THE ADAPTIVE COUPLING OF COMPUTATIONAL FLUID DYNAMICS WITH WHOLE-BUILDING THERMAL SIMULATION

Ian Beausoleil-Morrison

CANMET Energy Technology Centre, Natural Resources Canada, Ottawa Canada
(ibeausol@nrcan.gc.ca)

ABSTRACT

A adaptive controller was devised and implemented within the ESP-r simulation program to support the conflation of CFD with dynamic whole-building thermal simulation. This controller manages all interactions between the thermal and CFD modelling domains. It incorporates the latest turbulence modelling advancements applicable for room air flow simulation and possesses a suite of handshaking and thermal boundary condition treatments. The controller is based upon a double-pass modelling approach. Each time-step that the thermal domain handshakes with CFD, an investigative simulation is performed to approximate the room's flow and temperature field. Using these estimates the controller determines the nature of the flow (forced, buoyant, mixed, fully turbulent, weakly turbulent) adjacent to each surface. This information is used to select suitable boundary condition treatments for each surface. A second CFD simulation is then performed using the refined modelling approach to more accurately resolve the room's air flow and temperature distribution, and to predict surface convection.

INTRODUCTION

Computational fluid dynamics (CFD) has been widely and successfully applied in the prediction of room air motion for a quarter century. As analysis is usually restricted to single rooms, the modeller must supply boundary conditions in the form of internal wall conditions (surface temperatures or heat flow) and air flows entering or leaving the room. This presents a fundamental dilemma as a room does not exist in isolation. Wall temperatures and air flows through openings are dynamic and dependent on external weather conditions, states prevailing throughout the rest of the building, and the operation of plant equipment, these in turn depending on conditions within the room. This is a significant issue because a CFD solution is really nothing more than the extrapolation of boundary conditions into the domain interior: erroneous boundary conditions invariably lead to erroneous air flow predictions. Therefore, boundary conditions must be dynamically established if CFD is to do more than predict room air flow under very specific design conditions.

The logical solution to this quandary is to integrate CFD modelling with whole-building simulation. The whole-building model can supply CFD with realistic and time-varying boundary conditions. While CFD has the potential to predict the details of flow and temperature fields within particular zones, thus enabling flow visualisation, studies on pollutant dispersion, thermal comfort assessments, and enhanced modelling of convection heat transfer at internal building surfaces. This synergy led Negrão (1995) to integrate a CFD model into the ESP-r building simulation program (ESRU 2000). With this the two modelling domains operate in tandem, "handshaking" on a time-step basis.

Although this work represents a significant advancement, the static coupling between building simulation and CFD is not sufficient: the simulation program must be given the ability to adapt modelling approaches to prevailing conditions. Therefore, building upon the platform provided by Negrão, this paper presents an approach for adaptively coupling CFD with whole-building thermal simulation.

The paper commences with a brief summary of this *adaptive conflation controller*'s building blocks. Issues such as the applicability of CFD, turbulence modelling, and handshaking mechanisms are discussed. However, treatment is succinct due to space limitations. The interested reader is referred to Beausoleil-Morrison (2000) for a full discussion of these issues. The bulk of the paper focuses upon the adaptive conflation controller itself and demonstrates its application. Conclusions are then drawn and recommendations made for future work.

BUILDING BLOCKS

In essence, CFD involves the solution of a set of non-linear partial differential equations using numerical techniques, the equations expressing fundamental physical laws—the conservation of mass, momentum, and energy. Dealing with turbulence (the presence of random fluctuations which exists in most flows of practical interest) complicates matters considerably. Turbulent fluctuations enhance the transport of momentum, heat, and pollutants, and must be considered in the formulation and solution of the equations of motion.

The influence of the turbulent fluctuations on time-mean flow quantities can be approximated by filtering the equations of motion with respect to time (a process known as *Reynolds' averaging*) and by eliminating the high-frequency fluctuating quantities using the eddy viscosity concept. With this, the conservation of momentum, and energy can be expressed as (Rodi 1980),

$$\frac{\partial}{\partial t}(\rho \bar{u}_i) + \frac{\partial}{\partial x_j}(\rho \bar{u}_j \bar{u}_i) = -\frac{\partial \bar{P}}{\partial x_i} + \qquad (1)$$

$$\frac{\partial}{\partial x_j}\left[(\mu + \mu_t)\left(\frac{\partial \bar{u}_i}{\partial x_j} + \frac{\partial \bar{u}_j}{\partial x_i}\right)\right] - \rho\beta(\bar{T}_\infty - \bar{T})g$$

$$\frac{\partial}{\partial t}(\rho c_p \bar{T}) + \frac{\partial}{\partial x_j}(\rho c_p \bar{u}_j \bar{T}) = \qquad (2)$$

$$\frac{\partial}{\partial x_j}\left[\left(k + \frac{c_p \mu_t}{\sigma_t}\right)\frac{\partial \bar{T}}{\partial x_j}\right] + q'''$$

Where x_j are dimensions in the Cartesian coordinate system; $\bar{u}_j$ are time-mean velocity components, t is time; $\bar{P}$ is time-mean pressure; $\bar{T}$ is time-mean temperature; q''' is heat generation; ρ is density; β is the thermal expansion coefficient; g is the gravity vector; c_p is specific heat; and k is conductivity. μ is the molecular viscosity and μ_t is the turbulent or eddy viscosity. σ_t is the turbulent Prandtl number, normally treated as a constant.

Each of these equations (tensor notation is used to collapse three equations into equation 1) contains transient, convection, diffusion, and source terms. These can be thought of as a balance of the forces (or energy flows) that affect a single velocity component (or temperature).

The impact of the turbulent fluctuations upon the mean flow quantities is manifested through the diffusion terms (the terms in square brackets). If $\mu_t/\mu \gg 1$ the turbulent diffusion terms will dominate the molecular terms (μ in equation 1; k in 2). Whereas, if $\mu_t/\mu \approx 0$, molecular effects will prevail. Consequently, the ratio μ_t/μ can be thought of as an indicator of "how turbulent" a flow is locally. Therefore, the distribution of μ_t throughout the flow domain must be established in order to calculate the impact of turbulence. This is the job of the turbulence model.

Numerous turbulence models exist but most applications of CFD for room air flow and heat transfer simulation have employed the standard $k - \varepsilon$ turbulence model with log-law wall functions (Launder and Spalding 1974). This approach was formulated for fully developed turbulent flows, which contrasts with the flow regimes commonly found in rooms. Although there have been many successful *air flow* predictions, the technology has proven deficient at predicting *convective heat transfer* at solid surfaces. Poor surface convection predictions are the result of the inability of the log-law wall functions to resolve the near-wall regions in room air flows (the wall functions assume the form of the velocity and temperature profiles within the boundary layer) And, due to the $k - \varepsilon$ model's overprediction of μ_t in low flow regions.

Numerous alternatives are available but none have proven to be suitable universal replacements for $k - \varepsilon$. One which is particularly interesting for room air flow modelling is the Chen and Xu zero-equation model (1998). New near-wall treatments for the $k - \varepsilon$ model suitable for room air flow prediction are now available, and more are under development. Yuan et al (1993), for example, developed wall functions appropriate for natural convection flows. In contrast to the universal nature of the log-law wall functions however, these are only appropriate for buoyancy-driven flow over vertical surfaces. The Chen and Xu zero-equation model and the Yuan wall functions were incorporated into ESP-r to support the adaptive conflation controller. The code was structured so that either the $k - \varepsilon$ or the zero-equation turbulence model could be invoked at any given time-step. And when the $k - \varepsilon$ is active, structured so that the Yuan wall functions could be applied to some surfaces and the log-law wall functions to others.

An algorithm was also developed to adaptively calculate surface convection coefficients (h_c). Working with a base of 28 h_c correlations, this algorithm (known as the *ACA*) assigns appropriate equations to each internal surface and adapts the selection in response to the room's evolving flow regime. The ACA is described in a companion paper.

The adaptive controller is built upon a handshaking mechanism referred to as *surface conflation*: the thermal and CFD solution domains operate independently but exchange information at the internal surfaces at the beginning of each simulation time-step. The thermal domain establishes boundary conditions for CFD (temperatures or heat fluxes), while the CFD model calculates h_c coefficients for the thermal domain.

A number of boundary condition treatments were devised to support the surface conflation approach. These give the simulator three basic options for resolving the pivot point between the thermal and CFD modelling domains:

1) CFD calculates the air-to-surface heat transfer using an appropriately selected turbulence and near-wall model;
2) the thermal domain calculates the heat transfer using an h_c supplied by the ACA;
3) the CFD and thermal domains cooperatively calculate the heat transfer with support from the ACA.

A Dirichlet boundary condition is used for option 1. With this, the thermal domain prescribes the

turbulence model	handshaking mechanism	CFD thermal boundary condition	applicability
$k-\varepsilon$ model & log-law wall functions (for momentum eqs)	one-way	• Dirichlet • CFD calculates q_{conv} with log-law wall function	• Predicting flow and temp field • Not suitable for buoyancy-driven flow • Not suitable for flows strong affec by q_{conv}
		• Neumann • thermal domain $\longrightarrow T_{surf}$ • thermal domain $\longrightarrow T_{room-air}$ • ACA $\longrightarrow h_c$	• Predicting flow and temp field • Suitable for flows strongly affected by q_{conv}
		• coop Neumann • thermal domain $\longrightarrow T_{surf}$ • CFD $\longrightarrow T_{room-air}$ (avg) • ACA $\longrightarrow h_c$	• Predicting flow and temp field • Suitable for flows strongly affected by q_{conv} • Useful when room stratified
		• coop Robin • thermal domain $\longrightarrow T_{surf}$ • CFD $\longrightarrow T_P$ (local) • ACA $\longrightarrow h_c$	• Predicting flow and temp field • Suitable for flows strongly affected by q_{conv} • Useful when room stratified
	conditional two-way	• Dirichlet • CFD calculates q_{conv} with log-law wall function	• Predicting flow and temp field • Enhancing surf conv calcs • Not suitable for buoyancy-driven flow • Not suitable for flows strong affec by q_{conv} • Next-to-wall points must be properly placed
		• coop Robin • thermal domain $\longrightarrow T_{surf}$ • CFD $\longrightarrow T_P$ (local) • ACA $\longrightarrow h_c$	• Predicting flow and temp field • Enhancing surf conv calcs • Suitable for flows strongly affected by q_{conv} • Useful when room stratified
$k-\varepsilon$ model & Yuan wall functions	one-way	• Dirichlet • CFD calculates q_{conv} with Yuan wall function	• Predicting flow and temp field • Only suitable for buoyancy-driven flow • Only suitable for vertical surfaces
	conditional two-way	• Dirichlet • CFD calculates q_{conv} with Yuan wall function	• Predicting flow and temp field • Enhancing surf conv calcs • Only suitable for buoyancy-driven flow • Only suitable for vertical surfaces
Chen & Xu 0-eqn model	one-way	• Dirichlet • CFD calculates q_{conv}	• Predicting flow and temp field • Suitable for quick indication of flow • Less suitable for buoyancy-driven flow • Less suitable for flows strong affec by q_{conv}
	conditional two-way	• Dirichlet • CFD calculates q_{conv}	• Predicting flow and temp field • Enhancing surf conv calcs • Less suitable for buoyancy-driven flow • Less suitable for flows strong affec by q_{conv} • Next-to-wall points must be properly placed

Table 1: Turbulence, handshaking, and boundary condition options for the adaptive conflation controller

temperature of the internal surface and CFD resolves the boundary layer using its calculated temperature, velocity, and turbulence fields in conjunction with a near-wall treatment. A Neumann boundary condition is used for option 2. With this, the surface convection heat transfer is imposed on the CFD domain through the thermal domain's calculated wall and room air temperatures. Two methods were devised to support option 3. One uses a Neumann boundary condition while the other a Robin boundary condition. The former uses the same approach as for option 2 except that the CFD domain supplies the room air temperature (mass-averaged from the previous solver iteration). The latter makes use of local h_c values and the air temperatures predicted at the next-to-wall grid points.

THE ADAPTIVE CONFLATION CONTROLLER

The previous section laid the foundation for dynamically controlling the conflation between the thermal and CFD modelling domains. Supported by the ACA, a suite of handshaking schemes, boundary condition treatments, and turbulence models have been implemented to provide optional approaches for resolving the pivot point between the two modelling domains.

Upon examining a number of alternatives, ten viable combinations of these approaches were selected for the adaptive conflation controller. These are summarised in Table 1. The $k-\varepsilon$ model with log-law wall functions is used in six of the schemes, although four of these do not employ the thermal wall function. Two schemes use the $k-\varepsilon$ model with the Yuan wall functions and two use the Chen and Xu zero-equation model.

The last column of the table describes, in broad terms, the applicability and limitations of each scheme. As can be seen, a number of the schemes are useful for predicting the flow and temperature fields, while others can perform this function as well as calculate surface convection for the thermal domain. Some schemes are restricted to certain flow regimes (e.g. buoyancy-driven flow) while others are more generally applicable. Three of the schemes are particularly useful for dealing with stratified rooms.

Control schemes

Given the reality that thermal conditions outside the CFD domain are dynamic, the schemes laid out in

Table 1 can not be fully exploited unless the simulation program is supported by a controller that monitors thermal and air flow conditions and dynamically select an appropriate scheme for the prevailing conditions.

A plethora of structural and logical options exist for devising the controller. Two viable solutions were devised and implemented to demonstrate the concept but it should be noted that no attempt has yet been made to optimize an approach. One solution is referred to as the *one-way adaptive conflation controller*. It is suitable when the objective of the analysis is to predict the flow and temperature field within the room, perhaps for the purposes of visualizing the flow, assessing thermal comfort, or studying pollutant dispersion. The second solution is called the *conditional two-way adaptive conflation controller*. It is suited to performing these tasks as well as enhancing the thermal simulation by providing surface convection estimates for the thermal domain. The latter is described here. The structure of the one-way adaptive conflation controller is very similar with one notable exception. With it, the link from CFD to the thermal domain is broken to protect the integrity of the thermal simulation.

Figure 1 illustrates the logic flow. This controller is based on the surface conflation method and makes use of a double-pass modelling approach. An investigative CFD simulation is first performed to approximate the room's flow and temperature field. These results are appraised by the controller to determine the nature of the flow. Boundary condition treatments and other simulation parameters appropriate for the prevailing flow conditions are selected, and a second CFD simulation performed using the refined modelling approach. In other words, two CFD simulations are performed between the time-steps of the thermal domain. In this implementation, the zero-equation model is employed in the investigative analysis, and the $k-\varepsilon$ model is used in the refined CFD simulation. There is an appraisal stage following convergence of the refined CFD simulation, prior to h_c coefficients being returned to the thermal domain. These principle steps are outlined below.

Investigative CFD simulation

The purpose of the investigative simulation is not to accurately quantify the flow and temperature field, but rather to approximate the flow regime in order to establish appropriate modelling approaches. The zero-equation model is a good candidate for this task because of its stability and computational efficiency. Additionally because the zero-equation model gives realistic μ_t predictions, it is a useful tool for characterizing the nature of turbulence throughout the room.

An approximate solution is adequate for the assessment. Therefore, the controller moderates the

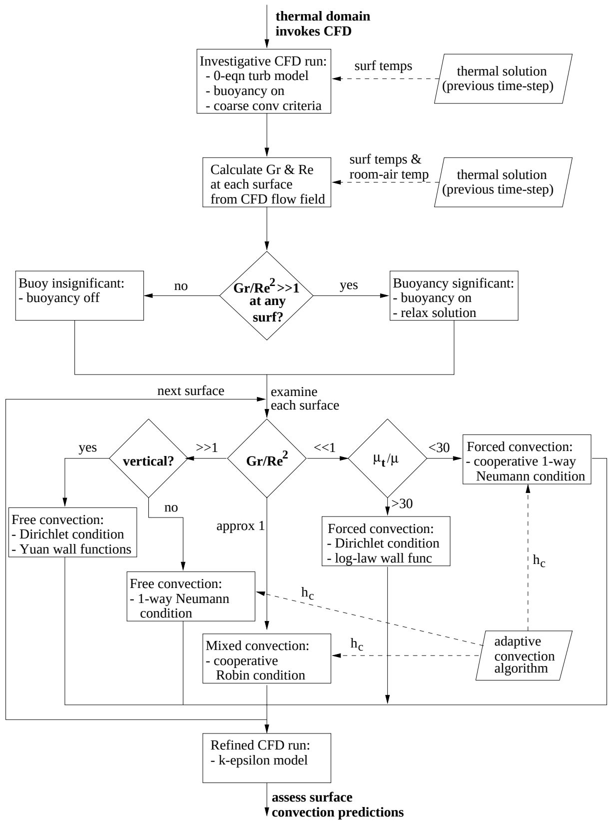

Figure: 1 Conditional two-way adaptive conflation control

convergence criteria so that CFD may resolve the flow pattern with fewer solver iterations. Once the investigative simulation converges, the controller appraises the predicted flow and temperature field by calculating relevant dimensional groupings at each surface.

Dimensionless groupings

The Grashof and Reynolds numbers were selected as the relevant physical groupings. The Grashof number indicates the ratio of the buoyancy force to the viscous force acting on the fluid, and thus measures "how buoyant" the flow is adjacent to the surface. The Reynolds number serves a similar role in gauging "how forced" the flow is adjacent to a surface, as it indicates the ratio of the inertial to viscous forces. These dimensional groupings are given by,

$$Gr_L = \frac{g\beta\rho^2 \,|\, T_{wall} - T_{room-air} \,|\, L^3}{\mu^2} \tag{3}$$

$$Re_L = \frac{\rho V_{ref} L}{\mu} \tag{4}$$

The length scale (L) is automatically calculated by the adaptive conflation controller when the problem's geometry is read by the simulation program, prior to the time-step simulation. It is equal to the height of vertical surfaces and the ratio of area to perimeter for horizontal surfaces. The surface

(T_{wall}) and room-air ($T_{room-air}$) temperatures in equation 3 are taken from the thermal domain's solution. The reference velocity (V_{ref}) in equation 4 is determined from the investigative CFD solution. It is equal to the greatest velocity predicted at any of the CFD grid points adjacent to the surface.

The nature of the flow in the vicinity of each surface is then determined. When $Gr_L/Re_L^2 \ll 1$ it is concluded that free convection is overwhelmed by forced convection effects. When $Gr_L/Re_L^2 \gg 1$ free convection effects dominate. And when Gr_L and Re_L^2 are of the same order of magnitude, both forced and free effects are significant.

Refining simulation parameters

If buoyant effects are found to be important at any surface, the controller adjusts the CFD model to consider the buoyancy term in the z-momentum equation (the last term in equation 1); otherwise, this buoyancy term is dropped. Including buoyancy tends to destabilize the CFD iterative solution procedure because it provides a stronger coupling between the sequential solutions of the momentum and energy equations. Consequently, when buoyancy is considered, the controller automatically reduces the relaxation factors for the velocity, pressure, temperature, and turbulence variables to encourage convergence.

The controller then establishes an appropriate boundary condition for each surface (independently). If the Gr_L/Re_L^2 ratio indicates that free convection effects are dominant at a vertical surface, then the two-way Dirichlet condition with the Yuan wall functions is used. Since no technique currently exists to satisfactorily resolve the heat transfer for free convection at horizontal surfaces, the one-way Neumann condition is applied and the ACA calculates h_c using an appropriate correlation. When both free and forced effects are important, the two-way cooperative Robin condition is employed, this to ensure local temperature effects are considered. Again, the ACA supplies h_c.

When forced effects dominate the controller examines the μ_t/μ ratios of the next-to-wall grid points to gauge "how turbulent" the flow is adjacent to each surface. Based upon the guidance provided by Baker et al (1994), if the average μ_t/μ ratio at these points is less than 30 it concludes that the flow is (locally) weakly turbulent. In such conditions the $k-\varepsilon$ model cannot accurately predict the surface convection using the next-to-wall grid point temperatures, irregardless of the near-wall treatment (Beausoleil-Morrison and Clarke 1998). Therefore to protect the integrity of the thermal domain, the controller applies the cooperative one-way Neumann condition. And if the average μ_t/μ ratio of the grid points adjacent to the surface is greater than 30, the conditional two-way Dirichlet condition with log-law wall functions is used.

This completes the preparation for the second CFD simulation. The CFD model is then invoked by the controller to refine the flow and temperature fields. The approximate flow, temperature, and μ_t fields predicted by the investigative run are used to initialize the solution variables to accelerate convergence, and to provide greater numerical stability.

Once the refined CFD solution converges the controller generates images of the flow, then returns control to the thermal domain. When the thermal domain calls upon the adaptive conflation controller at next time-step, the process elaborated above is repeated.

Assessing the CFD-predicted h_c

Where the Dirichlet and Robin boundary conditions are used, the CFD domain returns h_c coefficients to the thermal domain. As previously discussed, the accurate calculation of surface convection is a highly problematic task for CFD. Therefore, to avoid distorting the thermal domain's heat balance equations, the controller employs a two-stage screening process to assess (and where necessary rejects) the CFD-predicted h_c coefficients. In this way, the h_c coefficients predicted at some surfaces may be accepted for use by the thermal domain, while those at other surfaces are rejected.

The $k-\varepsilon$ model with log-law wall functions can lead to realistic surface convection predictions when the next-to-wall grid points are appropriately placed. There is no consensus on how the gridding should be placed, but based on the recommendations of Niu and van der Kooi (1992), Chen (1995), Schild (1997), and Awbi (1998), the controller considers the gridding to be appropriate when the dimensionless spacing to the next-to-wall grid point (y^+) is in the range of 10 to 30. Therefore, following convergence of the refined CFD simulation, the y^+ values of the next-to-wall grid points are calculated. If they lie outside this range, the CFD h_c predictions at these surfaces are rejected.

As a second level of screening, the CFD h_c predictions are compared to those from the ACA. The CFD results are rejected if they are not within the same range as the empirical values: $\frac{1}{5} \cdot h_{c,ACA} \leq h_{c,CFD} \leq 5 \cdot h_{c,ACA}$ where $h_{c,ACA}$ and $h_{c,CFD}$ are the convection coefficients predicted by the ACA and CFD, respectively.

EXAMPLE APPLICATION

A model of a hypothetical office building was created to demonstrate and test the operation of the adaptive conflation controller. This example makes use of the one-way handshaking control logic. The building is conditioned with a constant volume HVAC system that varies its supply air temperature in response to each office's heating and cooling demands. A single office of the building is illustrated in Figure 2. The locations of the HVAC system's supply air diffuser and extract can be seen. Heated or cooled air is supplied to the room as a horizontal jet, the diffuser being located high on the wall opposite the window. Return air is extracted at the same wall, but at a location near the floor. The goal of the system is to provide an even distribution of air and to extract pollutants from the room.

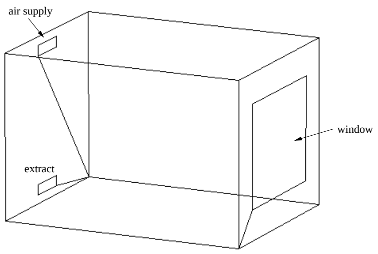

Figure 2: ESP-r model of office

A CFD domain was established for a single office. A coarse mesh of 12x10x12 grids was used, this being sufficient to characterize the general flow pattern. Ottawa weather data was used in the simulation. The period from 10h00 to 13h00 on January 9 was chosen for the analysis as this provides the controller with an interesting and challenging range of operational conditions over a short period of time. Cold temperatures (-10°C to -20 °C) and high solar radiation (up to 900 W/m^2) result in a rapid warming of the room's internal surfaces and a quickly changing load profile over the analysis period.

The HVAC system maintains the zone air temperature at a constant 20°C. Throughout most of the morning the internal surfaces are colder than the room air. This results in negative buoyancy forces adjacent to the surfaces. By about 11h00, however, the surfaces have warmed to above the room air temperature due to increasing solar insolation, this reversing the direction of the buoyancy forces adjacent to the surfaces.

The loads placed on the HVAC system vary considerably over this period. At 10h00, for example, the zone requires heating to maintain the set-point temperature. As a result, the HVAC system supplies air at about 25°C. The heating load diminishes with time. In fact, by 11h30 the internal and solar gains balance the heat losses so the system supplies air at the room temperature. And after this time, the HVAC system must supply cool air to the zone to extract excess heat. The temperature difference between the supply air and the room air produces a buoyancy force on the incoming jet. The jet will tend to rise when warm air is supplied, and will tend to detach from the ceiling and drop when it is cool. Therefore, the direction of the buoyancy force changes quite suddenly between 11h00 and 12h00.

Adapting modelling to the flow

Simulations were performed using a 10-minute time-step in order to capture the impact of the rapidly varying conditions. Following a short start-up period to condition the model, the thermal domain invoked the adaptive conflation controller on January 9 at 10h00. The surface temperatures predicted by the thermal domain were passed to the CFD domain as boundary conditions. Likewise, the supply air temperature calculated by the the thermal domain was mapped to the CFD domain's air inlet.

At this point, the controller performs an investigative CFD run with the zero-equation model using coarse convergence criteria. A rapid estimate of the flow field results (illustrated in Figure 3).

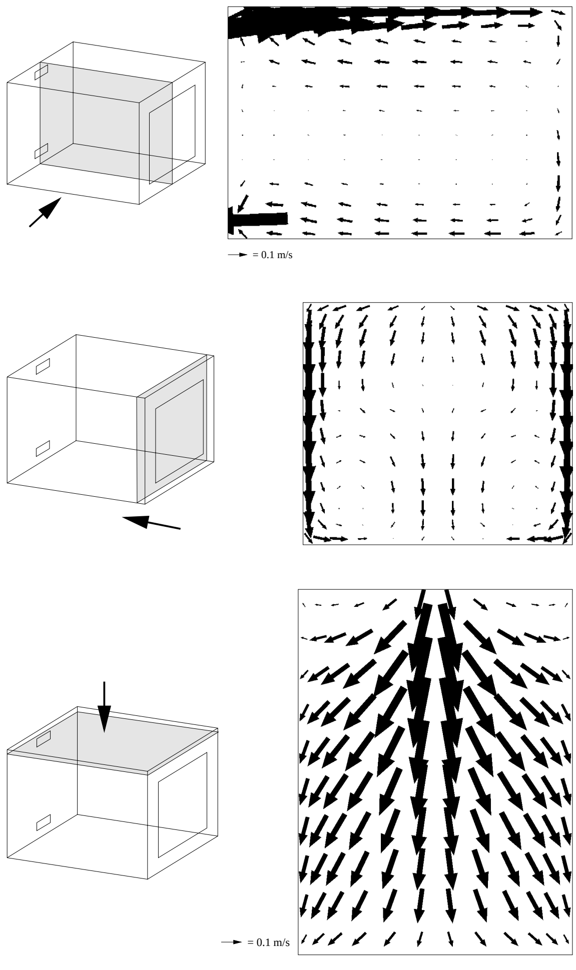

Figure 3: Flow predicted by investigative CFD run at 10h00

Three views of the flow are displayed. There is a strong flow along the ceiling as can be seen in the first and third views. The incoming jet spreads across the ceiling and reaches the opposite wall (the one with the window). The jet's warm temperature has helped it to adhere to the ceiling. The second view indicates there is also strong flow down the side walls. The impact of the buoyant forces at the window can be seen in the this image. In this case the buoyant force assists the mechanically driven jet. The window surface is 13.5°C whereas the surrounding wall is 17.3°C. This results in a greater buoyancy force adjacent to the window, and thus greater velocity. Flow is quite diffuse and even across the floor (view not shown).

The controller then assesses the flow and temperature field. It calculates the Grashof and Reynolds numbers for each surface based on the temperatures and velocities solved at the next-to-surface grid points. By comparing these dimensionless groupings the controller determines that flow at all vertical surfaces is primarily buoyant ($Gr_L/Re_L^2 \approx 10-20$). As a result, it selects the Dirichlet boundary condition with the Yuan wall functions for these surfaces. As Gr_L and Re_L^2 were approximately the same at the ceiling and floor, the controller concludes that flow at these surfaces is mixed. Therefore, the cooperative Robin boundary condition is employed at the horizontal surfaces (the ACA supplies h_c coefficients).

Next, to resolve the flow in greater detail the controller adjusts the convergence criteria, turns buoyancy on, reduces the relaxation factors, then invokes the CFD solver using the $k-\varepsilon$ model. The flow and temperature fields predicted by the investigative run are used as the initial conditions for the refined run. This was found to greatly enhance the stability of the calculations. Indeed, attempts to resolve this flow with the $k-\varepsilon$ model without these initial conditions invariably led to divergence.

Control returns to the thermal domain following the convergence of the refined CFD run at 10h00. And once the thermal domain simulates the thermal state for the 10h10 time-step, it passes control back to the adaptive conflation controller. Another investigation CFD simulation is performed and its flow predictions assessed. At 10h10, due to the lower supply air temperature and higher surface temperatures, the Gr_L/Re_L^2 ratios are lower at all surfaces. As a result, the controller retains the boundary conditions at the floor, ceiling, front wall, back wall, and window, but adjusts the treatment at the side walls. Here, the Yuan wall functions are replaced with the cooperative Robin condition as flow here is now considered mixed.

Predicting the evolution of the flow

The process of investigative run, flow assessment, and boundary condition adjustment is repeated each time-step to predict the flow in response to the changing thermal conditions. Figure 4 illustrates the evolving flow in the room using a side view at the cross-section of the diffuser. Note that the time space between views is not even, this to focus attention on the period (11h00 to 12h00) during which changes occur most drastically.

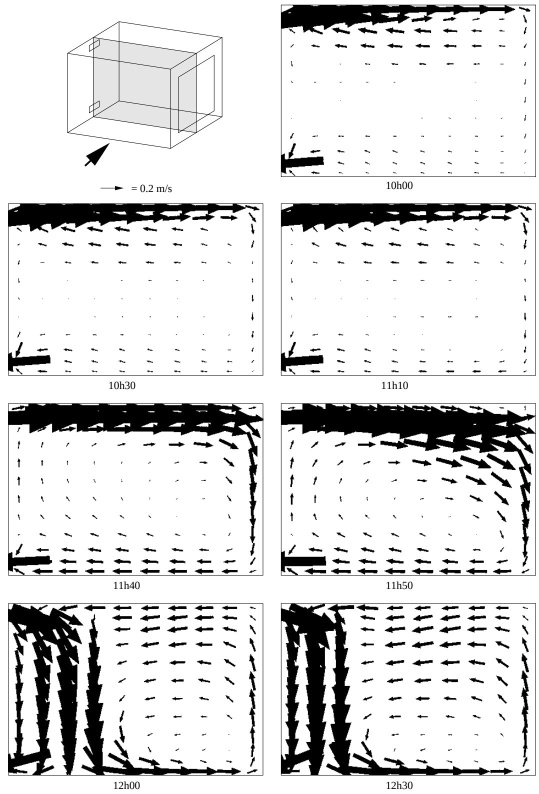

Figure 4: Evolving flow pattern

The impact of the buoyancy of the incoming jet can be easily seen in this figure. The supply air is 5°C warmer than the room air at 10h00, and thus rises to the ceiling. The jet does not rise as much at 10h30 and 11h10, when the supply air is only 2 to 2.5°C warmer than the room air. The supply air temperature drops (slightly) below the room air temperature at 11h40. The jet still penetrates to the back wall at 11h50, but is clearly detaching from the ceiling. Finally, at 12h00, the buoyancy of the supply air (now 3.5°C cooler than the room air) causes the jet to completely detach from the ceiling. Flow spills to the floor, where it separates, causing a recirculation region in the right half of the room.

CONCLUDING REMARKS

This paper has described the *adaptive conflation controller* which has been implemented in the ESP-r program to support the coupling of CFD with whole-building thermal simulation. The controller incorporates the latest turbulence modelling advancements applicable for room air flow simulation and possesses a suite of handshaking and thermal boundary condition treatments. The job of this controller is to monitor the evolving thermal and air flow conditions in the room and dynamically select an appropriate combination of modelling approaches for the prevailing conditions.

The controller makes use of a double-pass modelling approach. Each time-step that the thermal domain handshakes with CFD, the controller performs an investigative simulation to approximate the room's flow and temperature field. Using these estimates, the controller calculates dimensionless groupings to determine the nature of the flow (forced, buoyant, mixed, fully turbulent, weakly turbulent) adjacent to each internal surface. This information is used to select suitable boundary condition treatments for each surface. A second CFD simulation is then performed using the refined modelling approach to more accurately resolve the room's air flow and temperature distribution, and to predict surface convection. In order to protect the thermal domain, a two-stage screening process is used to assess (and where necessary reject) the CFD-predicted surface convection estimates.

An example was presented to demonstrate the controller. The objective of this analysis was to examine an HVAC system's performance at ventilating an office. Air motion in the room was influenced not only by the HVAC system's flow rate, but also by buoyancy forces resulting from temperature differences between the supply air, internal surfaces, and the room air. This example illustrated how the adaptive conflation controller can resolve the impact of realistic (dynamic) operating conditions upon room air flow. Appropriate boundary conditions were selected for the CFD domain in response to the local flow conditions, and this selection evolved with the flow. This case was a challenging test for the controller. The mechanical forces of the jet were in close balance with buoyancy, and this balance altered over a relatively short period of time. Notwithstanding, the controller was able to resolve the evolving flow in the room and provide information on the ventilation system's effectiveness.

Two possible control schemes have been developed for the adaptive conflation controller. As these were based strictly on intuition and experience, there is much room for refinement and optimization. In particular, the criteria used to assess the CFD-predicted surface convection estimates should be examined, critically reviewed, and optimized. Although the investigative CFD simulation was found to add little computational burden, savings could be realized by using a coarser grid or by initiating the investigative run only when boundary conditions change significantly between time-steps. New or different criteria could be used to establish boundary conditions for the refined CFD simulation. Suggestions include: measuring stratification; sensing the operational state of plant components; sensing flow through network air flow connections; and using different dimensionless groupings to assess the nature of the flow. Adapting the convergence criteria to the objective of the analysis could realize significant computational savings. As well, operating on the flow field predictions of the investigative CFD run, a grid tuning algorithm could be used to optimally locate the next-to-wall grid points for the purposes of predicting heat transfer with the log-law wall functions.

REFERENCES

Awbi H.B. (1998), 'Calculation of Convective Heat Transfer Coefficients of Room Surfaces for Natural Convection', *Energy and Buildings*, 28 219-227.

Baker A.J., Williams P.T., and Kelso R.M. (1994), 'Development of a Robust Finite Element CFD Procedure for Predicting Indoor Room Air Motion', *Building and Environment*, 29 261-273.

Beausoleil-Morrison I. and Clarke J.A. (1998), 'The Implications of Using the Standard $k-\varepsilon$ Turbulence Model to Simulate Room Air Flows which are not Fully Turbulent', *Proc. ROOMVENT '98*, (1) 93-100, Stockholm Sweden.

Beausoleil-Morrison I. (2000), *The Adaptive Coupling of Heat and Air Flow Modelling within Dynamic Whole-Building Simulation*, PhD Thesis, University of Strathclyde, Glasgow UK.

Chen Q. (1995), 'Comparison of Different $k-\varepsilon$ Models for Indoor Air Flow Computations', *Numerical Heat Transfer*, B(28) 353-369.

Chen Q. and Xu W. (1998), 'A Zero-Equation Turbulence Model for Indoor Airflow Simulation', *Energy and Buildings*, 28 137-144.

ESRU (2000), *The ESP-r System for Building Energy Simulations: User Guide Version 9 Series*, University of Strathclyde, Glasgow UK.

Launder B.E. and Spalding D.B. (1974), 'The Numerical Computation of Turbulent Flows', *Computer Methods in Applied Mechanics and Engineering*, 3 269-289.

Negrão C.O.R. (1995), *Conflation of Computational Fluid Dynamics and Building Thermal Simulation*, PhD Thesis, University of Strathclyde, Glasgow UK.

Niu J. and van der Kooi J. (1992), 'Grid-Optimization for $k-\varepsilon$ Turbulence Model Simulation of Natural Convection in Rooms', *Proc. ROOMVENT '92*, (1) 207-223, Aalborg Denmark.

Rodi W. (1980), *Turbulence Models and their Applications in Hydraulics—A State of the Art Review*, International Association for Hydraulic Research, Delft The Netherlands.

Schild P. (1997), *Accurate Prediction of Indoor Climate in Glazed Enclosures*, PhD Thesis, Norwegian University of Science and Technology, Trondheim Norway.

Yuan X., Moser A. and Suter P. (1993), 'Wall Functions for Numerical Simulation of Turbulent Natural Convection Along Vertical Plates', *Int. J. Heat Mass Transfer*, 36 (18) 4477-4485.