Mitigation of dense gas releases within buildings: validation of CFD modelling

S. Gilham *, D.M. Deaves, P. Woodburn

WS Atkins Science and Technology, Woodcote Grove, Ashley Road, Epsom, Surrey KT18 5BW, UK

Abstract

When an accidental release of a hazardous material is considered within a safety case or risk assessment, its off-site effects are generally assessed by calculating the dispersion of vapour from the site. Although most installations handling flammable materials will be in the open air, many types of plant, particularly those handling toxics, are enclosed, partly to provide some form of containment and hence, to mitigate the effects of any release. When such a release occurs within a building, the gas or vapour will undergo some mixing before emerging from any opening. The degree of mixing will depend upon the building geometry and the nature of the ventilation, which in turn may be modified by the leak. This situation is considered in this paper, with specific application to calculating the rate of release of a dense vapour from a building. The paper describes the application of computational fluid dynamics (CFD) techniques to modelling the release and mixing processes within buildings. Examples of validation calculations for simple geometric arrangements, as well as more complex geometries representative of an industrial site, are described. The results demonstrate the capabilities of CFD for this application but highlight the need for careful modelling of the near-wall flows and heat transfer, and need for an accurate fluid dynamics and thermodynamic representation of the release source. © 2000 Elsevier Science B.V. All rights reserved.

Keywords: CFD; Gas build-up; Building mitigation; Validation; Risk assessment

1. Introduction

In many industrial installations, it is possible for toxic or flammable materials to be released accidentally as gas or rapidly vaporising liquid within buildings. The impact on the surrounding population of such a release can be assessed by undertaking gas

---

* Corresponding author. WS Atkins Safety and Reliability, Woodcote Grove, Ashley Road, Epsom, Surrey KT18 5BW, UK. E-mail: kcook@waatkins.co.uk

dispersion calculations using a suitable source term. Determination of the source term is frequently undertaken either by ignoring the effects of the building, or by assuming complete mixing and unmodified natural ventilation within the building. To improve such estimates, research has been undertaken into this problem, using both advanced computational fluid dynamics (CFD) techniques and simpler 'zone' type models (see Deaves et al. [1]).

This paper describes the application of a commercial CFD code, STAR-CD, to contaminant dispersion, focusing on toxic vapour releases in a chlorination room and also in two more simple test geometries. The capabilities of the models are assessed through comparisons with experimental data, and the sensitivity of the model predictions to parameters such as the wall heat transfer model, the turbulence model and the method used to characterise the release source are also presented.

2. The physical processes and their modelling

2.1. Physical processes

The situation considered is an accidental release of a toxic material as a result of a flange, pipework or storage tank failure. This generally results in either a vapour release (from the tank vapour space or vapour pipework) or a flashing liquid release. In contrast to flammable materials where rapid dilution to below the LFL is desirable, the principal need for toxic releases is to contain the release. For this reason, many toxic materials are stored and processed within relatively well-sealed buildings which can temporarily contain the release and thus mitigate its effect. Additionally, the work described here has focused on materials which form denser-than-air vapours such as ammonia and chlorine. A typical example would be the chlorination room of a water treatment works.

In general, the degree of mitigation provided by the containment building, and thus the rate at which the toxic material is released to the atmosphere, is determined by the ventilation characteristics of the building and the details of the release. For relatively well-sealed buildings, the extent of mixing is principally driven by the release characteristics, the release momentum and its geometry determining the extent of the mixing. The location and flow characteristics of any adventitious openings in the building then determine the transient release signature from the building. In these circumstances, the mixing and the consequent release from the building are driven by the source itself and its complex fluid dynamic and thermodynamic behaviour.

For the industrial situation, there is an additional complication that the buildings in which a release is likely to occur are large and have complex geometries, typically including storage tanks, bunds and extensive pipe work runs.

For all but the most severe accidental releases, where all containment is lost, the geometrical scales of the release are likely to be small compared with those of the building. The release typically lasts of the order of tens of minutes, and is characterised by a two-phase flashing high-speed liquid jet, which through impingement can also result in the formation of an evaporating pool, although some releases may be of vapour only. Both the mixing induced by the release and its thermodynamic behaviour are

important in determining the building release characteristics. This range of physical processes, geometrical scales and time scales presents a number of significant modelling challenges: resolution of flow features which range from the details of the release to the flow within the room; a range of turbulence levels which can potentially vary from high Reynolds number flows (in the jet) to low Reynolds number or transitional flow; geometrical features which vary in scale from room scale to the scale of the pipework; transient dense gas build-up with the potential for significant stratification; two-phase flow and heat transfer.

In principal, all these processes can be incorporated within a CFD model, albeit with varying levels of confidence. An alternative to using CFD is to use zone-type models. These, however, require some gross simplifications and assumptions to be made and are thus not able to incorporate the complete range of physical processes described above. This approach has, however, been considered in addition to the CFD modelling described here, and the results are presented by Deaves et al. [1].

2.2. Computational model

The code used for all the calculations in this paper was the STAR-CD code, developed by Computational Dynamics [2]. The code is an unstructured finite-volume code capable of both steady-state and transient calculations. Hybrid differencing was used for all equations except pressure.

The Reynolds-averaged Navier–Stokes momentum equations for this model have the following form:

$$\frac{\partial \rho U_i}{\partial t} + \frac{\partial \rho U_i U_j}{\partial x_j} = -\frac{\partial p}{\partial x_i} + \frac{\partial}{\partial x_j}\left(2\mu S_{ij} - \rho\overline{u_i' u_j'}\right) - \rho g \delta_{i2}. \tag{1}$$

The turbulence modelling strategy for this project was based on the standard $k - \varepsilon$ turbulence model (Launder and Spalding [3]) which has been frequently used for this type of work. Limited buoyancy effects are included in this model (Rodi [4]). The equations for k and ε are as follows:

$$\frac{\partial k}{\partial t} + U_i\frac{\partial k}{\partial x_i} = \frac{\partial}{\partial x_j}\left(\frac{\nu_t}{\sigma_k}\frac{\partial k}{\partial x_j}\right) + P + G - \varepsilon, \tag{2}$$

$$\frac{\partial \varepsilon}{\partial t} + U_i\frac{\partial \varepsilon}{\partial x_i} = \frac{\partial}{\partial x_j}\left(\frac{\nu_t}{\sigma_\varepsilon}\frac{\partial \varepsilon}{\partial x_j}\right) + \frac{\varepsilon}{k}\left[C_{\varepsilon 1}P + C_{\varepsilon 3}G\right] - C_{c2}\frac{\varepsilon^2}{k}, \tag{3}$$

where

$$-\overline{u_i' u_j'} = 2\nu_t S_{ij} - \frac{2}{3}k\delta_{ij} \quad \nu_t = C_\mu\frac{k^2}{\varepsilon},$$

$$P = \nu_t S_{ij}\frac{\partial U_i}{\partial x_j} \quad G = \frac{g}{\rho}\frac{\nu_t}{\sigma_t}\frac{\partial \rho}{\partial y} \quad S_{ij} = \frac{1}{2}\left[\frac{\partial U_i}{\partial x_j} + \frac{\partial U_j}{\partial x_i}\right],$$

U_i are the mean velocity components in the three Cartesian co-ordinate directions (x_i); k, ε are the turbulence energy and dissipation rate; P, G are the shear and buoyancy production rates of k; ν_t is the kinematic turbulent (eddy) viscosity; S_{ij} is the shear strain tensor; ρ is the density; and σ_t, σ_k, σ_ε are the turbulent Prandtl numbers for temperature, k and ε.

For some of the simulations, a two-layer variant of this model (Rodi [5]) was used, in which the near-wall region is simulated via a k transport equation and an algebraic prescription for the turbulent length scale, given by the van Driest [6] mixing length:

$$l_m = \kappa\, y\,(1 - e^{-D}), \tag{4}$$

where l_m is the mixing length, κ has value 0.42, y is the distance to the wall, and D is given by:

$$D = \frac{y^+}{26}, \tag{5}$$

y^+ is a non-dimensional near-wall spacing.

3. Application of CFD to simple test geometries

Before applying CFD to the complex and realistic case of an accidental release in a chlorination room (described in Section 4), a number of simpler test cases were studied. This was, in part, due to the lack of availability of relevant datasets prior to the experiments described here. Two cases are described: flow and dispersion in the IEA test room (Moser [7]) and flow and dispersion in the SILSOE test room (Gilham et al. [8]).

3.1. IEA test room

An extensive series of tests has been conducted as part of the International Energy Agency's collaborative project 'Annex 20, Air Flow Patterns within Buildings'. In total, six different experimental programmes were conducted, and from these two cases were considered: isothermal flow within the room and isothermal flow with a dispersing containment. The room was a simple rectangular construction, 4.2 m long, 3.6 m wide and 2.5 m high. The test room was force-ventilated with ventilation rates ranging from 1 to 6 air changes per hour (ACH), and for the contaminant dispersion tests, a momentum-less source of a dense contaminant was introduced at the centre of the room. The room geometry is shown in Fig. 1.

For the isothermal flow test case with 3 ACH, the calculation was carried out in half the room (a plane of symmetry was assumed for the calculations) and used the two-layer, buoyancy-extended $k - \varepsilon$ turbulence model with standard buoyancy correction terms described in Section 2.2. It was not possible to resolve the details of the ventilation inlet diffuser, which was therefore simulated indirectly using appropriate sub-models. The predicted velocity distribution on the vertical symmetry plane is
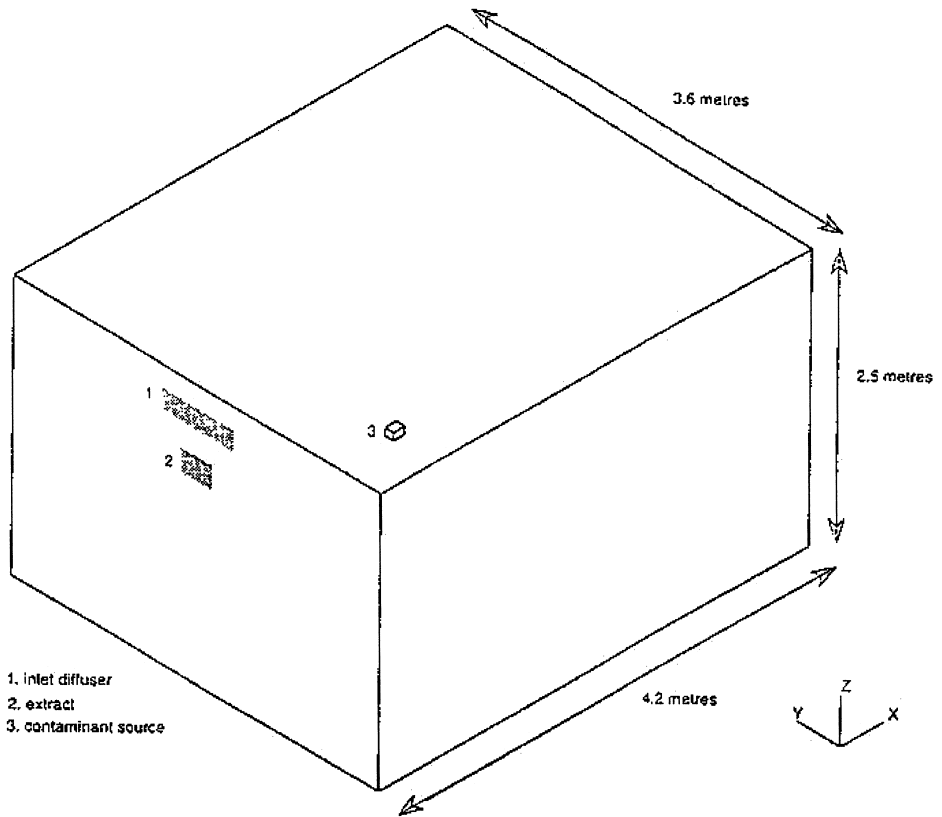
Fig. 1. Schematic diagram of the IEA test room.

compared with the experimental measurements in Fig. 2; the good agreement achieved for this case with this approach is illustrated. The flow is dominated by the wall jet along the ceiling generated by the inlet diffuser which is directed towards the ceiling and causes the recirculating flow seen in Fig. 2. Calculations were also carried out using the standard $k - \varepsilon$ turbulence model with a coarser calculation mesh; these also demonstrated reasonable agreement, although the near-wall flow was less well-predicted.

For the contaminant dispersion calculations, the ventilation flow rate was 1.5 ACH and the slow release (0.025 1/s) of dense contaminant was introduced as a momentumless source using an appropriate sub-model. The calculation was again carried out using the two-layer $k - \varepsilon$ turbulence model. The computational mesh was designed to resolve the dense plume and its interaction with the floor boundary layer. Fig. 3 compares the computed and measured normalised contaminant concentrations on the room symmetry plane. The dense contaminant forms a layer on the floor and a stratified flow results. The computation reproduces this behaviour well although the near floor concentrations are slightly overpredicted. Computations were also carried out using the standard buoyancy-extended $k - \varepsilon$ turbulence model, which caused a more smeared shear layer due to increased vertical dispersion.
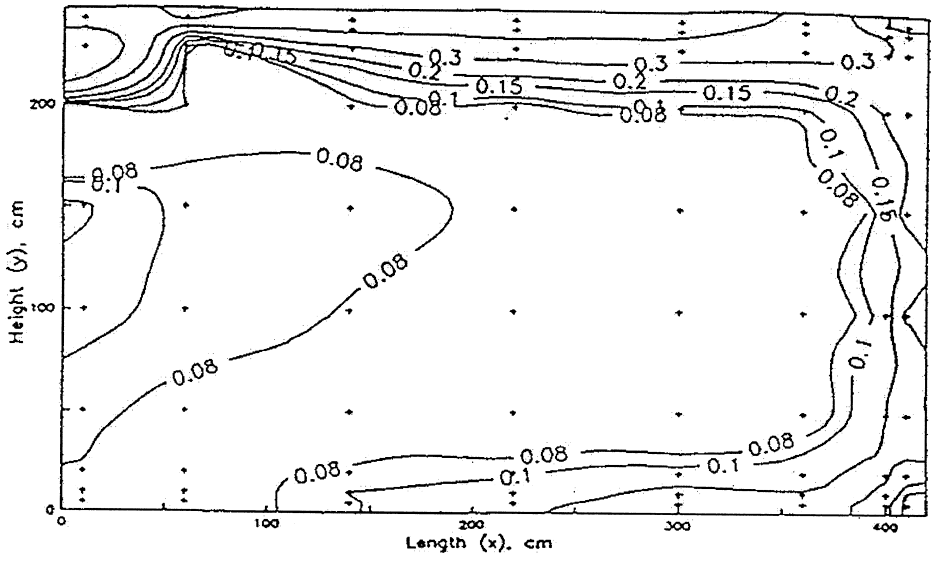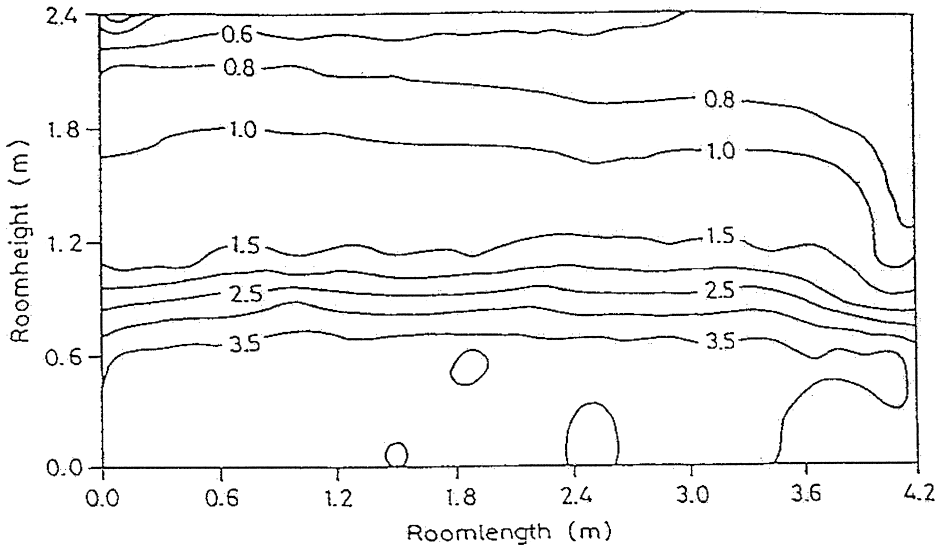
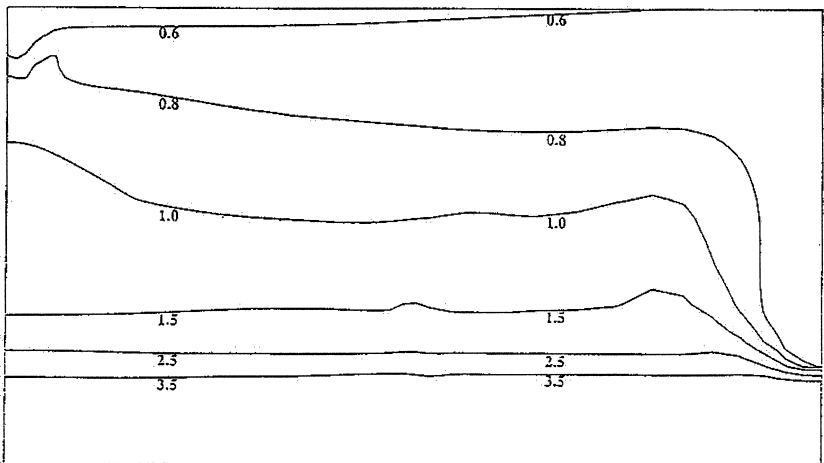(a)   Experimental results
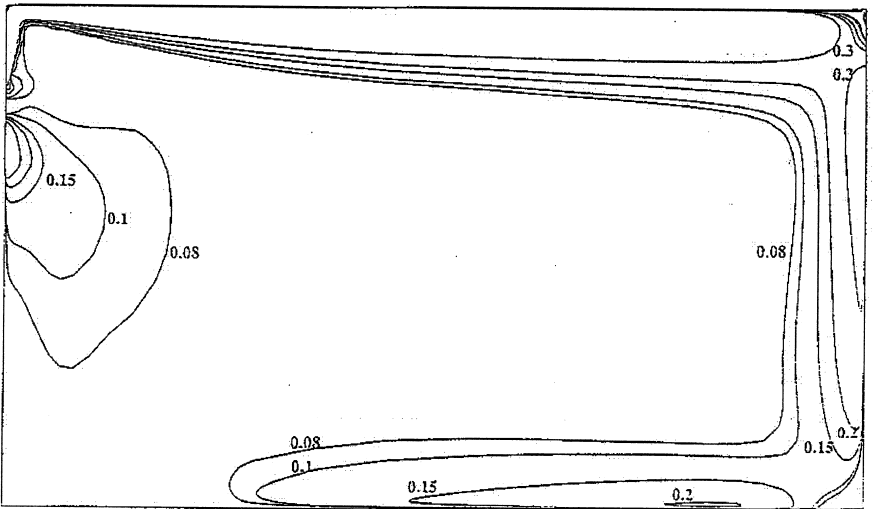
(b)   Computed results, two layer turbulence model

Fig. 2. Velocity magnitude contours on the symmetry plane of the IEA test room.

(a) Experimental results

(b) Concentration contours on symmetry plane

Fig. 3. Contaminant concentration contours on the symmetry plane.

3.2. SILSOE test room

The SILSOE test room [8] was a 2.4-m cube with vents positioned on the centreline of one of the walls to prevent pressurisation of the room. A schematic diagram of the test room showing the location of the sensors is shown in Fig. 4. Carbon dioxide was introduced through a 2-in. diameter pipe at the centre of the same wall as the vents. Tests were carried out for two release rates and with different combinations of ventilation openings.

For the 600 l/min release rate, and with only the top outlet open, the calculation was carried out within half the room, and employed the standard buoyancy-extended $k - \varepsilon$ turbulence model. Care was also taken to resolve the near source region. Fig. 5 compares the predicted and measured concentration time histories throughout the duration of the release and shows good qualitative and quantitative agreement between the experimental results and the computations. There are slight differences between the vertical concentration profiles; the measurements indicate a uniform concentration over the lower 1 m during the first 150 s of the release. After this time, the 1-m
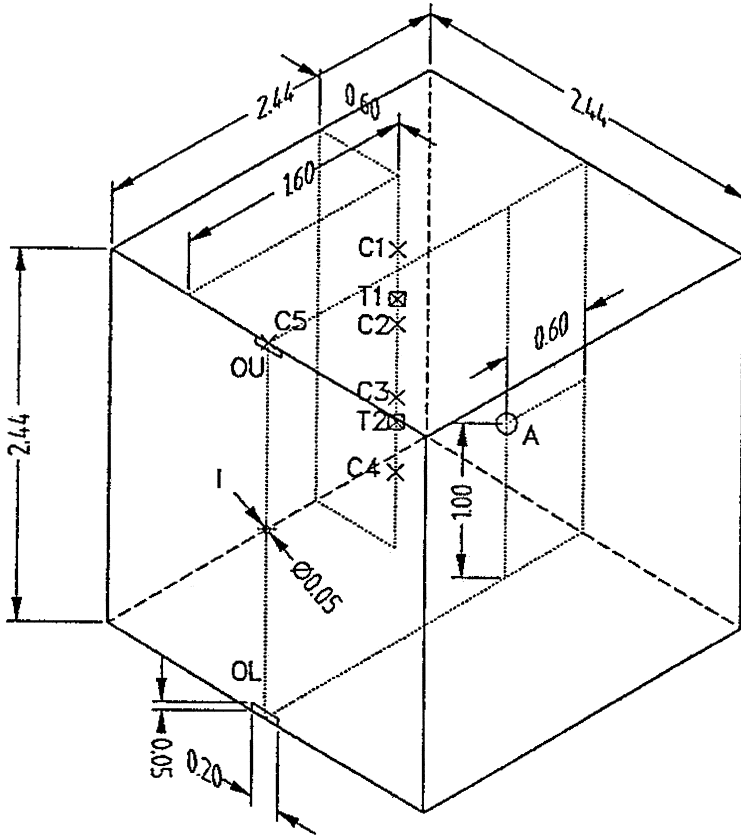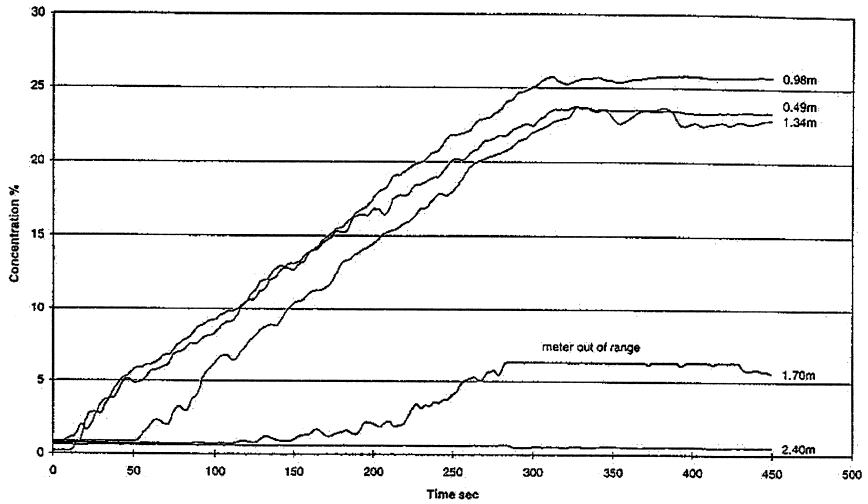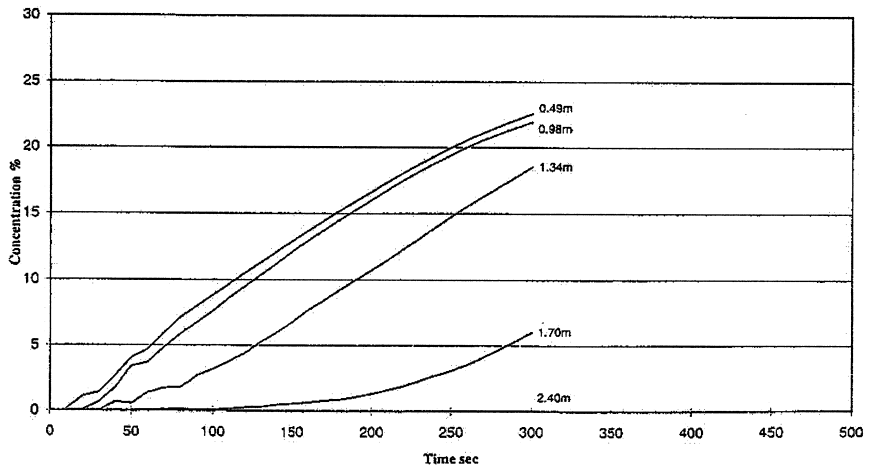

Fig. 4. Schematic diagram of the SOILSOE test room (I — inlet; OU and OL — upper and lower outlets; C1 to C5 — concentration sensors; T1, T2 — temperature sensors; A — ultrasonic anemometer).

a) **Experimental Results**

b) **Computed Results**

Fig. 5. CO_2 concentration time history for the 600 l/min release in the SILSOE test room.

concentrations exceed those at 0.5 m; this behaviour is not reproduced in the computations. There is also some evidence that the calculations predicted a greater smearing of the vertical concentration profile.

3.3. Implications for industrial cases

Although the calculations carried out for the simpler cases did not incorporate all the features of the industrial case, they were instructive in identifying many of the important

parameters. The most significant parameters over which the user of a CFD code has control are the mesh density and its distribution. It was evident that considerable care needs to be exercised when constructing the mesh, particularly in resolving source and near-wall regions. The buoyancy extended $k - \varepsilon$ turbulence model proved to be robust and gave reasonable predictions, although improved performance was achieved using a two-layer turbulence model, particularly in regions of low Reynolds number flow. The $k - \varepsilon$ model tended to overpredict the vertical mixing and this was particularly evident for cases where significant stratification occurred. Using the two-layer turbulence model gave improved solutions, principally as a result of better resolution of regions where the flow velocities were low. It was also evident from the calculations that sub-models can be employed successfully to define flow features such as inlets and sources which could not be resolved in detail, provided that care is taken to ensure that they contain appropriate physics.

4. Experimental measurements in the model scale chlorination room

To address the deficiency in experimental data noted above, an experimental study was carried out by Bettis and Allen [9] in a one-third linear scale model of a chlorination room of a water treatment works. The test room geometry was designed to incorporate many of the complex geometrical features of an actual industrial site. Freon-22 was used as the working fluid because of its similarities with chlorine. Three releases were considered: two jet releases and an evaporating pool release. Temperature, gas concentration and velocity measurements were taken for each of the three releases.

4.1. Geometry of test room

Fig. 6 shows the geometry of the one-third scale test room. It consisted of two horizontal cylindrical storage tanks surrounded by a bund. The tanks were arranged horizontally and were surmounted by a walkway. Two large vertical cylinders, which were representative of scrubber units in a full scale room, stood adjacent to the storage area. Also present was a smaller horizontal tank which represented a typical pressure relief vessel. The test room, constructed within a concrete bunker, was 7.31 m long by 4.87 m wide by 2.44 m high and had a free air volume of 52.74 m^3.

The test room was unventilated with the exception of the full height vertical extractor duct located in one of the corners. The duct was square in cross-section and bifurcated above the ground before terminating in two openings just above the floor. The room was sealed to ensure that no uncontrolled ventilation occurred.

The external walls of the room were constructed from 12-mm thick hardwood stage panels, the ceiling and internal walls from 13-mm thick MDF and the cylinders and internal detail from laminated cardboard. The floor of the concrete bunker also formed the floor of the enclosure, and, since it was partially buried in a hillside, was maintained at a relatively low, steady temperature.
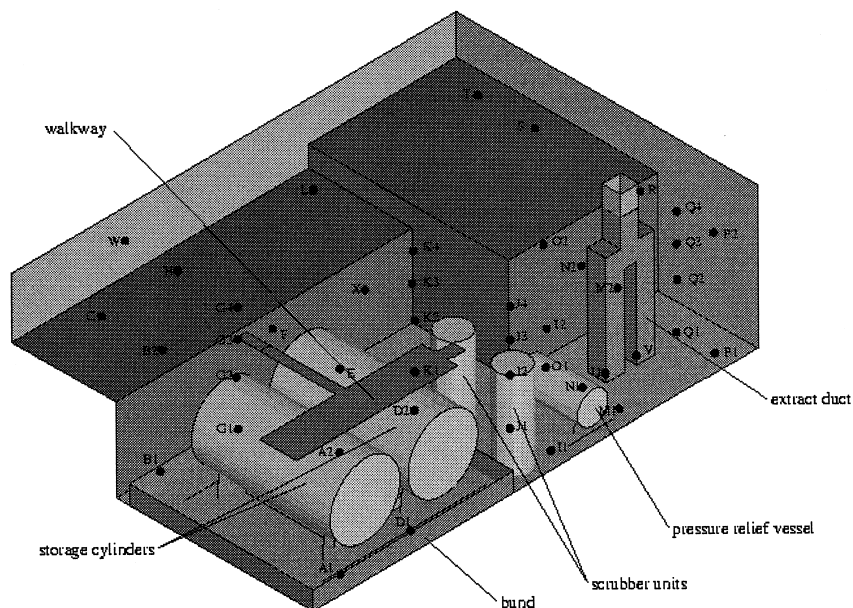
Fig. 6. Geometry of the test room (●, measurement locations).

4.2. General details of the trials

A number of preliminary tests were undertaken by Bettis and Allen [9]. These included the characterisation of the jet sources in the laboratory, and 'dry runs' in which the release mechanisms and measurement setup were tested. There were then three 'production runs', detailed in Table 1.

The experiment utilised a total of 44 thermocouples, 21 gas detectors, seven 'hot wire' anemometers, one Laser Doppler Anemometer and one ultrasonic anemometer. Forty of the thermocouples, 20 of the gas detectors and six of the hot-wires were arranged in 19 vertical arrays, with one, two or four measuring points in each array. An isometric view of the room showing the measurement locations is given in Fig. 6.

4.3. Jet sources

For the two jet releases, the source nozzles were located above the walkway and the jet was directed along the axis of the tanks towards the far wall. In both cases, the

Table 1
Details of experimental trials

Trial number [4]	Release type	Release rate and duration	Duration of measurements [s]
3	small jet	0.0117 kg/s over 1800 s	2000
4	large jet	0.0320 kg/s over 1300 s	1500
5	evaporating liquid pool	pool filled at 0.052 kg/s over 360 s	3600

pressurised freon-22 exited the nozzle as a mixture of vapour and liquid. Mass flow rates for each case were calculated from the time history of readings from load cells on which the freon-22 bottles were placed. For the large release (0.032 kg/s), the nozzle diameter was 2.8 mm and for the small release (0.0117 kg/s), a pinhole nozzle 1 mm in diameter was used.

Prior to commencing the measurement programme within the test room, detailed measurements were made of the jet source characteristics within a separate test chamber. For each release rate, measurements of temperature and velocity were made across the jet at a single downstream location (750 mm downstream for the large jet, 500 mm for the small jet). In both cases, the downstream location of the measurements was chosen to be beyond the distance of minimum centreline temperature, in order to be outside the flashing region of the jet.

4.4. Evaporating pool release

The evaporating pool release was formed by releasing freon-22 into a 1.45-m^2 horizontal steel tray located centrally beneath the two storage tanks. The pool was generated by injecting the liquid freon-22 into four 100-ml plastic beakers filled with ceramic beads to make a porous bed. The beakers removed the momentum of the freon-22, allowing it to spill over into the pool without the momentum or the entrainment of the free jet. This resulted in a release which occurred in two distinct phases: an initial phase which occurred as the pool is filled with liquid freon-22, and a secondary phase after the filling was complete.

Both the freon-22 cylinders and the tray containing the pool were mounted upon load cells in order to evaluate the mass flow rate of freon-22 into the pool during filling, and mass flow out of the pool due to evaporation. The filling phase lasted 6 min, during which time 18.7 kg of liquid freon-22 were injected into the pool, although much of this evaporated immediately due both to flashing as the liquid depressurised and to boiling as the liquid absorbed enthalpy from its surroundings. The maximum amount of freon-22 contained in the pool was 10.2 kg which occurred at 360 s, indicating a mean evaporation rate of 0.024 kg/s over the initial period. After the filling was complete, the residual freon-22 evaporated at a much lower rate, which then further decayed.

5. The CFD model of the chlorination room

CFD model predictions were undertaken before the experiments were undertaken. These preliminary calculations were used to examine the suitability of the initially proposed release conditions and to provide data to assist in the location of the instrumentation. The experiments were then carried out and the CFD calculations repeated with the measured source conditions. These 'base case' calculations were then compared with the measurements. Having compared the results, additional calculations were undertaken to examine the sensitivity of the predictions to variations in a number of model parameters: near-wall mesh spacing, heat transfer modelling, jet source modelling, and turbulence model.

5.1. Computational mesh

An illustration of a computational mesh used for the initial computations is given in Fig. 7. This shows the mesh around the solid surfaces of the room. For the majority of the calculations, the mesh was as shown in the figure and was constructed from 55,693 computational cells. For the pool release, however, a more refined mesh was adopted in which a small amount of mesh embedding was used around the evaporating pool source and near the enclosure floor, increasing the number of cells for that release to 87,450. For the jet release cases, the average cell volume was 0.001 m^3 and the average cell linear dimension was 0.1 m, while for the evaporating pool release, the average cell volume was 0.0006 m^3 and the average linear cell dimension 0.085 m.

5.2. Physical models

For all the computations, turbulent, non-isothermal flow was assumed. The standard $k - \varepsilon$ turbulence model detailed in Section 2.2 was used. The simulations were non-Boussinesq, and densities were calculated using the ideal gas law. Although the releases were two-phase at their source, two-phase flow was not modelled. The releases were introduced into the computational domain at a point when all the liquid droplets within the release had evaporated and were modelled as appropriate sources of freon-22 vapour, momentum and enthalpy. The dispersion of the freon-22 vapour was modelled using a scalar transport equation, and the fluid properties were calculated locally from a volume-weighted average of the background fluid, assumed to be dry air, and the freon-22. The releases were modelled as time-varying sources and the development of
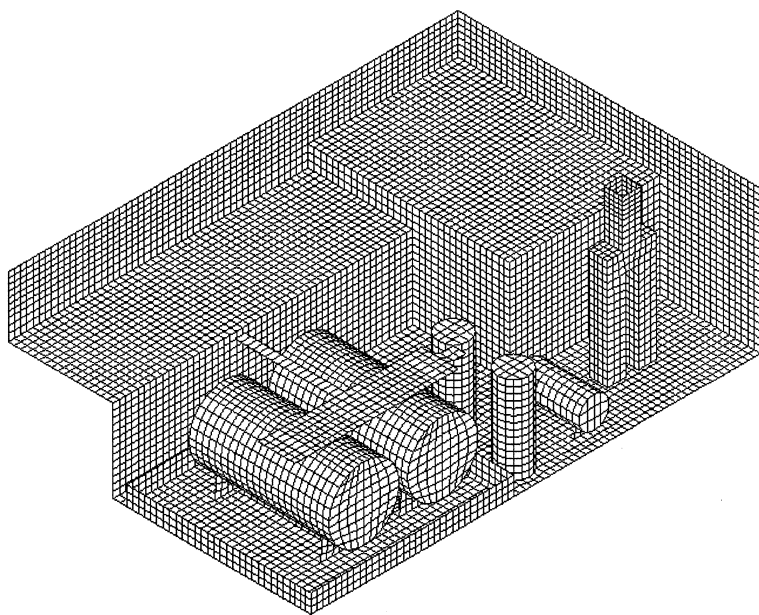

Fig. 7. Isometric view of the calculation mesh.

the vapour cloud within the test room was simulated using a fully transient solution scheme.

5.3. Boundary and initial conditions

For the velocity field, no-slip boundary conditions were introduced using a wall function, in which the velocity at the cell nearest the wall was determined from empirical relationships.

The wall heat transfer boundary conditions were specified in terms of the wall material thermal conductivity, the wall thickness and the external temperature. These boundary conditions assume one-dimensional conduction through the wall whose exterior surface was maintained at a constant temperature (i.e. the ambient air temperature in the bunker in which the experiment was set up). The heat flux was calculated automatically by the CFD code as a function of the fluid temperature at the wall–fluid interface. Isothermal wall boundary conditions were imposed on the concrete floor of the test room. For the base case calculations, the standard forced convection, high Reynolds number, and wall function formulation were used to calculate the heat transfer at the walls.

Quiescent, isothermal initial conditions were assumed.

5.4. Discretisation

All runs were transient, using the implicit PISO algorithm with first-order backward differencing for the temporal discretisation. Time steps for the large and small jet releases were 0.25 and 0.5 s, respectively, and were 0.1 s for the evaporating pool release. STAR-CD's 'self-filtered central differencing' scheme was used for the spatial discretisation of velocity, turbulence, temperature and scalar distribution for the jet release calculations. In the evaporating pool release calculations, the flow was buoyancy-driven and consequently, it is much more difficult to obtain converged solutions than in the convection-dominated jet releases. The more diffusive, but more stable, upwind differencing scheme was therefore used for all variables in the evaporating pool release.

5.5. Jet source characterisation

The measured jet velocity profile was used to evaluate the jet volumetric flow rate and was combined with the temperature data to provide estimates of the jet concentration and entrainment. This allowed the momentum introduced into the room by the jet source to be estimated. The flowrates and flow duration of the jet source in the experiments are given in Table 1.

Although STAR-CD allows the direct modelling of multi-phase flows, for flashing releases, this approach requires an extremely fine mesh resolution which is expensive both in the number of mesh cells and in the number of time steps required because of the resultant restrictions placed on the time step size. Instead, the properties of the jet were specified at the downstream location where the characterisation measurements were made in the experiments, which was well into the vapour-only region of the jet. The jet was modelled using sub-grid scale sources of momentum and freon-22 calculated from the jet characterisation data, and an enthalpy sink representing the energy

extracted from the room as the liquid fraction of the jet at exit from the nozzle evaporated, calculated from the flash fraction and the heat of vaporisation as:

$$H_{\text{sink}} = (1 - X) M_{\text{R}22} h_{\text{fg}}, \tag{6}$$

where H_{sink} is the enthalpy sink, X is the flash fraction, $M_{\text{R}22}$ is the freon-22 mass release rate, and h_{fg} is the heat of vaporisation of freon-22. The flash fraction was calculated, assuming an isentropic expansion from saturation pressure at storage conditions at 10°C.

5.6. Evaporating pool source characterisation

The pool release occurs in two distinct phases: an initial phase over the first 6 min during which the pool was being filled with liquid freon-22, and a secondary phase in which the remaining liquid evaporated from the pool. The data for the experimental test are shown in Table 1. The rate of evaporation during the initial phase (approximately 0.023 kg/s) was much higher than during the secondary phase (approximately 0.005 kg/s at the start of the phase) due to the quantity of freon-22 that flashed off immediately as it depressurised. This, combined with the underlying evaporation which occurred as a result of the enthalpy gained by the freon-22 from the surrounding environment, produced the initially high release rate. During the secondary phase, the evaporation rate decreased exponentially to approximately 0.001 kg/s after 1 h.

The evaporating pool was modelled assuming that the refrigerant spreads to fill the bunded area instantaneously. The pool was then simply represented by a layer of cells into which the contaminant was directly injected at the appropriate rate as a sub-grid scale source. The temperature of the injected fluid was specified as the boiling temperature of freon-22 at atmospheric pressure, which is consistent with the approach used for the jet. An isothermal boundary condition was applied at the base of the pool and it was assumed that the energy required to evaporate the liquid fraction was taken from the release apparatus. An enthalpy sink was therefore unnecessary.

6. Results of CFD calculations for the chlorination room

6.1. Jet sources

Despite the considerable density of the freon-22, the momentum of the large jet was such that its trajectory dipped very little, and therefore impinged on the wall directly opposite the release point. This is illustrated in Fig. 8, which shows the velocity flow field plotted on a vertical plane through the jet source after 600 s. The jet rapidly reached a pseudo-steady-state, in which the flow field remained virtually unchanged while the temperature and the concentration varied monotonically. A generally well-mixed flow field resulted, in which there was no indication of any significant stratification due to the density of the freon-22.

Fig. 9 compares the temperatures predicted by the CFD model with the experimental measurements at sensor array K (located outside the bund) as shown in Fig. 6. While the
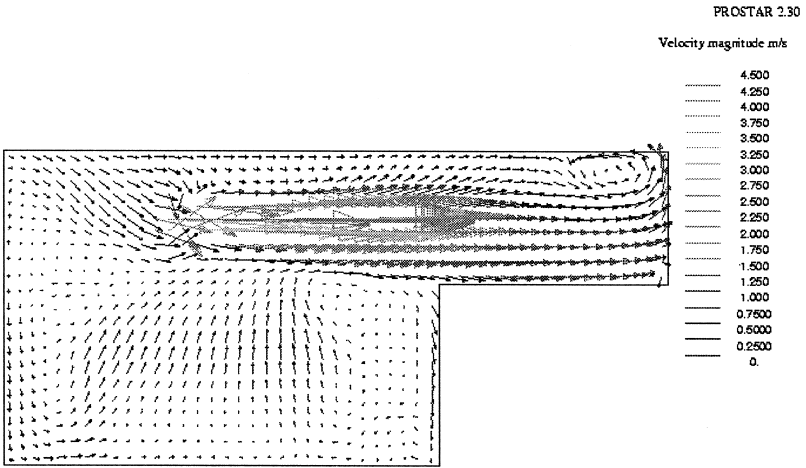
Fig. 8. Velocity vectors on cross-plane through jet (large jet case).

CFD prediction shows similar qualitative behaviour to the experimental data, the predicted temperatures are consistently lower. This underprediction is approximately 4°C towards the latter part of the transient at this location. It is likely that this is due to the wall heat transfer boundary conditions, and this is discussed further in Section 6.3. The initial discrepancy between the predicted and measured temperatures is due to slightly non-uniform conditions in the experiments. While the initial temperature in the predic-
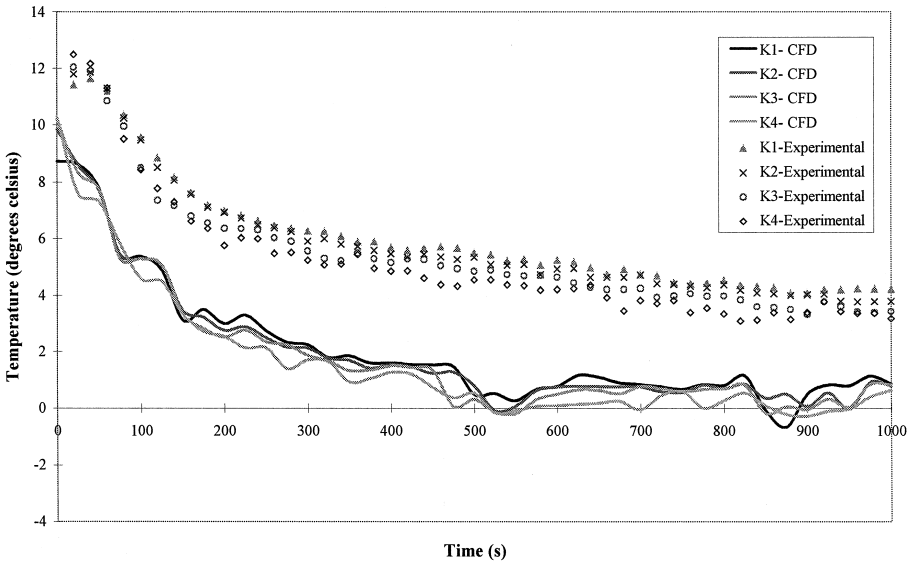

Fig. 9. Temperature time histories at sensor location K (large jet case).

tions was taken as 11°C, some of the experimental sensors recorded initial temperatures between 12° and 13°C.

Fig. 10 compares the volume concentrations of freon-22 from the CFD predictions with the experimental measurements at sensor array K. The agreement is much better than for the temperature measurements, although the predicted concentrations are generally slightly lower than those measured. The generally well-mixed nature of the flow, as noted above, is apparent in both the CFD and the experimental sensor measurements; little variation of concentration with height is observed.

The flow field for the small jet release was markedly different from that observed for the large jet release discussed above. The flow field after 500 s is plotted on a vertical plane through the centre of the jet in Fig. 11. The jet is significantly affected by the buoyancy of the dense freon-22, and dips sufficiently that it impinges on the corridor wall approximately midway across the room, before also impinging on the far wall. This is consistent with observations made during the characterisation phase of the experimental work, where the 1-mm jet was observed to dip significantly. The magnitudes of the predicted velocities are also considerably reduced for this small jet case as a result of its reduced momentum. In common with the large jet release, the flow field rapidly assumed a pseudo-steady-state and thereafter varied very little with time.

Unlike the large jet case, a degree of stratification occurred with the small jet experiments. This stratification was seen in the CFD predictions, although comparison of the predicted and measured temperature fields showed that the temperature levels were again lower. Stratification was also evident in the concentration field and this was predicted by the CFD model, although there was evidence that the turbulent mixing was overpredicted. These differences are similar to the large jet case and are further discussed in Section 6.3.
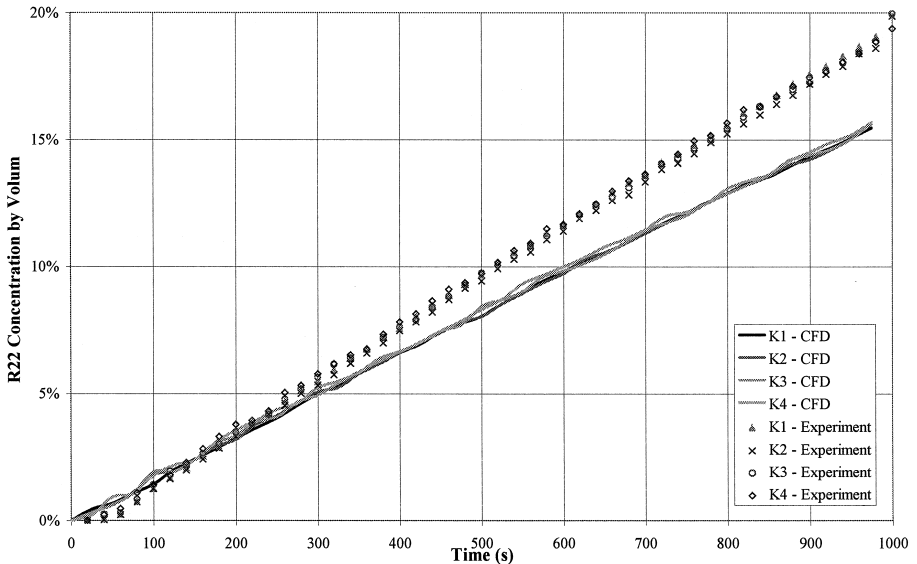

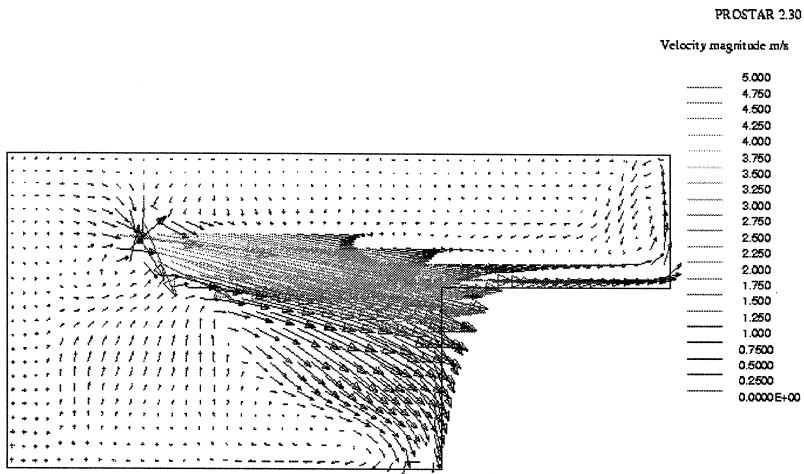Fig. 10. Concentration time histories at sensor location L (large jet case).

Fig. 11. Velocity vectors on vertical cross-plane through jet (small jet case).

It should be noted that only a model as detailed as the present model would have predicted the difference in flow fields caused by the large and small jets, the change in impingement which caused this as well as the resultant stratification, and the consequent increase of the release rate from the room.

6.2. Evaporating pool source

The velocities generated by the evaporating pool were considerably smaller than those observed in the jet releases. The upward velocity of the evaporating freon-22 was almost negligible and the flow field was dominated by the slumping of the dense freon-22 vapour, firstly spilling over the edge of the source tray and filling the bund, and then flowing over the floor once the bund had been filled, and resulting in the formation of a dense freon-22-rich layer next to the floor of the test room. As the freon-22 cloud spread out from the source region, its temperature increased through entrainment of the surrounding fluid and heat transfer at the surfaces.

The extract vent was located within the stratified dense floor layer and thus, the flow leaving the room had a much higher concentration of freon-22 than it would have had if the mixing in the room had been complete. The release rate from the extract duct is partly mitigated, however, by the reduction in the temperature of the room which caused air to be drawn into the room initially rather than expelled.

Figs. 12 and 13 compare the time histories of temperature at sensor locations D (located inside the bund) and K (located outside). At sensor D1, located close to the floor, the experimentally observed trend is reproduced computationally, though the temperature is significantly underpredicted. Well above the floor, at sensor D2, the temperature remains constant at its initial value in both the CFD predictions and the experiments. At sensor array K (just outside the bunded area), both the experimental measurements and the computational results suggest a simple stratification with the
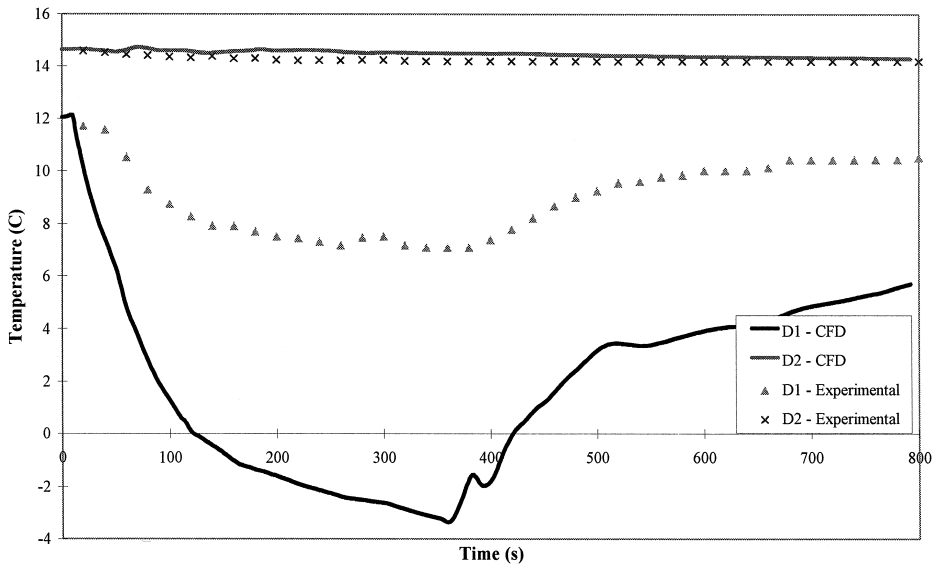
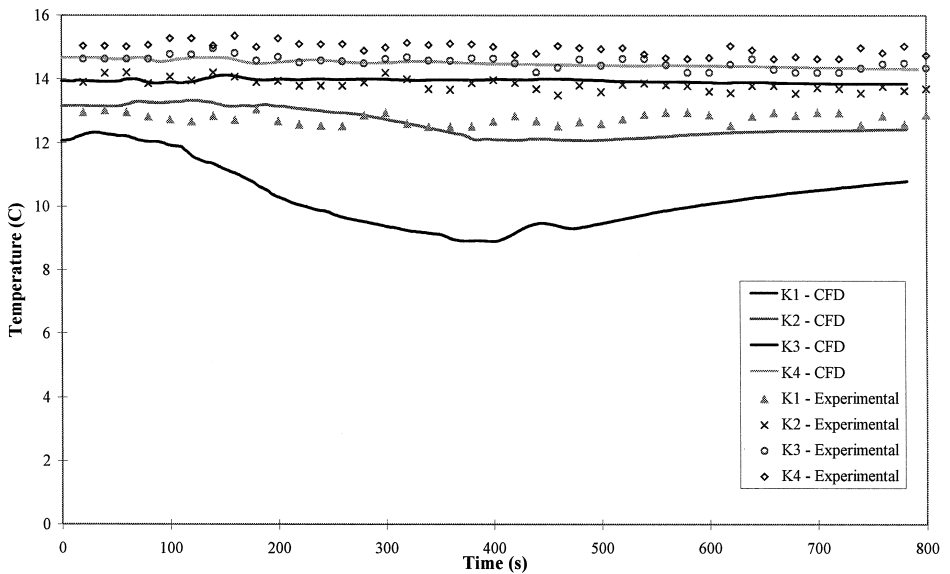Fig. 12. Temperature time histories at sensor location D (evaporating pool case).

Fig. 13. Temperature time histories at sensor location K (evaporating pool case).

temperature increasing with height. However, the computations predict that the temperature at the lower two sensors falls during the release, while the experimental measurements show that the temperatures remain constant during the release; this was the case for all measurement stations outside the bund. The underprediction of the temperatures within the bund indicates that the initial heating of the freon-22 around the source through entrainment and heat transfer at the wall is not modelled accurately.

Fig. 14 compares the CFD model predictions and the experimental time histories of freon-22 volumetric concentration at sensor array D. In both the experiments and simulations, the concentration at sensor D1 began to rise immediately and increased rapidly with time as the bund filled with the dense freon-22 vapour, although the predicted concentrations were lower than measured and the build-up was delayed slightly. Sensor D2 was located directly above D1, but above the freon-22 layer, and neither the CFD model nor the experiments record significant freon-22 concentration at D2. The generally lower predicted concentrations suggest that the computed layer is deeper than that which occurred during the tests, implying greater predicted mixing.

Fig. 15 shows the concentrations at sensor array K, located just outside the bund wall. The concentration measured by the upper two sensors in both experiments and calculations was approximately zero. At sensor K1, the observed trend was reproduced but the concentrations were again underpredicted. There was a delay before significant amounts of freon-22 were detected, due to the bund wall containing the freon-22 and slowing its dispersion; this is reproduced in the computations. The delay time is, however, underpredicted, which suggests that the calculated slumping velocity is greater than that which occurred during the experiment. No freon-22 was detected at the sensor K2 which was
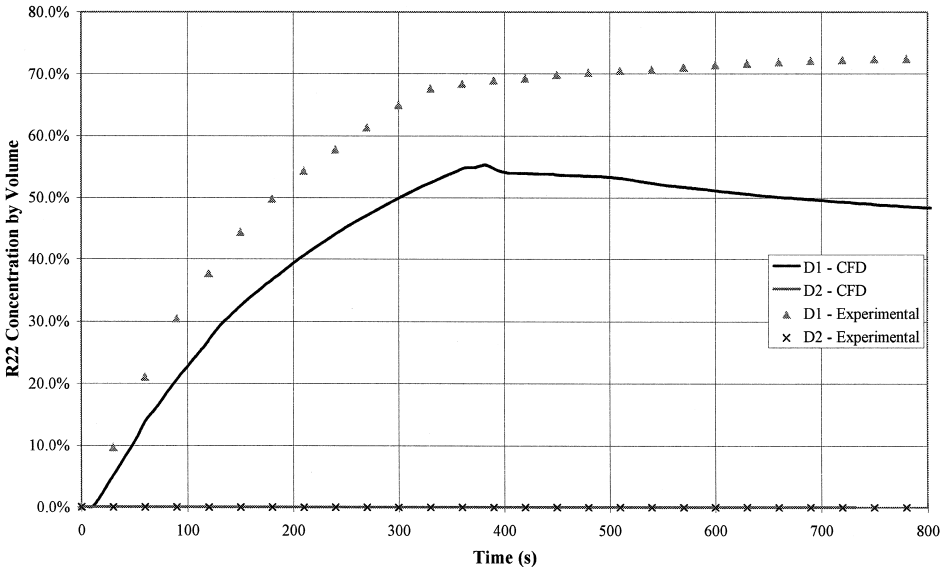

Fig. 14. Concentration time histories at sensor location D (evaporating pool case).
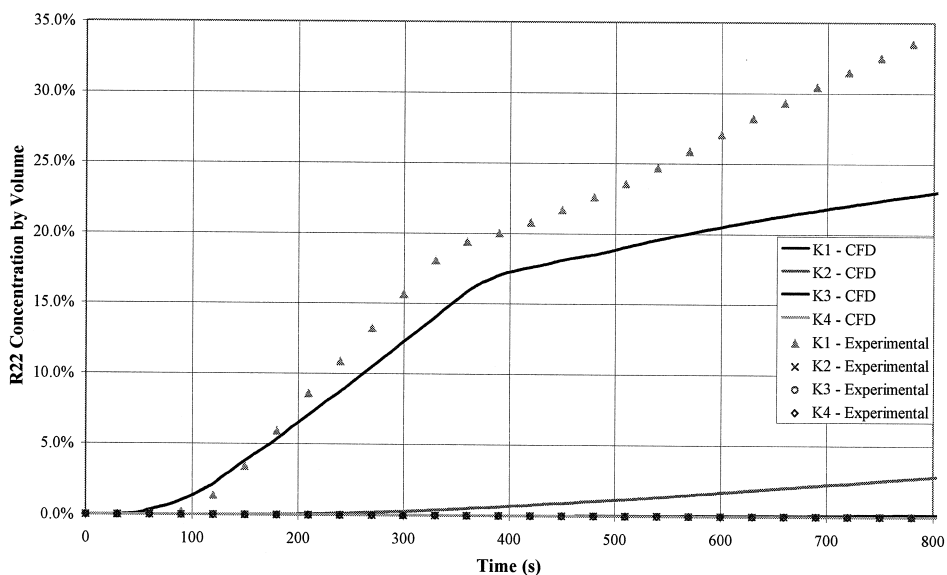
Fig. 15. Concentration time histories at sensor location K (evaporating pool case).

0.830 m above the floor; the computations predicted a small concentration at this location, suggesting that the model over predicts the growth of freon-22 layer.

6.3. Sensitivity studies

In order to identify the causes of the discrepancies between CFD predictions and the experimental measurements, sensitivity studies were carried out for both large jet and pool cases. The calculations described in Sections 6.1 and 6.2 were termed the 'production runs' and provided base cases against which the effects of a number of model refinements were assessed. The results of these sensitivity tests are described below.

6.3.1. Effect of near-wall mesh resolution

One of the main features of the results of the production runs was the underprediction of the room temperatures. For the test room, which was well-sealed, the temperature was principally determined by the balance of the heat transfer through the walls and the energy required to evaporate the liquid fraction of the jet or the pool. The accuracy of the heat transfer predictions at the wall is thus critical in achieving the correct balance. The mesh adopted for the production runs was relatively coarse (see Fig. 7) and employed a uniform mesh node spacing which gave average y^+ values of 208 with a maximum of approximately 1500. Keeping all the model parameters other than the mesh the same as for the production run, the large jet calculation was repeated using a refined mesh of 87,602 cells, which reduced the average y^+ values to 105 with a maximum value of 500 (still outside the ideal range of 12 to 50, but significantly closer).
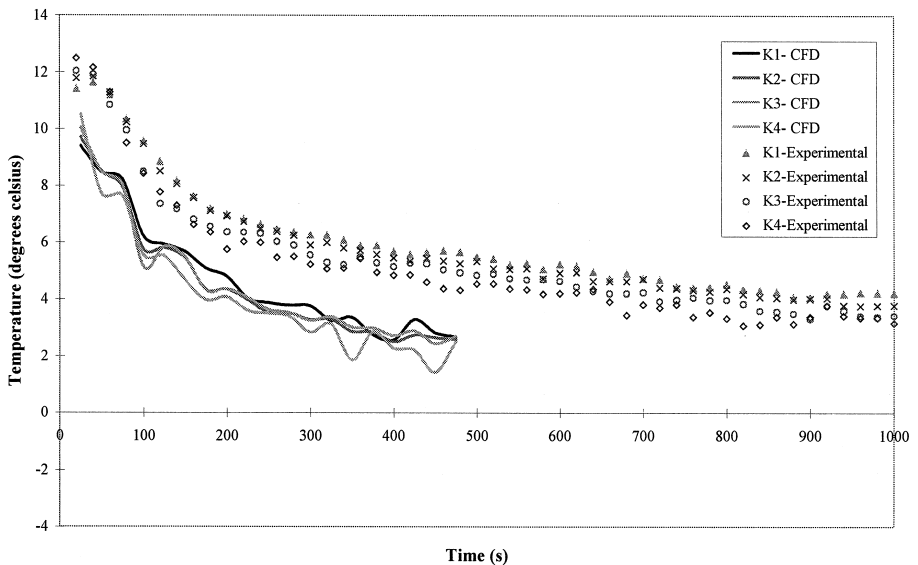
Fig. 16. Temperature time histories at sensor location K (large jet release with near-wall mesh refinement).

Fig. 16 compares the simulated temperatures with those observed in the experimental tests at sensor array K. The predicted temperatures are again underpredicted. However, the near-wall refinement increased the predicted temperatures by approximately 2°C, bringing them much closer to the observed values. The improved temperature prediction had little effect on the concentration predictions; the generally good agreement was maintained.

6.3.2. Effect of heat transfer modelling

Initial calculations, carried out with an adiabatic boundary condition on all the walls and internal structures within the room, showed the predicted room temperatures to fall very rapidly to unrealistically low levels, which highlighted the importance of the thermal boundary conditions. Using isothermal boundary conditions, the predicted temperatures increased to more realistic levels, while for the production runs, the thermal resistance of the wall was introduced.

The heat capacities of the walls and of the objects within the room were not included in the computations; the conduction model assumes a quasi-steady temperature difference. Order-of-magnitude calculations showed that the heat capacity of the internal and external walls would initially provide heat to the room at a comparable rate to that required to evaporate the liquid in the jet. The faster decline of temperatures in the CFD predictions is principally due to ignoring the heat capacity of the walls and objects within the room. It should be noted, however, that this is a transient effect and would not affect the final steady-state temperatures.

An additional pool calculation was also carried out in which the wall boundary conditions were modified to include the effects of free convection, which, since the flow velocities are very low, is likely to be significant. The pool release production run was

repeated with a modified wall boundary condition, in which a mixed free and forced convective heat transfer coefficient was calculated. The wall function was used to determine the forced component and a classical empirical correlation was used to evaluate the free component; the components were then averaged to determine the local heat transfer coefficient. Introducing the revised heat transfer coefficient initially increased the predicted temperatures, but by 200 s, had little effect on the sensor temperatures. The concentration field was essentially unaffected by the modification.

The effect of incorporating the model for natural convection heat transfer was also assessed for the large jet case; its contribution to the overall heat transfer rates was estimated to be approximately 10%. The calculated temperatures in the large jet case were therefore not particularly sensitive to the details of the natural convective heat transfer modelling.

Overprediction of the heat transfer in the region of jet impingement, due to the use of the $k - \varepsilon$ turbulence model (known to overpredict heat transfer rates in impinging jets), was estimated to cause an increase in the overall heat transfer rates of less than 2%.

The results were not found to be particularly sensitive to the factors detailed above, and their effects individually were small. However, cumulatively, the effects of these factors may explain a significant proportion of the differences between the predictions and measurements.

Simple estimates showed that the calculated room temperature was not sensitive to the specified flash fraction, to natural convective heat transfer to the freon-22 delivery flow line outside the room, or to the heat released due to the condensation of water vapour as the room air was cooled by the freon-22 jet.

6.3.3. Effect of turbulence model

In a parallel study, Cowan et al. [10] used both a $k - \varepsilon$ turbulence model and a Reynolds stress turbulence model (RSM) to study the small jet release case. The principal findings of this study were that the turbulence model made little difference qualitatively to the flow predictions; the flow field was dominated by the trajectory of the jet and its impingement on the opposite wall, and the $k - \varepsilon$ turbulence model with standard values for the constants was well tuned for round jets.

The effect of the turbulence model was therefore relatively small for this jet case. However, it is the case that other more stratified flows, or flows controlled by buoyancy, heat transfer or turbulent diffusion, may show much greater differences using different turbulence models. The overprediction of the depth of the freon layer in the pool calculation is likely to be due to an overestimation of the turbulent mixing by the isotropic $k - \varepsilon$ model and an RSM may give improved results for this case. This may also occur where the high momentum of a jet release is destroyed close to the source by impingement.

7. Discussion of results

7.1. Temperature prediction

A feature common to all the production runs was the underprediction of the room temperatures. For the jet releases, it was found that the predictions could be improved

significantly if care was taken in resolving the near-wall flow and heat transfer. The importance of the wall heat transfer was a consequence of the nature of the test; the room itself was well-sealed and the room temperatures thus represented the balance between the rate of heat transfer through the walls and floors against the energy required to evaporate the liquid fraction of the release. This highlighted the importance of the heat transfer calculations and also the need to model accurately the thermodynamics of the jet itself.

The drop in temperature within the room also has an effect on the ventilation, since air is initially drawn into the room as the temperature falls and the air density increases. For the cases studied, the release rate was such that the temperature only fell by a few degrees, suggesting that the heat transfer through the room walls was sufficient to provide the energy required for evaporation. It is possible, however, that as the size of the release increases, the room temperature could fall to the boiling point of the release material. At this point, the evaporation of the liquid fraction of the jet would cease and the vapour release rate would be limited to the flash fraction, suggesting that larger leak rates in a well-sealed room could be heat-transfer-limited. In this case, including the thermal inertia of the equipment within the room may be important.

The modelling of the jet source is also critical in developing an accurate model to account for the situation in which the room temperature falls below the boiling point. For a model to be applied across the full spectrum of potential release rates, feedback between the room temperatures and the vapour release rate is required.

For the pool release, the temperatures were also underpredicted, probably due to the definition of the source conditions and the simulation of the flow immediately around the pool source rather than the modelling of the room wall heat transfer. Inclusion of the natural convective heat transfer model showed little change in results, particularly close to the source. This result suggests that the temperatures close to the source were dominated by the source conditions and mixing immediately around the source. The rate of heat transfer to the freon-22 pool from the structures around it is likely to have a significant effect on the temperatures in the cloud around the pool, though this heat transfer was not measured and is difficult to estimate. The overflow of the vapour from the pool was possibly not sufficiently well-resolved in the simulations, and the consequent entrainment of ambient air was underpredicted. A combination of the two effects is likely to account for the underprediction of the temperatures in this area. Away from the source, subsequent mixing with the ambient air increases the temperature of the freon-22 layer. The calculations captured this behaviour although they predicted a deeper layer than observed in the experiments, probably as a result of an overprediction of the vertical mixing rate.

7.2. Concentration prediction

For the jet releases, the freon-22 concentrations were generally well-predicted despite the inaccuracies in the temperature predictions. The principal reason for this was that for both cases, the flow field was momentum-dominated, giving a well-mixed flow within the room. Buoyancy forces are small and thus, any errors in their predicted magnitude have a minimal effect on the results. The slightly lower predicted concentrations are

likely to be due to an overprediction of the turbulent mixing. The release rate from the room was well-predicted, with predicted concentrations at the low-level sensors closest to the bottom of the extract duct within 10% of the measured concentrations.

For the pool release, the flow was buoyancy-dominated and strongly influenced by the conditions close to the source. The calculations showed larger differences between the computed and measured concentrations, possibly because of the underprediction of the temperature. However, if the rate of contaminant release from the room is considered, the concentration measurements at low-level sensors closest to the extract duct show that the essential features of the flow are well-represented, in particular, the timescales. While this might be surprising when the differences in temperatures are considered, the plume spreading rate is governed by buoyancy flux rather than density difference, so the correct buoyancy flux at the source should give good predictions of the spreading rate.

7.3. Flow prediction

Little quantitative information was available on the flow within the room for any of the releases. The jets were observed to impinge on the far wall of the enclosure and to generate a recirculating, well-mixed flow within the room. Also, a small difference was observed between the trajectories of the two jets, with the smaller jet dipping before impinging on the enclosure wall. Both of these observed characteristics were predicted.

For the pool release, the flow was driven only by the density difference between the cold dense freon-22 introduced into the pool and the air within the room. This evidently resulted in a strongly stratified flow with a high concentration cold layer flowing out from the pool across the floor. The CFD predictions of the pool reproduce this behaviour, at least qualitatively.

8. Conclusions

If an accidental release of toxic material occurs within a building, its effects can be mitigated by the presence of the building. It is possible to predict the extent of this mitigation, and thus to characterise a release, using CFD. This paper presents a number of example applications of CFD which were designed to gain a better understanding of the capabilities, and indeed the limitations, of CFD for this category of application.

In addition to a number of simpler test cases, which were instructive in identifying many of the critical modelling parameters, a series of experiments and CFD calculations was also carried out for a dense toxic release within a model chlorination room. Encompassing a wide range of scales, both geometric and time, and complex physical processes, this presented a considerable modelling challenge.

The simulations carried out, however, showed that CFD can give good predictions of contaminant release rates from buildings. For both jet and evaporating pool releases, good qualitative and quantitative predictions were achieved, although the room temperatures were generally underpredicted. The sensitivity of the results to the modelling details, however, meant that considerable care was required in detailing the model so

that it accurately reflected the physics of the problem. In particular, care was needed in defining the leak source conditions, both fluid dynamically and thermodynamically, and in modelling the near-wall flow and heat transfer.

Each of the source conditions simulated produced different flow fields and mixing regimes within the chlorination room. These complex mixing, heat transfer and thermo-dynamic processes were well-predicted by the CFD and were unlikely to be reproduced in such detail by a less complex model.

Acknowledgements

References

[1] D.M. Deaves, H. Spencer, S. Gilham, Mitigation of dense gas releases in buildings: use of simple models, J. Hazard. Mater. 71 (2000) 129–157.
[2] Computational Dynamics, STAR-CD Version 3.05, Computational Dynamics, 1998.
[3] B.E. Launder, D.B. Spalding, The numerical computation of turbulent flow, Comp. Meth. Appl. Mech. Eng. 3 (1974) 269.
[4] W. Rodi, Influence of buoyancy and rotation on equations for turbulent length scale, Proc 2nd Symp. Turbulent Shear Flows, 1979.
[5] W. Rodi, Experience with two-layer models combining the $k - \varepsilon$ turbulence model with a one-equation model near the wall, AIAA 91-0216, 1991.
[6] E.R. van Driest, On turbulent flow near a wall, J. Aeronaut. Sci. 23 (1956) 1007–1011.
[7] A. Moser, The message of Annex 20: air flow patterns within buildings, 12th AIVC Conference, Air Movement and Ventilation Control Within Buildings, Ottawa, Canada, 1991.
[8] S. Gilham, D.M. Deaves, R.P. Hoxey, C.R. Boon, A. Mercer, Gas build-up within a single building volume — comparison of measurements with both CFD and simple zone modelling, Journal of Hazardous Material 53 (1997) 93–114.
[9] R.J. Bettis, J.T. Allen, Improved validation of source term models: results of experimental study, Health and Safety Laboratory Report IR/L/FR/97/6, 1997.
[10] I.R. Cowan, D.M. Deaves, P. Gallagher, Application of Reynolds stress modelling to gas build-up and external dispersion, HSE Contractor Report AM5096-R01, 1998.