for opposing wind. The scaled volume flux out of the top opening Q_T/Q_H (where Q_H is defined in (1.7)) is plotted as a function of W (as defined in (1.5)) with $V = 1$ (as defined in (1.1)). A positive flux (forward flow) solution exists for all values of W . Two negative flux (reverse flow) solutions exist only for $W^3 V > 27/4$; one is unstable (dashed line) and the other is stable (full line). Dotted vertical lines indicate the specific cases considered in detail later, whose time-dependent properties are plotted in figure 3.

If we denote the density of the ambient fluid ρ and that of the upper layer $\rho - \Delta\rho$, then at equilibrium, the reduced gravity $g' = g\Delta\rho/\rho$ in the upper layer is equal to its value within the plume at a height h above the floor. Thus, we have

$$g' = g'_p = C^{-1} B_0^{2/3} h^{-5/3}, \tag{1.3}$$

where we have assumed that the plume properties are determined by the classical similarity solutions of MTT. The volume flux Q_T through the top opening balances the volume flux Q_p in the plume at height h , so that

$$Q_T = A^* [g'(H - h)]^{1/2} = Q_p = C B_0^{1/3} h^{5/3}, \tag{1.4}$$

where once again Q_p is given by the MTT similarity solutions. Linden *et al.* (1990) also conducted salt bath experiments, which confirmed the existence of a steady-state ‘displacement ventilation’ flow.

If there is a prevailing wind blowing around a building, depending on the position of the openings, it can either assist the action of buoyancy (Hunt & Linden 1999, 2001) or oppose it (Hunt & Linden 2000, 2005; Li & Delsante 2001). In the latter situation, multiple steady states can exist for the same wind forcing, heat load and room geometry. For the case of an isolated source of buoyancy, Hunt & Linden (2005, henceforth HL05) showed that a stratified steady state with flow in the forward (buoyancy-driven) direction theoretically exists for all wind strengths. Since the reduced gravity in an ideal plume is singular as $h \rightarrow 0$ (1.3), if the interface approaches the floor of the box, the buoyancy of the upper layer can be arbitrarily large, forcing the fluid out of the top opening irrespective of the strength of the opposing wind. However, they demonstrated that for sufficiently strong wind, two wind-driven steady states (one stable and the other unstable) can also exist. In these states, flow is in the ‘reverse’ direction, with cool ambient air blown into the top opening and warm air expelled from the bottom one. The interior of the box is fully

contaminated with plume fluid and assumed to be a single well-mixed layer, leading to the existence of a steady-state ‘mixing ventilation’ flow.

This multiplicity of solutions is illustrated in figure 1(b), in which we plot volume flux out of the top opening as a function of wind strength for a single effective opening area (cf. figure 2 of HL05). We here define the dimensionless opening parameter V as in (1.1) and the dimensionless wind strength W as

$$W = \frac{\Delta p / \rho}{g'_H H}. \quad (1.5)$$

(Note that in the notation of HL05, $W = Fr^2$.) Here, Δp is the difference in ambient pressure between the windward and leeward sides of the building, and g'_H is the reduced gravity in the plume when it first reaches the top of an empty box,

$$g'_H = g'_P|_{h=H} = C^{-1} B_0^{2/3} H^{-5/3}. \quad (1.6)$$

Volume flux through the top opening is scaled by Q_H , the volume flux in the plume when it first reaches the top of an empty box,

$$Q_H = Q_P|_{h=H} = C B_0^{1/3} H^{5/3}. \quad (1.7)$$

With respect to these newly defined parameters, the condition for the existence of multiple steady states becomes

$$W^3 V > 27/4. \quad (1.8)$$

The existence of multiple steady states for a given forcing was confirmed experimentally by HL05. The state adopted by the system was observed to be dependent on the time-history of the flow, with the transitions between states exhibiting hysteresis.

If a room contains a distributed rather than localised heat source at floor level, plume-based displacement ventilation is not possible and it is then often assumed that the interior is a single zone of uniform temperature regardless of the direction of flow through it. For this case, the multiplicity of solutions in the presence of opposing wind was established by Li & Delsante (2001), Li *et al.* (2001) and Heiselberg *et al.* (2004), while time-dependent models for investigating transitions between them have been developed by Yuan & Glicksman (2007), Yuan & Glicksman (2008) and Lishman & Woods (2009). Thus far, however, the more complex transient dynamics of the localised source scenario have not been investigated.

Kaye & Hunt (2004, henceforth referred to as KH04) examined transient behaviour in an emptying filling box in the absence of wind. Their model adopts a ‘quasi-steady approach,’ in which the upper layer of plume fluid is assumed to be of uniform density at all times. This is effectively an assumption of instantaneous and ‘perfect’ mixing throughout this upper layer, valid in situations in which the time scale for mixing (due to diffusion and the secondary flows caused by entrainment into the plume) is much shorter than the time scale for movement of the interface. They studied the evolution of the system from an empty box to the steady state predicted by Linden *et al.* (1990) and observed that for a sufficiently small dimensionless opening area ($1/\mu$ in their notation, V in ours, as defined in (1.1)), the system is underdamped and the interface overshoots in equilibrium height. However, they also observed that while the size of overshoot varies with V , it is at most 3.7% of the total box height.

This overshoot phenomenon would therefore appear to be too small to be of interest to ventilation engineers. A similar ‘quasi-steady’ approach was employed by Bolster, Maillard & Linden (2008) and Bower *et al.* (2008) for investigating the response of an emptying filling box system at steady state to sudden changes in buoyancy source strength. In this case, there are again small ‘overshoots’ in interface height before the reduced gravity of the upper layer reaches its new equilibrium value. The presence of opposing wind can change matters dramatically, however, as shown later in this paper.

An alternative approach to modelling transient behaviour is to employ the type of numerical scheme developed by Germeles (1975). This has been used in the context of filling flows by Worster & Huppert (1983) and Caulfield & Woods (2002) and in the context of emptying filling flows by Flynn & Caulfield (2006*a*), Bolster *et al.* (2008) and Bolster & Caulfield (2008). This method assumes that horizontal surfaces of constant density are advected downward by a vertical bulk return flow that balances the upward plume volume flux. It therefore assumes that there is no (vertical) mixing of fluid by diffusion or turbulent eddies, and treats each ‘layer’ of fluid as a characteristic, whose location is then tracked with time. As such, the ‘perfect mixing’ model of KH04 and the ‘zero mixing’ model of Germeles (1975) (henceforth ‘KH model’ and ‘Germeles model’ respectively) constitute end members of the possible class of models for the evolution of density stratification in the interior.

Our aim in this paper is to investigate the transient behaviour of an emptying filling box flow in the presence of steady opposing wind. We develop two theoretical models for such a flow. In §2, we derive a perfect-mixing KH model, which consists of a coupled pair of nonlinear ordinary differential equations (ODEs) for interface height and reduced gravity. In §3, we present a zero-mixing Germeles model, a partial differential equation (PDE) scheme for plume fluxes and reduced gravity in the hot layer as a function of both time and vertical position. We apply our models to a set of initial conditions – an empty box – and present the results in §4. We discuss both the transient phenomena seen and the final state reached for all physically relevant regions of the V – W parameter space, including those in which HL05 have demonstrated that two stable equilibria exist. A phase space analysis of the equations governing the KH model is then used to predict which final state is attained for more general initial conditions. In §5, we describe the experimental procedures we employ to simulate these flows. We adopt the ‘salt bath modelling’ approach of Linden *et al.* (1990), in which dynamic similarity with ventilation flows is achieved on a reduced laboratory scale by using water as a working fluid and salinity differences to create the required buoyancy forces. We also make use of a novel method of flow visualisation by dye attenuation, in which the apparatus is backlit using an electroluminescent tape (Dalziel *et al.* 2008). In §6, we compare our experimental results with the predictions of our theoretical models before discussing the overall conclusions in §7.

2. Perfect-mixing KH model

We extend KH04’s transient model, adapting expressions for flux through openings to include the effect of wind. As well as the forward (buoyancy-driven) flow regime (regime A), we also consider two reverse (wind-driven) flow regimes (see figure 2). In the first reverse flow regime (regime B), the two-layer stratification is maintained. We assume here that if a cold ambient fluid is blown into the top opening, this fluid rapidly mixes with and cools the upper layer without overturning the whole




FIGURE 2. Schematic diagrams of the possible flow regimes. A: buoyancy-dominated forward flow with two-layer stratification. B: wind-dominated reverse flow with two-layer stratification still intact (note that no steady state can exist in this regime – it is purely transient). C: wind-dominated reverse flow with a fully contaminated interior.

box. However, if this regime persists for some time, the interface will inevitably descend to the floor of the box, so that the whole interior becomes contaminated with plume fluid and warm air is expelled from the bottom opening. Our second reverse flow regime (regime C), discussed in more detail below, describes the subsequent evolution of this fully contaminated state. The buoyancy-driven and (stable) wind-driven steady states shown in figure 1(b) thus correspond to regimes A and C, respectively.

2.1. Description of the problem

We consider a box of height H and uniform cross-sectional area A_C , containing a point source of pure buoyancy flux of strength B_0 . The box has one ceiling-level and one floor-level opening, which together have an ‘effective area’ A^* , as defined in (1.2). The balance between the ‘filling box’ flow driven by the turbulent plume rising from the buoyancy source and the ‘emptying’ flow through the openings results in a two-layer stratification with interface height h . We assume that outside the plume, each layer is well-mixed, and the only exchange of fluid between the layers occurs through the plume. The lower layer contains air of ambient density ρ , while the upper layer is of density $(\rho - \Delta\rho)$. The reduced gravity in the upper layer is $g' = \Delta\rho/\rho$. The terms Q_P and g'_P are the volume flux and the reduced gravity within the plume at height h , respectively, and Q_T is the volume flux out of the top opening. An opposing wind causes a difference Δp in the ambient pressure between the top and bottom openings.

2.2. Governing equations

For the upper layer, conservation of mass yields

$$A_C \frac{d}{dt} [(H - h)] = Q_P - Q_T, \quad (2.1)$$

and similarly conservation of buoyancy yields two equivalent expressions (using (2.1)):

$$A_C \frac{d}{dt} [g'(H - h)] = \begin{cases} Q_P g'_P - Q_T g', & \text{for } Q_T > 0 \\ Q_P g'_P, & \text{for } Q_T < 0 \end{cases} \quad (2.2a)$$

$$A_C (H - h) \frac{dg'}{dt} = \begin{cases} Q_P (g'_P - g'), & \text{for } Q_T > 0 \\ Q_P (g'_P - g') + Q_T g', & \text{for } Q_T < 0. \end{cases} \quad (2.2b)$$

Note that $Q_T > 0$ and $Q_T < 0$ need to be considered separately, as in one case, hot buoyant air is expelled through the top opening, while in the other, cold ambient air is brought in through it.

For the plume, Q_P and g'_P are given by (1.4) and (1.3), respectively (using the MTT model). For the top opening, following HL05, we have

$$Q_T = \begin{cases} A^* \left[g'(H - h) - \frac{\Delta p}{\rho} \right]^{1/2}, & \text{for } Q_T > 0, \\ -A^* \left[\frac{\Delta p}{\rho} - g'(H - h) \right]^{1/2}, & \text{for } Q_T < 0, \end{cases} \quad (2.3)$$

where Δp is the difference in external pressure between the two openings due to the presence of wind. Note that $\Delta p > 0$ corresponds to ‘opposing wind’.

If the interface descends right to the base of the box, the evolution equation is somewhat different. In this fully contaminated regime, the interface remains at the bottom of the box, i.e.

$$h = \frac{dh}{dt} = 0. \quad (2.4)$$

Conservation of buoyancy yields

$$A_C H \frac{dg'}{dt} = Q_P g'_P - Q_B g', \quad Q_B = A^* \left[\frac{\Delta p}{\rho} - g' H \right]^{1/2}, \quad (2.5)$$

where Q_B is the volume flux out of the bottom opening. Equations (2.1)–(2.5) together form a set of governing equations which can model flow in each of the regimes shown in figure 2.

2.3. Non-dimensionalisation – multiple time scales

We define the dimensionless variables

$$\hat{h} = h/H \quad \text{and} \quad \hat{g} = g'/g'_H, \quad (2.6)$$

where g'_H is given by (1.6). Non-dimensionalising time is more subtle, however, as there are three independent time scales present in this system. Each can be constructed by dividing the volume of the box by a particular scaling for flow rate. The buoyancy-driven ‘filling box’ time scale, T_{Bf} , the buoyancy-driven ‘draining box’ time scale T_{Bd} and the wind-driven time scale T_W are, respectively,

$$T_{Bf} = \frac{A_C H}{C B_0^{1/3} H^{5/3}} = \frac{A_C}{C B_0^{1/3} H^{2/3}}, \quad (2.7a)$$

$$T_{Bd} = \frac{A_C H}{A^* (g'_H H)^{1/2}} = \frac{A_C C^{1/2} H^{4/3}}{A^* B_0^{1/3}}, \quad (2.7b)$$

$$T_W = \frac{A_C H}{A^* (\Delta p / \rho)^{1/2}}. \quad (2.7c)$$

As we intend to examine the effect of varying opening size A^* and wind strength Δp while keeping B_0 fixed, we choose to non-dimensionalise time with respect to the filling box time scale, defining $\hat{t} = t/T_{Bf}$. In doing so, we recover the dimensionless parameters V and W as defined by (1.1) and (1.5), and now note that each is equal

to the square of a ratio of time scales:

$$\text{vent parameter: } V = \left(\frac{T_{Bf}}{T_{Bd}} \right)^2 = \frac{A_*^2}{C^3 H^4}, \quad (2.8)$$

$$\text{wind parameter: } W = \left(\frac{T_{Bd}}{T_w} \right)^2 = \frac{\Delta p / \rho}{g'_H H}. \quad (2.9)$$

(As already noted, V and W are related to μ as defined by KH04 and the Froude number Fr as defined by HL05, since $V = \mu^{-2}$ and $W = Fr^2$.)

2.4. Dimensionless governing equations

Our dimensionless governing

$$-\frac{d\hat{h}}{d\hat{t}} = \begin{cases} \hat{h}^{5/3} - |VP|^{1/2} & \text{(A)} \\ \hat{h}^{5/3} + |VP|^{1/2} & \text{(B)} \\ 0 & \text{(C)} \end{cases} \quad (2.10a)$$

$$\frac{d}{d\hat{t}}[\hat{g}(1 - \hat{h})] = \begin{cases} 1 - |VP|^{1/2} \hat{g} & \text{(A)} \\ 1 & \text{(B)} \\ 1 - |VP|^{1/2} \hat{g} & \text{(C)}, \end{cases} \quad (2.10b)$$

where the letters denote the relevant regimes shown in figure 2. Using this scaling, the direction of the flow is determined by the sign of the quantity P , where

$$P = \hat{g}(1 - \hat{h}) - W. \quad (2.10c)$$

Here P can be interpreted as a measure of the dimensionless pressure at the top opening, representing the resulting balance between the competing effects of buoyancy and wind. Note that $P > 0$ corresponds to the buoyancy-driven forward flow ($Q_T > 0$) and $P < 0$ to the wind-driven reverse flow ($Q_T < 0$). The three regimes shown in figure 2 can thus be classified in terms of the values of \hat{h} and P . Regime A, with forward flow, occurs when $\hat{h} \geq 0$ and $P \geq 0$. Regime B has reverse flow but continues to have stratification and occurs when $\hat{h} > 0$ and $P < 0$. Finally, regime C has reverse flow with a fully mixed interior and occurs when $\hat{h} = 0$ and $P < 0$. Given a set of initial conditions, the coupled pair of nonlinear ODEs (2.10) for \hat{h} and \hat{g} can be integrated numerically for any values of V and W , for example using a standard MATLAB fourth-order Runge–Kutta routine.

2.5. Equilibrium states

If we seek steady solutions to (2.10), we recover the equilibrium states derived by HL05. From regime A, we obtain the buoyancy-driven steady state which clearly exists for all values of W :

$$\hat{g}_{eq} = \hat{h}_{eq}^{-5/3}, \quad V = \frac{\hat{h}_{eq}^{10/3}}{[\hat{h}_{eq}^{-5/3}(1 - \hat{h}_{eq}) - W]}. \quad (2.11)$$

Solutions in regime C – the wind-driven steady states – correspond to roots of the following cubic:

$$\hat{g}_{eq}^3 - W\hat{g}_{eq}^2 + V^{-1} = 0. \quad (2.12)$$

A pair of real, positive roots – physical solutions – exists if and only if $W^3V > 27/4$, the criterion discussed in §1.

3. Zero-mixing Germeles model

If, on the other hand, we do not make the ‘perfect mixing’ approximation inherent in the KH model, the reduced gravity g' is a non-trivial function of both time and vertical position z within the box, with $g'(z, t) = 0$ for $z < h$ and $g'(z, t) > 0$ for $z > h$. Defining $\hat{z} = z/H$, the set of dimensionless governing equations in (2.10) can then be generalised to give

$$-\frac{d\hat{h}}{d\hat{t}} = \begin{cases} \hat{h}^{5/3} - |VP|^{1/2} & \text{(A)} \\ \hat{h}^{5/3} + |VP|^{1/2} & \text{(B)} \\ 0 & \text{(C)} \end{cases} \quad (3.1a)$$

$$\frac{d}{d\hat{t}} \int_0^1 \hat{g} d\hat{z} = \begin{cases} 1 - |VP|^{1/2} \hat{g}|_{\hat{z}=1} & \text{(A)} \\ 1 & \text{(B)} \\ 1 - |VP|^{1/2} \hat{g}|_{\hat{z}=0}, & \text{(C)} \end{cases} \quad (3.1b)$$

where

$$P = \int_0^1 \hat{g} d\hat{z} - W, \quad (3.1c)$$

the natural generalisation of P as defined in (2.10c) to situations where the reduced gravity is also a function of position in the upper layer. Note that the reduced gravity of fluid expelled from the top opening $\hat{g}|_{\hat{z}=1}$ differs from that of any fluid expelled from the bottom opening $\hat{g}|_{\hat{z}=0}$. (We assume that the openings are infinitesimally thin and therefore that fluid leaves or enters the box at its very top or bottom. The validity of this approximation for our experiments is discussed in §6.) Also, in expressions involving the total buoyancy in the box, $\hat{g}(1 - \hat{h})$ is replaced by $\int_0^1 \hat{g} d\hat{z}$. In particular, the volume flux out of the top opening is given by

$$\hat{Q}_T = \pm V^{1/2} \left| \int_0^1 \hat{g} d\hat{z} - W \right|^{1/2} = \pm |VP|^{1/2} \quad \text{for } P \geq 0. \quad (3.2)$$

Now, a model must be developed for the evolution of $g'(z, t)$. As discussed in §1, we consider the limit of ‘zero mixing’, in which the plume always rises through any hot layer to the very top of the box, depositing fluid at ceiling level. The plume volume flux $Q_P(z, t)$, specific momentum flux $M_P(z, t)$ and specific buoyancy flux $B_P(z, t)$ no longer obey the similarity solution used in (1.4), and therefore must be found by integrating the general plume equations of MTT:

$$\frac{\partial Q_P}{\partial z} = 2\pi^{1/2} \alpha M_P^{1/2}, \quad \frac{\partial M_P}{\partial z} = \frac{B_P Q_P}{M_P}, \quad \frac{\partial B_P}{\partial z} = -Q_P \frac{\partial g'}{\partial z}, \quad (3.3)$$

where α is (for simplicity) the ‘top-hat’ entrainment constant. Local continuity then dictates that the stratification within the box is advected vertically according to

$$\frac{\partial g'}{\partial t} + w \frac{\partial g'}{\partial z} = 0, \quad (3.4)$$

where $w(z, t)$ is the return velocity in the bulk flow required to ensure that the volume flux through any horizontal plane matches that through the top opening, under the assumption that the cross-sectional area of the plume is negligible compared with the cross-sectional area of the room A_C :

$$wA_C + Q_P = Q_T. \quad (3.5)$$

We now define dimensionless plume fluxes \hat{Q}_P , \hat{M}_P and \hat{B}_P :

$$\hat{Q}_P = \frac{Q_P}{Q_H} = \frac{Q_P}{C B_0^{1/3} H^{5/3}}, \quad \hat{M}_P = \frac{M_P}{M_H} = \frac{M_P}{\frac{1}{\pi} \left(\frac{9\alpha}{10} \pi B_0 H^2 \right)^{2/3}}, \quad \hat{B}_P = \frac{B_P}{B_0}. \quad (3.6)$$

Non-dimensionalising all other variables in the manner set out in §2.3 and substituting (3.5) into (3.4), we obtain the following equations:

$$\frac{\partial \hat{Q}_P}{\partial \hat{z}} = \frac{5}{3} \hat{M}_P^{1/2}, \quad \frac{\partial \hat{M}_P}{\partial \hat{z}} = \frac{4 \hat{B}_P \hat{Q}_P}{3 \hat{M}_P}, \quad \frac{\partial \hat{B}_P}{\partial \hat{z}} = -\hat{Q}_P \frac{\partial \hat{g}}{\partial \hat{z}}, \quad (3.7)$$

$$\frac{\partial \hat{g}}{\partial \hat{t}} + (\hat{Q}_T - \hat{Q}_P) \frac{\partial \hat{g}}{\partial \hat{z}} = 0. \quad (3.8)$$

This system is closed by (3.2) for \hat{Q}_T and is subject to the boundary conditions $\hat{Q}_P(0, \hat{t}) = \hat{M}_P(0, \hat{t}) = 0$ and $\hat{B}_P(0, \hat{t}) = 1$, for all \hat{t} .

We can integrate this for general initial conditions $\hat{g}(\hat{z}, 0)$ using a numerical scheme of the type developed by Germeles (1975). We discretise the density field in the bulk flow (outside the plume) into a finite number of layers and make the assumption that the time scale for the evolution of this field is much slower than the transit time for fluid rising through the plume (valid provided that the box is not extremely tall and narrow). At the start of each time step, we therefore integrate (3.7) over the whole height of the box, assuming a static (discretised) bulk density field. After evaluating $\hat{Q}_P(\hat{z}, \hat{t})$ and $\hat{Q}_T(\hat{t})$, the existing layer boundaries are then adjusted upward or downward according to (3.8), effectively acting as characteristics for the density field.

At the end of each time step, the top and bottom of the box can have new layers of fluid introduced or existing ones flushed out, depending on the direction and structure of the flow. Note that the presence of opposing wind necessitates a modification of the standard Germeles algorithm to account for the three different flow regimes illustrated in figure 2 and discussed at the beginning of §2. In regime A (forward flow with $\hat{Q}_T \geq 0$, $\hat{h} \geq 0$), a cool ambient fluid with $\hat{g} = 0$ is drawn in through the bottom opening. If a transition from regime (C) has just occurred, a new layer with $\hat{g} = 0$ is introduced. Otherwise, the existing bottom layer is enlarged. At the top, if $\hat{Q}_T > \hat{Q}_P$, then one or more existing layers are swept out of the box. If $\hat{Q}_T < \hat{Q}_P$, a new layer of plume fluid is deposited with $\hat{g} = \hat{g}_P(1, \hat{t})$. In regime B (reverse flow with a layer of the ambient fluid in the lower part of the box, $\hat{Q}_T < 0$, $\hat{h} > 0$), this bottom layer decreases in size. At the top of the box, a new layer is introduced with

$$\hat{g} = \frac{\hat{Q}_P(1, \hat{t}) \hat{g}_P(1, \hat{t})}{\hat{Q}_P(1, \hat{t}) - \hat{Q}_T}. \quad (3.9)$$

In regime C (reverse flow in which the whole box is contaminated with plume fluid, $\hat{Q}_T < 0$, $\hat{h} = 0$), one or more existing layers with $\hat{g} > 0$ are swept out of the bottom opening. A new layer is introduced at the top of the box in exactly the same manner as in regime B.

We assume here that any ambient fluid entering the box through the top opening mixes only with the fluid deposited by the plume in a very thin top layer. It is clear that this represents the opposite extreme to the ‘perfect mixing’ approach adopted in our KH model. The merits of each approach are discussed in more detail in §6. Note that if \hat{Q}_T suddenly becomes more negative, (3.9) results in a less buoyant (lower \hat{g})

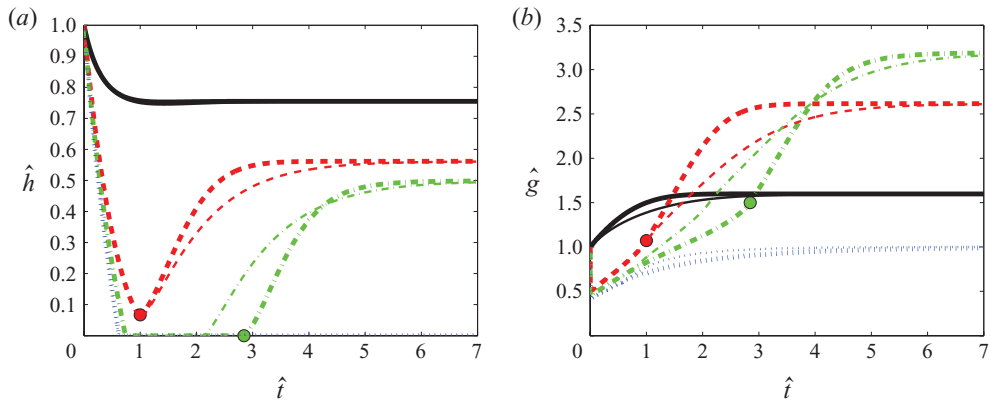


FIGURE 3. (Colour online) Theoretical predictions for (a) $\hat{h}(\hat{t})$ and (b) $\hat{g}(\hat{t})$ when $V = 1$ and $W = 0$ (plotted by black solid lines), $W = 1$ (red dashed lines), $W = 1.5$ (green dot-dashed lines) or $W = 2$ (blue dotted lines). Thick lines correspond to the results of the perfect-mixing KH model and thin lines correspond to the results of the zero-mixing Germeles model. Circular markers indicate points at which transition from reverse to forward flow occurs.

layer being introduced above more buoyant layers. As the Germeles approach does not incorporate any mechanism for mixing between layers, it may then predict that this statically unstable stratification persists, clearly a physically unrealistic outcome. This scenario, which could occur in the event of a sudden increase in W , is also touched upon in §6.

4. Theoretical results for box filling with steady opposing wind

Having developed two time-dependent emptying filling box models that incorporate wind, we now apply these to the simplest possible transient problem – the filling of a box initially full of ambient fluid in the presence of a constant opposing wind (constant W).

4.1. KH model results for $V = 1$

Following Baines & Turner (1969), we begin with an infinitesimally thin hot layer at the top of the box,

$$\hat{t} = 0 : \quad \hat{h} = 1, \quad \hat{g} = 1, \quad (4.1)$$

and use the perfect-mixing KH model to determine the time evolution of \hat{h} and \hat{g} for different values of V and W . The case $W = 0$ is exactly equivalent to the problem examined by KH04 – box filling with no wind – and we verify that our numerical code reproduces their results for all values of V considered. In addition, we verify (for $10^{-10} < V < 10^2$ and $10^{-1} < W < 10^2$) that at large times, both \hat{h} and \hat{g} asymptotically approach values corresponding to one of the steady states given in §2.5. Noting once again that any overshoots in the zero wind case are extremely small, we now consider finite W , restricting ourselves at first to a fixed opening size, with $V = 1$. Figure 3 shows the evolution of \hat{h} and \hat{g} for different wind strengths.

The size of the transient overshoot in \hat{h} increases dramatically with W , i.e. with non-zero wind. For $W = 1$, the overshoot is almost half the box height, far greater than the maximum possible in the absence of wind. It is clear from the plots of \hat{g} that early-time behaviour for finite W is qualitatively different from that for $W = 0$. In the former case, $P = \hat{g}(1 - \hat{h}) - W < 0$ for early times. Recalling §2.4, this means that the

system begins in regime B – reverse flow with a stratification maintained. For small wind strengths (see $W = 1$ in figure 3), transition to regime C does not occur (i.e. the interface does not reach the floor) and we can use (2.10b) to show that the system remains in reverse flow up until

$$\hat{t} = W, \quad \text{equivalent to} \quad t = \frac{\Delta p / \rho}{B / A_c}. \quad (4.2)$$

Thereafter, flow is in the forward direction (regime A) and \hat{h} and \hat{g} approach their buoyancy-driven steady-state values. In figure 3, the point when flow changes direction for the case $W = 1$ is clearly seen to be at $\hat{t} = 1$. It is important to appreciate that the time when flow changes direction does not correspond to the time of maximum overshoot and in general slightly precedes it. However, the lag between these two times is too small to be discernible for the example shown in figure 3.

For larger wind strengths (see $W = 1.5$ in figure 3), the interface descends right to the bottom of the box and a fully contaminated transient is observed. After some time, the box heats up enough for the system to switch to buoyancy-driven flow; a new interface starts to rise from the floor of the box and eventually the buoyancy-driven steady state is reached. For yet stronger wind, however (see $W = 2$ in figure 3), the system appears to remain fully contaminated and approaches the wind-driven steady state predicted by HL05.

Note that when $W > 0$, $d\hat{g}/d\hat{t}$ is singular at $h = 1$: an infinitely thin upper layer can change temperature infinitely quickly, immediately adjusting to a temperature at which the cooling due to incoming ambient air balances the heating from the plume. For our numerical simulations, initial layer depth is in fact of order $\epsilon \approx 10^{-6}$, and it can be shown that the initial adjustment of \hat{g} therefore also occurs on a time scale of order ϵ . This effect is clearly visible in figure 3(b). The behaviour of the model in this region need not be considered too carefully, however. In reality, extremely thin upper layers will immediately be broken down by opposing wind and thus layers of finite depth must be considered when quantitative comparison is made with experiments (discussed in further detail in §6).

4.2. Zero-mixing Germeles model results for $V = 1$

We apply the zero-mixing Germeles model to the same problem, starting with an empty box into which a thin top layer with $\hat{g} = 1$ is added in the first time step. Given that $\hat{Q}_p(1, 0) = 1$ and $\hat{Q}_T(0) = 0$, (3.8) dictates that the thickness of this first layer is equal to the time-step size $\Delta\hat{t}$. Note that $V = 0$ (no openings) represents the classical filling box problem of Baines & Turner (1969), investigated using a Germeles approach by Worster & Huppert (1983). We therefore verify that in this case our numerical code reproduces the latter's results for the evolution of the plume fluxes and reduced gravity profile $\hat{g}(\hat{z}, \hat{t})$. We also examine the same cases as in §4.1, varying W for a fixed opening size, $V = 1$, so that we can compare the results with those of the KH model in figure 3. Note that in figure 3(b), curves plotted for the zero-mixing Germeles model are of the mean value of \hat{g} in the upper layer, defined as

$$\bar{\hat{g}} = \frac{\int_{\hat{h}}^1 \hat{g} \, d\hat{z}}{(1 - \hat{h})}. \quad (4.3)$$

It is clear that the predictions of the two models are qualitatively very similar: an increased overshoot for small W , a fully contaminated transient for slightly larger W , and convergence to the wind-driven steady state for yet larger W . However, there

are some quantitative differences between the two sets of results. These are best understood by considering each of the three flow regimes in figure 2 separately.

In regime B, no hot fluid exits the box through either opening. Volume flux through the openings is determined by the total buoyancy in the box and not its distribution and thus the predictions of the zero-mixing Germeles model for \hat{h} and mean \hat{g} are identical to the predictions of the perfect-mixing KH model. This is evident in figure 2 for $W = 1$ (up until transition to regime A), $W = 1.5$ and $W = 2$ (up until transition to regime C). However, in regime A, under the zero-mixing Germeles model, warmer than average fluid is lost out of the top opening due to the stratification within the upper layer. Consequently, overall warming is slower than under the KH model. We see in figure 2 for $W = 0$, $W = 1$ (after the change in flow direction) and $W = 1.5$ (during final stages) that both \hat{h} and \hat{g} adjust more slowly to their equilibrium values under the zero-mixing Germeles model than they do under the perfect-mixing KH model. Finally, in regime C, under the zero-mixing Germeles model, cooler than average fluid is lost out of the bottom opening, and overall warming is thus faster than under the perfect-mixing KH model. We observe for $W = 2$ in figure 3(b) that \hat{g} reaches its equilibrium value more quickly and for $W = 1.5$ in figure 3(a) that transition to forward flow occurs earlier. We discuss the circumstances under which each model is likely to be more accurate in § 6, when comparison is made with experimental data for the evolution of both \hat{h} and the structure of $\hat{g}(\hat{z}, \hat{t})$.

4.3. Varying V and W : exploring the whole parameter space

In order to gain a clearer picture of the conditions under which each of the various observed transient behaviours occurs, we explore the wider parameter space $10^{-1} < W < 10^2$, $10^{-10} < V < 10^2$. Making the type of assumptions about building size and heat load discussed in Appendix B of Hunt & Linden (2005), this could correspond to wind speeds between 5 cm s^{-1} and 2 m s^{-1} , and effective opening areas approximately satisfying $10^{-3} < \sqrt{A^*}/H < 1$. For each point on a 60×120 logarithmic grid spanning this space, we perform a numerical simulation for the time evolution of \hat{h} and \hat{g} . In particular, we consider three important lengths: the minimum interface height reached \hat{h}_{min} , the final equilibrium interface height \hat{h}_{eq} and the difference between the two – the overshoot $\hat{h}_{eq} - \hat{h}_{min}$. These data are presented in figure 4. Note that as the two models produce qualitatively similar results, we here employ the less computationally costly KH model.

It is immediately clear from figure 4(c) that very large overshoots can occur in the presence of wind. In fact, it is theoretically possible to have an arbitrarily large overshoot – up to the full box height. In practice, our model ceases to be relevant once opening area approaches the order of the square of the box height ($V \approx 10^2$), but figure 4(c) illustrates that overshoots representing a significant fraction of the box height can occur at physically relevant values of V . Figure 4(a) shows that for any given value of V , as W is increased, the minimum height attained decreases smoothly down to zero. The black area corresponds to cases where the minimum height is exactly zero, i.e. those where a fully contaminated state (transient or permanent) is attained. The edge of this region is marked by a white curve, which is also reproduced in figures 4(b) and 4(c). Figure 4(b) shows that the final equilibrium height does not decrease smoothly to zero as V and W increase and is in fact discontinuous along a straight line in log–log space. To the left of this line (marked in white on all three panels of figure 4), the box always finishes in the buoyancy-driven steady state (associated with displacement ventilation, in regime A), to the right of it (the black region) the box finishes in the wind-driven steady state (HL05's stable solution, associated with




FIGURE 4. Colour maps of (a) the minimum interface height \hat{h}_{min} reached during evolution to equilibrium, (b) the final equilibrium interface height \hat{h}_{eq} and (c) overshoot size $\hat{h}_{eq} - \hat{h}_{min}$, in each case for $10^{-1} < W < 10^2$ and $10^{-10} < V < 10^2$. Regions coloured black correspond (a) to the interface reaching the floor and (b) and (c) to the final state being well-mixed and hence wind-driven. White lines demarcating these regions are drawn on all three plots. Circular white markers indicate three of the four sets of parameters investigated experimentally in § 6, while the fourth has $W = 0$ and thus does not appear on these logarithmic axes.

mixing ventilation, in regime C). The region between the straight and curved white lines represents the set of cases where a fully contaminated transient is observed (cf. $V = 1$, $W = 1.5$ in figure 3). The straight line corresponds to the relationship $W^3 V = 27/4$, precisely the criterion for the existence of wind-driven steady states (see (1.8)). This implies that for the particular initial conditions we consider, there is no ambiguity with regard to which of HL05's multiple steady states is adopted – whenever the (stable) wind-driven (regime C) steady state exists, the system locks on to it.

4.4. Phase diagrams

For constant V and W , the nonlinear dynamical system represented by the equations laid out in § 2.4 is autonomous – the perfect-mixing KH model's evolution equations for \hat{h} and \hat{g} have no explicit dependence on \hat{t} . It is therefore possible to generate phase diagrams to help understand the behaviour of the system and explain why HL05's wind-driven steady state is an attractor for the initial conditions considered. Figure 5 comprises four such diagrams, for the cases $V = 1$, $W = 0, 1, 1.5, 2$. At each point in the \hat{g} – \hat{h} space, a vector is plotted indicating the relative magnitudes and signs of $d\hat{g}/d\hat{t}$ and $d\hat{h}/d\hat{t}$ (which are computed directly from (2.10)). For simplicity, all such vectors are normalised so that the arrows depicted are of equal length. The dashed solid line shown on the three plots where $W > 0$ is the $P = 0$ curve, which separates regions of forward and reverse flow. The grey shaded areas hence correspond to regime A, the unshaded areas to regime B and the horizontal lines $\hat{h} = 0$, $P < 0$ to regime C (see figure 2). The solid red lines are trajectories corresponding to the initial conditions $\hat{h} = \hat{g} = 1$, plotted using the results shown earlier in figure 3. These curves are tangent to the corresponding vector fields at all times (as all system trajectories must be) and each one terminates at one of the fixed points corresponding to HL05's stable equilibrium states (given by (2.11) and (2.12) and shown by solid red circles).

The dotted black line on each diagram in figure 5 corresponds to the curve $\hat{g} = \hat{h}^{-5/3}$, along which the temperature of fluid in the hot layer is equal to the temperature of fluid injected into it by the plume. For any W (and indeed any V), HL05's buoyancy-driven equilibrium clearly must lie on the portion of this curve bordering the grey shaded region. The area above and to the right of this curve corresponds to

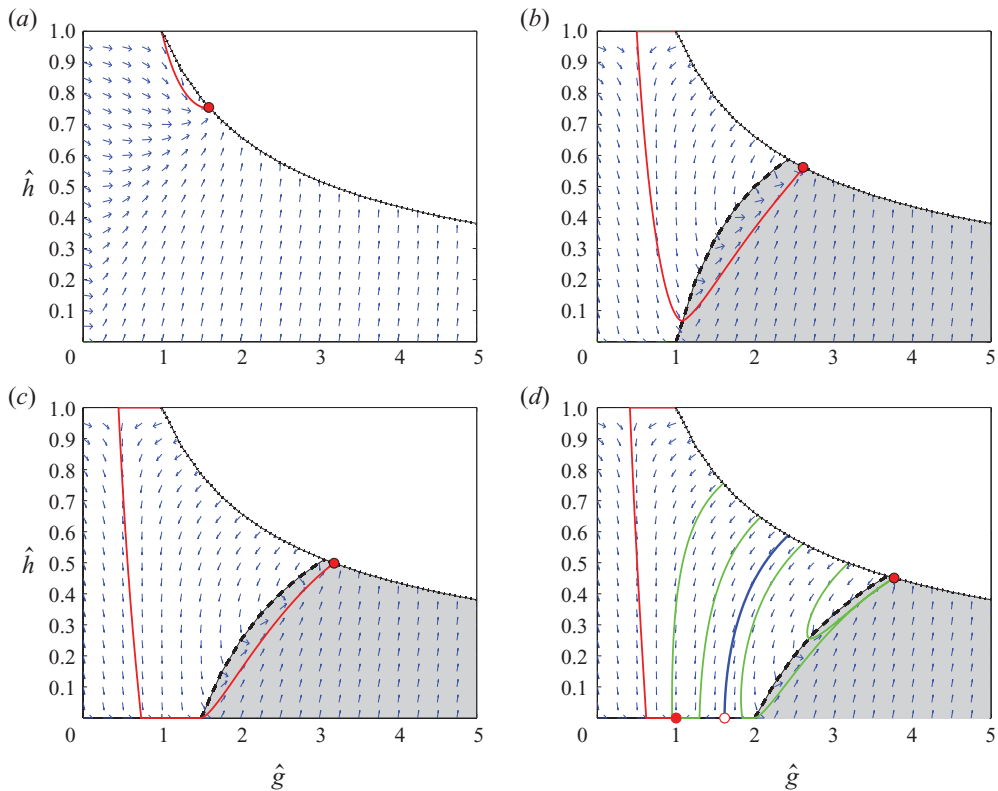


FIGURE 5. (Colour online) Phase diagrams for $V = 1$ and (a) $W = 0$, (b) $W = 1$, (c) $W = 1.5$ and (d) $W = 2$. Our model only applies below and to the left of the thin black line $\hat{g} = \hat{h}^{-5/3}$. Dashed black lines corresponding to $P = 0$ separate (shaded) regions of forward flow from (unshaded) regions of reverse flow. Solid red circles indicate stable equilibria. Solid red lines are trajectories starting from $\hat{h} = \hat{g} = 1$. The solid green lines on (d) are trajectories for four different initial conditions, equivalent to the equilibrium states attained for $W = 0, 0.5, 1$ and 1.5 . The hollow red circle on (d) indicates an unstable equilibrium, while the solid blue line leading into it represents the separatrix curve.

cases where the fluid being injected by the plume is colder and denser than that in the upper layer. In such cases, the assumption that the upper layer remains well-mixed is clearly questionable; it is more probable that the plume will intrude at the interface. These portions of the phase diagrams shown are not relevant to this study and thus are ignored. (An interesting extension to our work might be to use the approach of Bower *et al.* (2008) to model this ‘intruding regime’ with opposing wind.)

Figure 5(a) shows that in the absence of wind, the buoyancy-driven steady state is a global attractor. In their investigation of an overshoot criterion, KH04 linearise their system about this fixed point and establish numerically whether or not the corresponding Jacobian has complex eigenvalues. What they describe as ‘underdamping’ (when overshoot does occur) can now be interpreted in phase space as the existence of a stable focus, whereas ‘overdamping’ corresponds to a stable node. A similar analysis demonstrates that in all of the cases illustrated in figure 5, this steady state is indeed a stable focus.

The trajectories in figures 5(b) and 5(c) illustrate once again that reverse flow and a fully contaminated transient can occur before the system reaches the

buoyancy-driven steady state. The phase diagrams clearly indicate that the buoyancy-driven steady state is still a global attractor. However, in the large wind case (figure 5*d*), wind-driven steady states also appear; one is a stable node and the other a saddle point (unstable), marked by an open red circle. The red trajectory shows that the initial state $\hat{h} = \hat{g} = 1$ lies within the basin of attraction of the wind-driven node – the system locks on to this equilibrium, as discussed in §4.3. This remains the case for all V and W such that $W^3V > 27/4$.

The grey lines in figure 5(*d*) represent other possible system trajectories and illustrate that for different initial conditions, it is still possible to reach the buoyancy-driven equilibrium state, even when a wind-driven state exists. The solid blue line leads into the saddle point and is thus the separatrix curve dividing the basins of attraction of each stable steady state. We are able to locate this curve using a numerical algorithm for all relevant values of V and W . As W increases for fixed V , the buoyancy-driven equilibrium state moves down and to the right along the curve $\hat{g} = \hat{h}^{-5/3}$, while the stable wind-driven state moves to the left along $\hat{h} = 0$. The separatrix also migrates to the right, thereby reducing the basin of attraction of the buoyancy-driven equilibrium state and increasing that of the wind-driven one.

These results can be used to explain HL05's observation that the system has a hysteretic response to successive increments/decrements in wind strength. The initial conditions for the four green lines in figure 5(*d*) in fact correspond to the buoyancy-driven equilibrium states for $V = 1$, $W = 0, 0.5, 1$ and 1.5 (see figure 5*a–c*). We see therefore that while a sudden jump in wind strength from $W = 0$ or 0.5 to $W = 2$ would result in transition to a wind-driven equilibrium state, a jump from $W = 1$ or 1.5 to $W = 2$ would see buoyancy-driven equilibrium maintained. If wind is increased in sufficiently small increments, it is theoretically possible to remain in the buoyancy-driven state indefinitely. However, as the basin of attraction of this state becomes very small, it becomes unstable to the finite fluctuations seen in any experimental flow and transition to the wind-driven state is thus observed experimentally. Conversely, if W is reduced in sufficiently small decrements, the wind-driven state can be maintained in theory up to the critical point below which it cannot exist ($W^3V = 27/4$), and in reality up to a point close to this when its basin of attraction is too small to contain typical experimental fluctuations.

5. Experimental methods

Experiments are performed in a cuboid Perspex box (internally 29.5 cm long, 15 cm wide and 25 cm high) with rows of five circular holes (each of diameter 2 cm) cut into both the windward and leeward faces at high and low levels (20 holes in total). The total opening area is controlled by blocking some of these holes with rubber bungs. The effect of wind is recreated by positioning the box within the test section of a recirculating flume tank (265 cm long, 30 cm wide and 57 cm deep) filled with fresh water (see figure 6). By controlling the mean speed of flow in the tank, we vary the dynamic pressure difference Δp between the windward and leeward sides of the box, measured directly using an oil/water manometer. At all flow speeds, Δp is observed to fluctuate by up to 10 % of its mean value (see also HL05).

The buoyancy source is produced by the injection of a salt solution through a circular nozzle (0.5 cm in diameter) positioned in the centre of the top of the box. The nozzle used was specifically designed by Dr Paul Cooper to produce turbulent plumes even at small flow rates by forcing fluid through a chamber containing a sharp expansion and a fine gauze mesh (see Hunt & Linden 2001 and Cooper &




FIGURE 6. A schematic diagram of our experimental set-up, which is based very closely on that of HL05. Flow generated by the flume is turned by a reflector and passed through a honeycomb mesh, ensuring a roughly uniform horizontal velocity in the test section where the box is placed. A peristaltic pump supplies fluid to the plume source and an oil/water manometer measures the dynamic pressure drop across the box.

Linden 1996 for more details). Note that the Boussinesq approximation is valid, and so the resulting negatively buoyant saline plume can be used to model a positively buoyant thermal plume without altering the dynamics of the system. The bottom of the box now represents the ceiling of a room, and the top face represents the floor. The distance between the exit point of the nozzle and the bottom of the box is $H_m = 23.2$ cm and the distance between the centre of the lower holes and the bottom of the box is $d_c = 2$ cm. As discussed in detail by Hunt & Linden (2001), various adjustments can be made to the vertical distances in the box to account for the finite vertical extent of the openings. For simplicity, we choose to ignore these small corrections.

The apparatus as described thus far was used by Hunt & Linden (2001) and HL05. Unlike the latter, however, we use a peristaltic pump to supply the salt solution to the nozzle at a single flow rate of $1.33 \text{ cm}^3 \text{ s}^{-1}$, which is measured using an in-line flow meter and observed to fluctuate by no more than a few per cent. We also differ from previous studies by using sodium carbonate instead of sodium chloride as our solute, as chloride ions were found to cause complications for the flow visualisation technique we choose to employ. Sodium carbonate has a molecular diffusivity in water very similar to sodium chloride ($\sim 10^{-9} \text{ m}^2 \text{ s}^{-1}$) and so its use has no effect on the fluid dynamics of the experiments undertaken. Note that all mentions made of ‘salt’ and ‘salinity’ with respect to our experiments refer to sodium carbonate and concentration of the sodium carbonate solution. The density of the solution supplied to the plume and that of the fresh water in the tank is measured using an Anton Paar DMA5000 density meter. The reduced gravity of the source fluid used is 1.66 m s^{-2} .

To ensure opposing wind (noting that saline plumes are negatively buoyant), all the holes of the box are permanently blocked by bungs except for two of the lower holes on the windward side and two of the upper ones on the leeward side. In each case, the holes left unblocked are those furthest from the centre of the box, so that any rapid inflow or outflow is less likely to directly disturb the plume. By observing the rate at which dyed salt solution of a known concentration drains out of the box when these four bungs are removed (see Linden *et al.* 1990), we determine that the effective opening area $A^* = 4.6 \text{ cm}^2$ for this configuration. This estimate is consistent

with that obtained by using the conventional values $c_B = c_T = 0.6$ for the discharge coefficients in (1.2) (see HL05), and implies that the vent parameter $V = 0.026$.

Flows in the box are visualised by backlighting the apparatus and adding dye to the saline source fluid. The molecular diffusivities of the salt and dye are both a few orders of magnitude smaller than the kinematic viscosity $\nu \approx 10^{-6} \text{ m}^2 \text{ s}^{-1}$ and so dye concentration can be used as a surrogate for salt concentration. We therefore use measurements of light absorption to determine the cross-tank (line-of-sight integrated) density field in the box at any instant in time, employing techniques outlined by Cenedese & Dalziel (1998). Following Dalziel *et al.* (2008), we use electroluminescent LightTape (Electro-LuminX Lighting Corporation, Chester, VA, USA) together with the red food colouring ‘Fiesta Red’ (Allura Red AC, E129). A $54 \text{ cm} \times 54 \text{ cm}$ panel of tape attached to the rear side of the flume provides uniform backlighting of approximately 400 cd m^{-2} . Temporal fluctuations in luminosity are less than 1.5% and we take care to minimise reflections and the infiltration of light from any other sources. The small amount of heat generated by the tape does not cause any significant temperature increase inside the flume due to the large volume of fluid involved. We choose to use ‘Fiesta Red’ dye because the peak in its absorption spectrum corresponds quite closely with the near-monochromatic cyan light emitted by the tape. Very high absorption is thus achieved at high dye concentrations. We find that this dye forms a dark red precipitate when added to the concentrated sodium chloride solution. This chemical reaction may complicate the optical properties of the dye in dilute solutions too. It is for this reason that we choose to use sodium carbonate solution, for which we observed no such precipitation.

We measure the transmitted light intensity using a Jai CV-M4+CL 1.3 megapixel video camera fitted with a 135 mm f3.5 Nikon zoom lens, recording a 1380×1030 pixel image in 8 bit monochrome at 2 frames s^{-1} . The intensity response of the camera is linear with a small black offset ($b = 0.035$). To minimise the parallax error, we set the lens to maximum zoom and we position the camera approximately 3 m in front of the flume. The image recorded lies within the area illuminated by the tape and a 980×830 pixel portion of it is taken up by the interior of the box. This corresponds to approximately $0.3 \text{ mm pixel}^{-1}$ spatial resolution of our measurements. Other portions of it include a region above the box which is never contaminated by dye and a region where black masking tape on the front of the flume always prevents any light from reaching the camera. By applying a linear mapping that ensures that the measured intensities of these last two regions remain constant in time, we eliminate the effect of the small fluctuations in the LightTape’s luminosity.

We remove the effect of spatial variations in illumination and that of absorption due to everything other than the dye (the perspex walls of the box as well as the surrounding fluid-filled flume) by dividing any experimental image I by a reference image I_0 of the whole apparatus filled with fresh undyed water, taking account of the black offset of the camera as follows:

$$q = \frac{I - b}{I_0 - b}. \quad (5.1)$$

We assume here that the approximately monochromatic behaviour of the dye when illuminated by the tape permits use of the Lambert–Beer law. We determine the relationship between dye concentration and $\log(q)$ by conducting a series of calibration images *in situ* in the flume. As salt concentration may affect absorption independent of dye concentration, we ensure that the ratio between the two in the source fluid used for our experiments is maintained throughout this calibration

procedure. All image processing described in this section and the next is carried out using DigiFlow (Dalziel Research Partners, Histon, Cambridgeshire, UK).

While the theoretical results presented in §4 assume an ideal plume source, our experiments clearly involve a source with finite volume and momentum flux. Before making any comparison between the two, we apply an asymptotic virtual origin correction (see Caulfield 1991; Caulfield & Woods 1995; Hunt & Kaye 2001): a source of buoyancy flux alone positioned at a height z_v above the experimental plume nozzle would behave in the same way as the observed plume at sufficiently large distances from the source. In this case, we determine that $z_v = 1.2$ cm. Including this adjustment, we find $H = H_m + z_v = 24.4$ cm. It is this value that we use when calculating the scalings g'_H , Q_H , M_H and T_{Bf} . In particular, we find that the filling box time scale $T_{Bf} = 61$ s. (Note that the zero-mixing Germeles model can be adjusted to deal with a non-ideal source without recourse to a virtual origin correction, requiring only the new boundary conditions $\hat{Q}_P(0, \hat{t}) = 0.0076$, $\hat{M}_P(0, \hat{t}) = 0.0023$ for (3.7).) The ‘dead’ layer – the region of the box below the lower holes – is very small in our experiments and thus we choose not to include any correction for this effect.

With the box and flume initially filled with fresh water, we begin experiments by supplying saline fluid to the plume source. Approximately 30 s later, once a plume has descended to the bottom of the box and (after some initial spreading and ‘sloshing’) deposited a 3–4 cm thick layer of fluid, we switch on the recirculation mechanism in the flume. We record the subsequent flow in the box for a period of 20–30 min. This time interval is long enough for an equilibrium state to be reached while not long enough for significant contamination of the ambient by recirculated saline fluid.

6. Experimental results and comparison

We present experimental results for four different values of W , which we refer to as ‘zero wind’ ($W = 0$), ‘small wind’ ($W = 2.29$), ‘medium wind’ ($W = 3.98$) and ‘large wind’ ($W = 7.46$). In all cases, $V = 0.026$ and $T_{Bf} = 61$ s, as mentioned in §5. Note that all experimental data and images presented in this paper are vertically inverted to ease comparison with both theoretical results and real building ventilation flows. As saline fluid is thus observed to rise towards the top of the box, we refer to it as ‘hot’ relative to the ‘cold’ fresh fluid. In addition, because the configuration of two pairs of holes used in our experiments behaves equivalently to a combination of one high and one low opening with the same total effective opening area A^* , we refer to one pair of holes collectively as the top opening and the other pair as the bottom opening. This again eases comparison with our theoretical findings.

In the zero wind case, we observe behaviour of the type reported in KH04. After an initial period in which the saline plume spreads out and splashes back off the sides of the box, a sharp horizontal interface is formed between the layer of the plume fluid and the fresh water in the box. Small wave-like disturbances of the interface are observed as it evolves and remain after it reaches its equilibrium level. Little mixing is seen at the interface and thus the only interchange of fluid between the two layers is via the plume. Ambient fluid is drawn in through the bottom opening and hot fluid leaves through the top opening.

In the other three cases, as soon the flume is switched on, wind blows cold fluid in through the top opening. For small and medium wind, this descends and mixes into the hot layer. For large wind, the momentum of the inflow is seen to have a considerable effect on the flow structure within the hot layer, but it does not cause the breakdown of this layer or overturning of the box. Our modified ‘emptying filling’

theoretical modelling approach is therefore appropriate for all the wind strengths considered.

We note from figure 5 that the small, medium and large wind cases lie in regions of the V – W parameter space for which our theoretical models predict a large overshoot in interface height, a fully contaminated transient and a wind-driven final state respectively. In each case, we find that the qualitative type of behaviour predicted is indeed observed experimentally.

Note that in our experiments, reduced gravity $g'(z, t)$ clearly cannot exceed $g'_p(0, t)$, the fixed (finite) reduced gravity of the source fluid. We therefore deal henceforth with

$$c(z, t) = \frac{g'(z, t)}{g'_p(0, t)}, \quad (6.1)$$

the concentration of ‘heat’ (in reality of source fluid) at a given vertical position in the box. With this scaling, $c = 1$ at the source and $c = 0$ in the initially uncontaminated ambient fluid.

Our optical measurements yield data for the variation in c with respect to the horizontal as well as vertical position within the box. To obtain values of $c(z, t)$ appropriate for the bulk flow, we must take some sort of horizontal average over regions of the box that are away from the openings and the plume. We choose two regions, each the full height of the box and roughly one sixth of its width, one equidistant from the centreline of the plume and the left wall of the box, the other equidistant between the centreline and the right wall. In figure 7, we present experimental measurements of the evolution of $c_L(z, t)$ (for the left region) and $c_R(z, t)$ (for the right region) in the form of a set of colour time-series images. Curves corresponding to the predictions of both the perfect-mixing KH and the zero-mixing Germeles models for these cases are superimposed on the time-series images. The most significant source of error in the parametrisation of the theoretical models is a 10% uncertainty in experimental measurements of pressure drop (see §5) and therefore in W for all but the zero wind case. We hence plot curves in figure 7 using values of W 10% above and below the mean measured values to illustrate the ranges of behaviour consistent with our models.

We see in figure 7 that in contrast to shadowgraph or direct density measurement techniques, our method of flow visualisation not only elucidates the evolution of the interface position \hat{h} but also provides good spatial and temporal resolution of c_L and c_R within the hot layer. In the three cases where a buoyancy-driven final state is reached (zero wind, small wind and medium wind), we observe that the interface between the hot plume and cold ambient regions is blurred over a depth of 5%–10% of the box height. There is slightly more blurring on the right-hand side than on the left. Since ambient air is drawn in from the right, this can be accounted for by the fact that secondary flows within the lower region have higher velocities on this side and there is therefore greater shear of the interface here. We see that the theoretical equilibrium \hat{h} (upon which the choice of model and 10% variations in W have only a small effect) corresponds well with the centre of the blurred interface region in all these three cases, especially in the medium wind case.

Notably, there is also good agreement between theory and experiment for the transient evolution of \hat{h} . With the perfect-mixing KH model, we incorporate the fact that there is a short delay to allow a layer of plume fluid to form before wind is switched on in our experiments (see §5) by modelling an initial interval of box filling without wind. We cannot do the same with the zero-mixing Germeles model, however, as switching on wind after a delay results in the prediction of a statically unstable




FIGURE 7. Colour time-series images of the experimentally determined horizontally averaged concentration c_L (left, from a region to the left of the plume) and c_R (right, from a region to the right of the plume) for experiments with $V = 0.026$, $T_{Bf} = 61$ s and: (a) and (b) zero wind, $W = 0$; (c) and (d) small wind, $W = 2.29$; (e) and (f) medium wind, $W = 3.98$; and (g) and (h) large wind, $W = 7.46$. Solid black lines represent the perfect-mixing KH model results and dashed black lines represent the zero-mixing Germeles model results, plotted for $W \pm 10\%$ of each mean measured value.

stratification (see § 3). Consequently, we see that the perfect-mixing KH model predicts the initial descent of the interface more accurately in all cases. However, for both small and medium winds, it appears that the best overall agreement would be achieved by using the zero-mixing Germeles model with a value of W close to the upper limit of its error bounds. For zero, small and medium winds (particularly the latter two), we observe significant vertical variation in c_L and c_R within the hot layer while the system evolves towards equilibrium. This again suggests that the zero-mixing Germeles model may be more appropriate in these cases. As we would expect, there is little difference in this structure between the left- and right-hand sides. For medium wind, the momentum of the wind-driven inflow seen at early times possibly pushes plume fluid towards the right at the top of the box, resulting in very slightly higher concentration on this side, i.e. $c_R > c_L$.

For large wind, our theoretical models again accurately predict the initial descent of the interface to floor level. Thereafter, the box appears to remain well-mixed across its full depth, except for a band of fluid near the top of the box on the right-hand side in which c_R is significantly higher than elsewhere. Videos of this experiment reveal that wind blows a strong jet of fluid in through the top opening and across the top of the box, sweeping plume fluid towards the top right and ‘trapping’ some of it there. The assumption that the inflow and plume are isolated from each other is inaccurate in this circumstance, though because the inflow is not strong enough to drive a fully overturning circulation in the whole box, we still see some transient vertical stratification.

To investigate further the clear persistence of stratification in the upper layer, in figure 8, we compare experimentally determined concentration profiles $c(z, t)$ (plotted at 20 s time intervals, taking the mean of c_L and c_R in figure 7) with the corresponding predictions of our Germeles model. We note again that while the final states observed are close to two-layer or fully well-mixed profiles, there is significant vertical variation in c before these states are attained. In the experimental data, we see blurred first fronts rather than the sharp interfaces assumed theoretically, reminiscent of the diffusive profiles recently considered by Kaye *et al.* (2009). The final interface positions and hot layer concentrations for low and medium winds appear to be consistent with slightly higher values of W than those used to generate the theoretical profiles. Note that in the discussion of figure 7, we also conclude that values of W 5%–10% larger than the mean experimentally determined values yield the best fits for the evolution of \hat{h} . This suggests that the actual pressure difference between the two openings is slightly larger than that which is measured. A possible explanation for this inaccuracy is the fact that each free end of our manometer tubing must be positioned some 2–3 cm away from an opening in order to avoid disturbing the flow passing through it. The slightly lower-than-predicted final concentration in the bulk of the box for large wind is also consistent with an underestimation of W (given that equilibrium c decreases with increasing W for reverse flow), but could also be caused by the clearly visible region of hot plume fluid ‘trapped’ by the inflow, which our theory does not account for.

For zero wind, we observe the late-time development of an anomalous region of high concentration near the top of the box. As the flume’s recirculation mechanism remains switched off in this case, contaminated fluid leaving the box is not swept around and diluted across the full volume of ambient fluid in the tank, instead mostly remaining in the test section of the tank (see figure 5). In the last stages of our 20 min experiment, the resulting layer of the contaminated ambient fluid becomes deep enough to affect our optical measurements of c over a small part of the box’s depth.

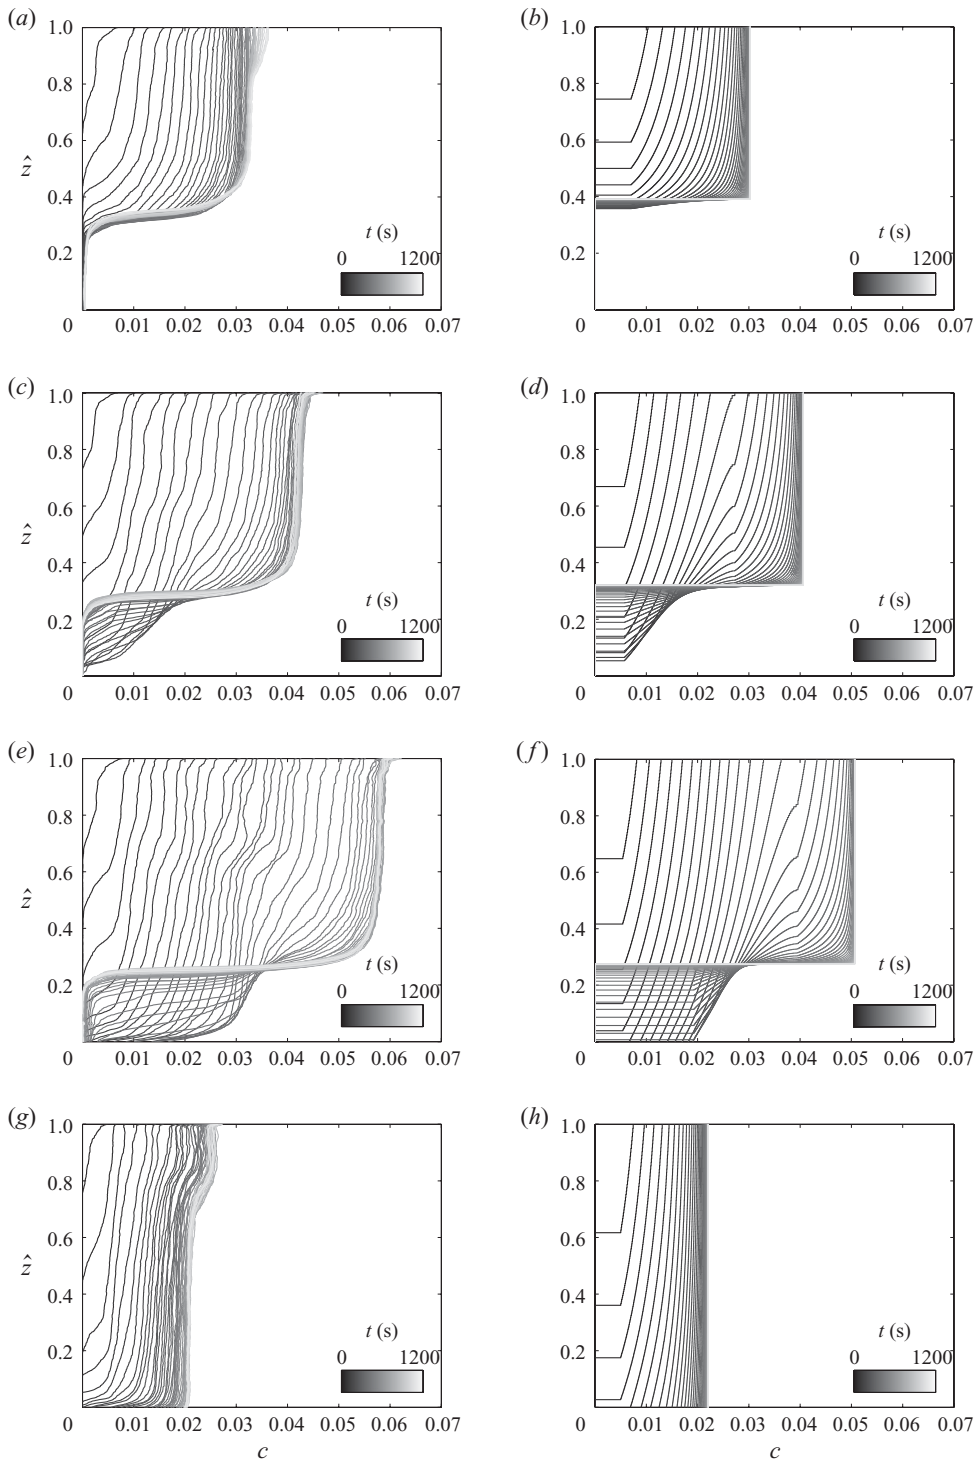


FIGURE 8. Experimentally determined concentration profiles $c(z, t)$ (left) and the theoretical predictions of the zero-mixing Germeles model (right), for $V = 0.026$, $T_{Bf} = 61$ s, and (a) and (b) zero wind, $W = 0$; (c) and (d) small wind, $W = 2.29$; (e) and (f) medium wind, $W = 3.98$; and (g) and (h) large wind, $W = 7.46$.

Overall, however, the transient evolution of $c(z, t)$ is captured well by the zero-mixing Germeles model. One feature of note is the form of profiles corresponding to times shortly after the transition to forward flow in the low and especially medium wind cases. The Germeles model accurately predicts that as the interface starts to rise from its minimum level, an inflection point in the concentration profile (caused by the switch from inflow to outflow at the top opening) propagates down from the top of the box. In our theoretical models, we make the idealised assumption that inflow and outflow occur at the very top and bottom of the box. In our experiments and of course in real buildings, the finite height of the openings mean that, for instance, hot air starts to be expelled from the bottom opening before the interface reaches all the way to the floor. This complication of the transition from regime B to regime C can explain the slight discrepancy between theory and experiment in figure 7 as \hat{h} approaches zero for medium and large winds.

Further complications arise during transitions from either regime B or C to regime A. We have assumed that flow through each opening is unidirectional at all times. However, when the dimensionless pressure P driving the flow (see (2.10c) and (3.1c)) is close to zero, exchange flow can occur through one or both openings. A close observation of our medium wind experiment reveals that exchange flow does occur but only intermittently over a period of less than 15 s. This does not significantly affect the overall dynamics of the system. However, in situations in which the two openings are larger and/or of different sizes, extended periods of exchange flow may indeed alter flow dynamics. For a detailed discussion of the circumstances under which exchange flow occurs in an emptying box, see Hunt & Coffey (2010).

When the flow rate Q_T is less than the finite volume flux of the plume source $Q_P(0, t)$, it is also theoretically possible to have a ‘blocked’ flow regime in which fluid is expelled from both openings (see Woods, Caulfield & Phillips 2003; Flynn & Caulfield 2009). Flynn & Caulfield (2009) explained why this regime cannot be accessed in reverse-to-forward (i.e. P increasing through zero) flow transitions such as those seen in our experiments. If, on the other hand, one wished to examine situations where gradual forward-to-reverse (decreasing P) transitions occur, perhaps as a result of gradually increasing W , then our models would need to be adapted to take transient blocking into account. Flynn & Caulfield (2006b) identified that blocking can cause a transient fully contaminated state in one chamber of a multi-chamber ventilation system forced by a buoyancy source with finite volume flux. We stress here that the fully contaminated state we observe is not due to blocking in this sense – it is wind-driven and (importantly for application to real buildings) can arise in the presence of ideal sources of buoyancy.

7. Conclusions

We have investigated the transient dynamics of emptying filling box flows driven by opposing wind and buoyancy. We have developed two theoretical models for the evolution of dimensionless interface height \hat{h} and reduced gravity \hat{g} in such flows, in which opening area and wind strength are characterised by the dimensionless parameters V and W , respectively. The perfect-mixing KH model, defined by (2.10), assumes instantaneous mixing within regions of the hot plume fluid and consists of a couple pair of nonlinear ODEs, while the zero-mixing Germeles model, defined by (3.2), (3.7) and (3.8), assumes zero vertical mixing and consists of a system of PDEs for \hat{g} as a function of both vertical position and time. We have identified that three distinct flow regimes can occur and have therefore taken account of these in our

models. Regime A has ‘forward’ (buoyancy-driven) displacement ventilation, regime B has reverse (wind-driven) flow with a two-layer stratification and regime C has reverse flow with a fully contaminated interior.

Applying our models to the scenario of an initially empty box, we have demonstrated that transient behaviour in the presence of steady opposing wind is qualitatively different from that seen without wind by KH04. Transient reverse flow can cause very large overshoots in \hat{h} and even a transient fully contaminated state (in which hot fluid is ejected from the bottom opening) before the two-layer equilibrium state of HL05 is attained. We have identified the regions of the V – W parameter space in which we expect these transients to occur. In regions where HL05 showed that both a buoyancy-driven and a wind-driven stable equilibrium can exist, we have established that the latter state is always attained (by an initially empty box). We have used phase space analysis of the KH model to explain why this is the case for these initial conditions but not certain others, determining the basin of attraction of each equilibrium state.

We have confirmed our theoretical predictions by means of an analogous salt bath laboratory study. For this, we have used an experimental set-up very similar to HL05 but employed the novel method of flow visualisation developed by Dalziel *et al.* (2008), in which electroluminescent tape provides the backlighting for optical measurements of dye concentration. As we observed precipitation of the dye (‘Fiesta Red’ food colouring) in highly concentrated sodium chloride solution, we have used sodium carbonate as our salt instead. These techniques have enabled us to achieve significantly better spatial and temporal resolution of the density distribution within the upper layer than previous experimental studies have managed.

By investigating four different values of W for a fixed value of V , we have established the experimental existence of each of the qualitatively different types of flow behaviour predicted theoretically. Furthermore, we have achieved good quantitative agreement between theory and experiment for the evolution of both \hat{h} and $\hat{g}(\hat{z}, \hat{t})$. Despite wind-driven inflow of cold fluid at the top opening causing some mixing within the hot layer, the zero-mixing Germeles model is seen to be more appropriate than the perfect-mixing KH model for the cases considered. We believe that this is not a general result, but depends on the extent to which the plume, once it arrives at the ceiling of the box, retains substantial momentum. In our experiments, little overturning (and vertical mixing) was observed as the plume spread to the full horizontal extent of the box and interacted with the sidewalls. However, as noted by Baines & Turner (1969) and investigated by Hunt, Cooper & Linden (2001) and Kaye & Hunt (2007), substantial overturning can occur in different geometries and with different source conditions.

We conclude by discussing the implications of our findings for the design of building ventilation schemes, noting that scenarios in which prevailing wind and buoyancy forces act in opposition to one another are not uncommon. If wind strength is large enough for multiple equilibria to exist, our analysis can be used to determine unambiguously which state is attained for given initial conditions. If ambient air is at a pleasant temperature, it may be desirable to establish buoyancy-driven displacement flow with the interface above the level of room occupants, whereas if ambient air is much colder than the target occupant temperature, wind-driven mixing ventilation may ensure greater thermal comfort. In either case, our analysis can be used to develop control strategies that ensure adoption of the preferred state. If the filling box time T_{Bf} (defined in (2.7)) for a space is sizeable (for instance 10 min or longer), occupant thermal comfort during transient phases must also be considered. Of significance

here are our findings that similar buoyancy-driven final states can have qualitatively different approaches to equilibrium. Merely ensuring that the equilibrium interface height is above the level of all occupants is thus not sufficient to guarantee a good displacement ventilation system, as the layer of hot (and contaminated) air may descend even as far as the floor before settling at its desired level.

We are extremely grateful to S. Dalziel for his advice on many aspects of our laboratory study and to D. Page-Croft and J. Milton for their technical assistance. Financial support for I.A.C. in the form of a CASE Industrial PhD grant was generously provided by the E.P.S.R.C. and Arup, facilitated by the Knowledge Transfer Network for Industrial Applied Mathematics. We particularly wish to thank G. Davies and D. Allwright for their advice, support and technical input.

REFERENCES

- BAINES, W. D. & TURNER, J. S. 1969 Turbulent buoyant convection from a source in a confined region. *J. Fluid Mech.* **37**, 51–80.
- BOLSTER, D. & CAULFIELD, C. P. 2008 Transients in natural ventilation: A time-periodically-varying source. *Build. Ser. Engng Res. Technol.* **29** (2), 119.
- BOLSTER, D., MAILLARD, A. & LINDEN, P. 2008 The response of natural displacement ventilation to time-varying heat sources. *Energy Build.* **40** (12), 2099–2110.
- BOWER, D. J., CAULFIELD, C. P., FITZGERALD, S. D. & WOODS, A. W. 2008 Transient ventilation dynamics following a change in strength of a point source of heat. *J. Fluid Mech.* **614**, 15–37.
- CAULFIELD, C. P. 1991 Stratification and buoyancy in geophysical flows. PhD thesis, DAMTP, University of Cambridge.
- CAULFIELD, C. P. & WOODS, A. W. 1995 Plumes with non-monotonic mixing behaviour. *J. Geophys. Astrophys. Fluid Dyn.* **79** (1), 173–199.
- CAULFIELD, C. P. & WOODS, A. W. 2002 The mixing in a room by a localized finite-mass-flux source of buoyancy. *J. Fluid Mech.* **471**, 33–50.
- CENEDESE, C. & DALZIEL, S. B. 1998 Concentration and depth fields determined by the light transmitted through a dyed solution. In *Proc. 8th Intl Symp. on Flow Visualization* (ed. G. M. Carlomagno & I. Grant), ISBN 0953399109, paper 061.
- COOPER, P. & LINDEN, P. F. 1996 Natural ventilation of an enclosure containing two buoyancy sources. *J. Fluid Mech.* **311**, 153–176.
- DALZIEL, S. B., PATTERSON, M. D., CAULFIELD, C. P. & COOMARASWAMY, I. A. 2008 Mixing efficiency in high-aspect-ratio Rayleigh–Taylor experiments. *Phys. fluids* **20**, 065106.
- FLYNN, M. R. & CAULFIELD, C. P. 2006a Natural ventilation in interconnected chambers. *J. Fluid Mech.* **564**, 139–158.
- FLYNN, M. R. & CAULFIELD, C. P. 2006b Transient blocking in multi-chamber natural ventilation. In *Proc. 6th Intl Symp. on Stratified Flows* (ed. G. N. Ivey), pp. 476–481. University of Western Australia Press.
- FLYNN, M. R. & CAULFIELD, C. P. 2009 Effect of volumetric heat sources on hysteresis phenomena in natural and mixed-mode ventilation. *Build. Environ.* **44** (1), 216–226.
- GERMELES, A. E. 1975 Forced plumes and mixing of liquids in tanks. *J. Fluid Mech.* **71** (3), 601–623.
- HEISELBERG, P., LI, Y., ANDERSEN, A., BJERRE, M. & CHEN, Z. 2004 Experimental and CFD evidence of multiple solutions in a naturally ventilated building. *Indoor Air* **14** (1), 43–54.
- HUNT, G. R. & COFFEY, C. J. 2010 Emptying boxes – classifying transient natural ventilation flows. *J. Fluid Mech.* **646** (1), 137–168.
- HUNT, G. R., COOPER, P. & LINDEN, P. F. 2001 Thermal stratification produced by plumes and jets in enclosed spaces. *Build. Environ.* **36** (7), 871–882.
- HUNT, G. R. & HOLFORD, J. M. 2000 The discharge coefficient – experimental measurement of a dependence on density contrast. In *Proc. 21st AIVC Conference, The Hague, Netherlands*.
- HUNT, G. R. & KAYE, N. G. 2001 Virtual origin correction for lazy turbulent plumes. *J. Fluid Mech.* **435**, 377–396.

- HUNT, G. R. & LINDEN, P. F. 1998 Time-dependent displacement ventilation caused by variations in internal heat gains: application to a lecture theatre. In *Proc. Roomvent '98, the 6th Intl Conf. on Air Distribution in Rooms, Stockholm, Sweden*, vol. 2, pp. 203–210.
- HUNT, G. R. & LINDEN, P. F. 1999 The fluid mechanics of natural ventilation – displacement ventilation by buoyancy-driven flows assisted by wind. *Build. Environ.* **34** (6), 707–720.
- HUNT, G. R. & LINDEN, P. F. 2000 Multiple steady airflows and hysteresis when wind opposes buoyancy. *Air Infiltration Rev.* **21** (2), 1–2.
- HUNT, G. R. & LINDEN, P. F. 2001 Steady-state flows in an enclosure ventilated by buoyancy forces assisted by wind. *J. Fluid Mech.* **426**, 355–386.
- HUNT, G. R. & LINDEN, P. F. 2005 Displacement and mixing ventilation driven by opposing wind and buoyancy. *J. Fluid Mech.* **527**, 27–55.
- KAYE, N. B., FLYNN, M. R., COOK, M. J. & JI, Y. 2009 The role of diffusion on the interface thickness in a ventilated filling box. *J. Fluid Mech.* **652**, 195–205.
- KAYE, N. B. & HUNT, G. R. 2004 Time-dependent flows in an emptying filling box. *J. Fluid Mech.* **520**, 135–156.
- KAYE, N. B. & HUNT, G. R. 2007 Overturning in a filling box. *J. Fluid Mech.* **576**, 297–323.
- LI, Y. & DELSANTE, A. 2001 Natural ventilation induced by combined wind and thermal forces. *Build. Environ.* **36** (1), 59–71.
- LI, Y., DELSANTE, A., CHEN, Z., SANDBERG, M., ANDERSEN, A., BJERRE, M. & HEISELBERG, P. 2001 Some examples of solution multiplicity in natural ventilation. *Build. Environ.* **36** (7), 851–858.
- LINDEN, P. F. 1999 The fluid mechanics of natural ventilation. *Annu. Rev. Fluid Mech.* **31** (1), 201–238.
- LINDEN, P. F. 2000 Convection in the environment. In *Perspectives in Fluid Mechanics* (ed. G. K. Batchelor, H. K. Moffatt & M. G. Worster), pp. 287–343. Cambridge University Press.
- LINDEN, P. F., LANE-SERFF, G. F. & SMEED, D. A. 1990 Emptying filling boxes: the fluid mechanics of natural ventilation. *J. Fluid Mech.* **212** (303), 309–335.
- LISHMAN, B. & WOODS, A. W. 2009 On transitions in natural ventilation flow driven by changes in the wind. *Build. Environ.* **44** (4), 666–673.
- MORTON, B. R., TAYLOR, G. & TURNER, J. S. 1956 Turbulent gravitational convection from maintained and instantaneous sources. *Proc. R. Soc. Lond. A* **234**, 1–23.
- STERN, N. H. 2007 *The Economics of Climate Change: The Stern Review*. Cambridge University Press.
- TURNER, J. S. 2006 Turbulent entrainment: the development of the entrainment assumption, and its application to geophysical flows. *J. Fluid Mech.* **173**, 431–471.
- WOODS, A. W., CAULFIELD, C. P. & PHILLIPS, J. C. 2003 Blocked natural ventilation: the effect of a source mass flux. *J. Fluid Mech.* **495**, 119–133.
- WORSTER, M. G. & HUPPERT, H. E. 1983 Time-dependent density profiles in a filling box. *J. Fluid Mech.* **132**, 457–66.
- YUAN, J. & GLICKSMAN, L. R. 2007 Transitions between the multiple steady states in a natural ventilation system with combined buoyancy and wind driven flows. *Build. Environ.* **42** (10), 3500–3516.
- YUAN, J. & GLICKSMAN, L. R. 2008 Multiple steady states in combined buoyancy and wind driven natural ventilation: the conditions for multiple solutions and the critical point for initial conditions. *Build. Environ.* **43** (1), 62–69.