An Investigation Of The Airflow In Mushroom Growing Structures, The Development Of An Improved, Three-Dimensional Solution Technique For Fluid Flow And Its Evaluation For The Modelling Of Mushroom Growing Structures

James J. Grant

Diploma (Hons.) in Applied Physics

Submitted in Partial Fulfilment for the Degree of Doctor of Philosophy

to

The School of Mathematical Sciences,
Dublin City University

Thesis Supervisors:
Dr. D. Williams
Prof. A. Wood

September 2002

REFERENCE

I hereby certify that this material, which I now submit for assessment on the programme of study leading to the award of PhD, is entirely my own work and has not been taken from the work of others save and to the extent that such work has been cited and acknowledged within the text of my work

Signed: _____ (Candidate)

ID No.: _____

Date: _____

This thesis is dedicated, with love, to my wife, Siobhán. Her support is deeply appreciated.

Acknowledgements

Contents

Abstract

This thesis is an examination of the airflows in mushroom growing rooms. An experimental investigation of the nature of the flows in Irish tunnels showed them to be of low magnitude at the crop but controllable in principle for single layer growing. It was found that stratification of the airflow in growing tunnels could cause severe reductions in cropping surface airspeed and the operation of the heating system was identified as the main source of this. An alternative air distribution system was shown to have the potential to overcome the effects of heating. Airflow for three level growing systems in tunnels was found to be non-uniform and the use of wall-mounted deflecting plates was shown to have the potential to correct this.

The provision of air flow solutions for the wide range of new growing systems would be difficult using empirical methods alone and therefore a modelling approach was sought to complement and aid the experimental work.

The initial modelling work was carried out in two dimensions with TEACH-T code (SIMPLE flow solver) to calculate the turbulent flow. The code was extended to three dimensions because it was not possible to model usefully in a two-dimensional approximation.

Convergence times for the SIMPLE solver were found to be excessively long. Trial applications of multi-level acceleration produced approximately 15% savings in computational effort so a new solver was investigated. The CELS (Coupled Equation Line Solver) method had been reported as superior to SIMPLE in two dimensions and already has a multi-level technique to accelerate convergence, i.e. Additive Correction Multigrid (ACM).

CELS was first applied in two dimensions in order to test its usefulness with the turbulence model in the equation set. Improvements in the time to convergence, relative to SIMPLE, justified its extension to three dimensions. The Additive Correction Multigrid technique also produced significant improvements and this was extended to three dimensions.

CELS3D is essentially a plane solver applied to a three-dimensional grid and a number of procedures for its application were investigated. All produced savings relative to the SIMPLE solver. The QUICK differencing scheme was incorporated in the TEACH-based code and CELS3D was tested with various geometries and values of the Reynolds number. The best results gave a 79% reduction in the time to convergence of the solver. The ACM technique in three dimensions was investigated but no useful savings in computational effort were made.

In the application to mushroom growing structures, the principles of the application of CELS3D to flows around obstructions in the flow domain were examined and the difficulties identified. A solution was found but its implementation proved impractical for all but the simplest cases.

Nomenclature

A	area or coefficient of discretised equations
a	coefficient of discretised equations
A	coefficient matrix
b	source term for discretised equations
B	source term for CELS equations
C	coefficient of pressure term
c_1, c_2, c_D, c_μ	constants of turbulence model
D	term in pentadiagonal solution formulation, conductance
E	wall function constant
e	error vector
F	term in pentadiagonal solution formulation, flow rate
G	generation term in k equation
I	linear interpolation vector
i,j,k,l,m,n	indices
k	turbulence energy
L	length scale
p	pressure
r	residual vector
S	source term
t	time
u,v,w	velocity components
u	solution vector
v	approximation to solution vector
x,y,z	Cartesian co-ordinates
α	residual reduction factor
Γ	diffusion coefficient
ε	dissipation rate
ϕ	general dependent variable
λ	molecular diffusivity
κ	wall function constant
ν	molecular/kinematic viscosity
ν_t	turbulent or eddy viscosity
ρ	density
σ	diffusion constant

Chapter 1

Introduction

1.1   Background

The work presented in this thesis was prompted by the study of airflows in mushroom growing structures. Mushrooms, as a crop, do not photosynthesise and have no specialised fluid transport system equivalent to the green plants vascular system. Water is transported by capillary action between the cells and fungal strands that make up the organism and osmotically by the cells themselves and, therefore, in order to gain control over the growth of the organism it is necessary to gain control of the evaporative conditions at the cropping surface, i.e. the crop micro-climate. Control of microclimate for mushrooms means manipulation of the evaporating power of the air that is usually defined as the product of the vapour pressure deficit and the air speed (Edwards, 1973). Thus, achieving correct airflow across the developing crop is vital to the success of the production process.

The mushroom crop, during a ten week production cycle, moves through a number of phases of development. In the first two phases (three to four weeks) of vegetative development where strands of fungal mycelium colonise the compost that supplies its nutrients and water, air speed at the crop is not a critical quantity. Air at this time functions largely as a medium for the removal of heat produced by the metabolic reactions in the mycelium. The fruiting process (reproduction) that produces the mushrooms is initiated by applying a shock to the crop (reduced temperature, humidity and carbon dioxide concentration) and then evaporation, and hence air speed, becomes extremely important. Essentially, a period of high evaporation for two to three days

induces the formation of small fruiting structures, called pinheads, and then evaporation is reduced but tightly-controlled while the pinheads develop over a period of one week into full-sized mushrooms.

During this last period there is a relatively narrow band of evaporating power within which the micro-climate must be maintained. If the rate of evaporation is too low then conditions on and around the mushroom cap are too moist and they favour the development of bacterial populations that produce browning (bacterial blotch) and pitting of the cap surface. At the other side of the band, if the evaporation rate is too high then surface tissues rupture under water stress and then dry and produce a flaky appearance called scaling. This browns quickly in storage.

Remaining within the required band of evaporation rate produces a bright white mushroom with an unblemished surface. This high quality produce is in demand among large multiple retail outlets and they will pay a premium price for the product. In Ireland, the profitability of a mushroom growing enterprise depends on producing a high percentage of top grade mushrooms and the competitiveness of the industry as a whole has been built on the quality of the product.

1.2   Growing structures for mushroom production

In achieving the correct evaporating conditions, controlling vapour pressure deficit is relatively easy as moisture can be added to air by injecting steam or fine mist or removed by passing air over a chilled heat exchanger. Setting the correct air speed is more difficult and, in particular, ensuring uniformity of conditions over the entire surface of a crop demands that airflow be well understood for the various cropping structures employed commercially.

There are a number of different commercial mushroom-growing systems in Europe and in the United States. All the production is carried out indoors and there are a variety of growing systems and associated structures. Some growing rooms are approximately

2

square or rectangular in section and others are curved, polyethylene-covered tunnels with a variety of cross-section shapes that can deviate markedly from a semi-circle.

The square geometries are used traditionally for multi-tier growing and these pose the greatest difficulties in the provision of uniform air flow at all points on a cropping surface. For example, an American growing room is shown schematically in figure 1.1 and the general airflow directions are indicated (Lomax, 1993). Air is delivered at the top of the building through apertures in the sides of steel ducting and it then flows across the top of the structure and down the side. The crop is grown on a tiered system of supports for the substrate that are called shelves or trays and the cropping takes place on the upper surface of these shelves. The efficiency of the production process depends to a great extent on the air flow through the tiers but these structures have posed a great deal of difficulty in trying to optimise air speed and achieve uniform air flows. Making distributed measurements in these situations is difficult and time-consuming and an airflow calculation approach could yield insights and solutions as effectively as a more empirical approach.
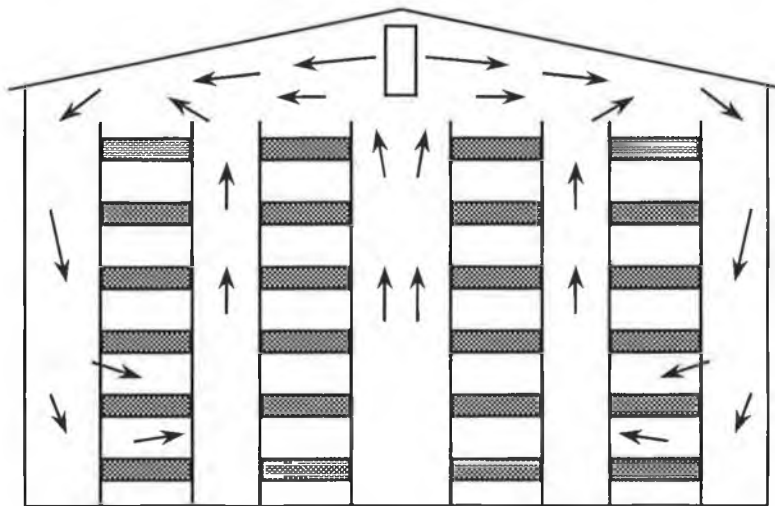
Figure 1.1. Schematic of an American mushroom growing room.

The structures and the airflow in Irish growing tunnels are simpler in principle but there are features that would benefit from a modelling approach. Figure 1.2 is a schematic diagram of the type of structure used in mushroom growing in Ireland. It is a double-clad plastic tunnel with insulation between the two layers of polyethylene and the markings on the end wall represent the relative size of bags of mushroom compost as used in these tunnels. The floor is usually covered with these with allowance made for walkways.

The interior of an empty mushroom growing tunnel is pictured in figure 1.3. Air is delivered by means of a polyethylene duct that is designed for uniform output along its length and the general airflow pattern is shown in figure 1.4. The air jets from the duct strike the tunnel wall and then travel along it until the flow reaches the crop at the top surface of the bags. Air then moves across the cropping surface and begins rising to be entrained in the jets. Measurement of this air flow shows it to be a rapidly varying quantity and of low magnitude (chapter 2). Typical air speeds just above the cropping surface range from 10 to 30 cms^{-1} and this low velocity makes the flow difficult to measure in the commercial situation.
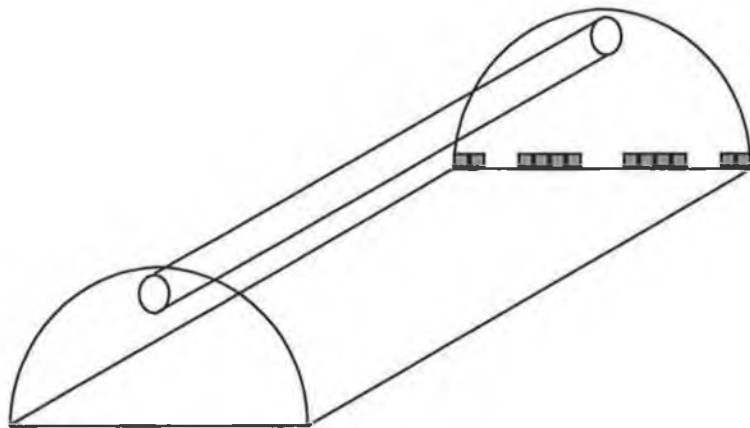
Figure 1.2. Schematic of an Irish mushroom-growing tunnel.

Figure 3. The interior of an Irish mushroom growing tunnel.

A model would allow experimentation with possible means of improving uniformity of flows and promising results could then be applied to the real situation. In this way, the model could be used to guide practical experiments and measurement effort could be reduced to validating the model rather than examining each new situation empirically.

Figure 1.5 shows the variety of tunnel shapes that can be encountered in practice. The effect of these variations in geometry could be studied with a suitable model. The tunnels shapes are not simple, i.e. do not fit neatly into a cylindrical co-ordinate system.
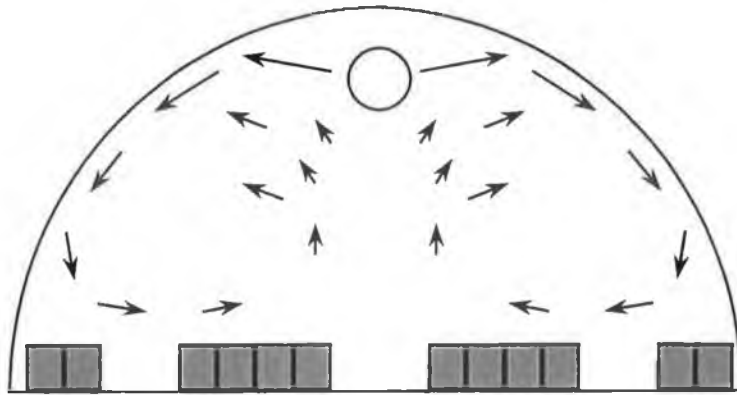
Figure 1.4. The general airflow pattern in a typical bag and tunnel system of growing.
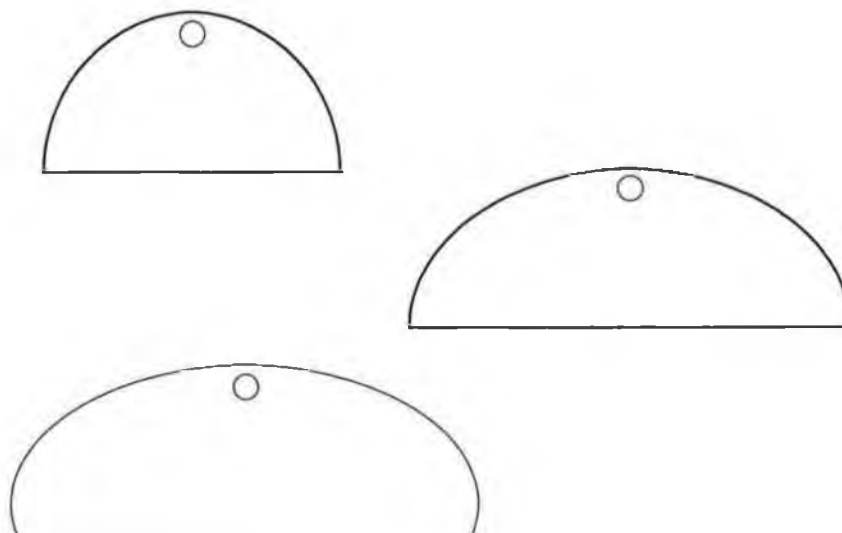
Figure 1.5. The variety of tunnel shapes in use in commercial production.

Airflow in these types of structure has been studied (Bowman, 1987) and there have been a number of control studies (Hayes, 1991, Meath, 1993, Murray, 1995 and Murray et al., 1995) but, to the author's knowledge, no attempt has been made to apply flow calculation techniques to the problem.

There have been a number of publications concerning airflow in mushroom growing but they have largely been aimed at a grower readership and technically detailed publications are limited.

1.3 Publications related to mushrooms

Atkins (1965a and 1965b) in two popular articles drew attention to the importance of airflow in mushroom growing and discussed it in terms of the magnitude, velocity, distribution and conditioning. He quoted figures for the volume of air delivered to the growing room and indicated possible speeds that could result across the crop but noted that the latter were notional because the distribution of air was not uniform within the growing rooms. Schroeder (1968) examined the air supply and distribution using an experimental ventilation system and concluded that its performance was satisfactory on the basis of temperature uniformity within the growing room. Schroeder et al. (1974) described the use of carbon-dioxide concentration as a variable for automatic control of mushroom tunnel ventilation and drew attention to the energy savings to be made by minimising the introduction of fresh air.

Intensification of mushroom production and the effects of this on the microclimate of the growing room were discussed by Storey (1968). The changes in production focused attention on the rate of carbon dioxide production by the crop and rate of evaporation when quantities of compost (growing substrate) and the total production surface area were increased. Cooling and evaporation in particular became very important aspects of mushroom growing and Storey used Piche evaporimeters to give some guidance on overall rates of moisture loss from the substrate. Flegg (1974) discussed the use of the Piche evaporimeter for water management in mushroom growing. He found that it gave a reliable indication of evaporative conditions in a mushroom house and, when

calibrated, could be used to estimate the amount of water evaporating from the casing layer. A simple and approximate adjustment was given to allow for the effect of compost temperature on evaporative loss of water from the casing

The physiological effects of microclimate factors were discussed in two publications by Tschierpe (1973a and 1973b) in which he considered the microclimate requirements of each phase of the cropping process.

Bishop (1979) drew attention to the lack of appreciation in the industry of the importance of ventilation and control of the microclimate in the growing room. He summarised ventilation requirements and gave guidelines for inlet and exhaust positioning, design of air distribution ducts and the mixing box characteristics of fresh and re-circulated air. A wide-ranging review of available information on mushroom room ventilation was published by Edwards (1975a and 1975b). He described the functions of airflow in mushroom growing and examined the quantitative basis for practises current at the time. The importance of correct design of air delivery and distribution systems was stressed as well as factors like the ratio of total air volume to cropping surface area. A large volume of air relative to the cropping surface area provides a buffer against sudden changes in conditions.

Van Soest (1979) gave some rules of thumb for ventilation requirements and noted that every degree increase in temperature above 16°C results in a 20% increase in the production rate of carbon dioxide that further increases the fresh air requirement. He discussed ventilation arrangements for shelf growing in rectangular rooms with particular emphasis on maximising cooling.

Bailey (1982) published a report on a computer-based design technique for air distribution by polyethylene ducts. This method has been widely used in design of mushroom house air distribution. Uniform delivery of air is a pre-requisite for efficient air distribution within a growing space. A microcomputer based control system for mushroom tunnels was described by Burrage et al. (1988). The air-conditioning system to be controlled was described as well as the actuators and sensors. The software and relatively simple control approach were discussed. The importance of

humidity control for the improvement of mushroom quality was highlighted by Barber and Summerfield (1989) who described the role of intercellular water in the whiteness of the mushrooms. A low rate of evaporation was observed to result in a damp mushroom cap surface and reduced whiteness by a process likened to the difference in appearance between dry and damp blotting paper.

Bowman (1991) consider the problem of uniform air distribution in mushroom cropping house with a tiered growing system and found that improved control of air distribution is possible if conditioned air is supplied via a permeable distribution duct and air circulation established by means of wall-mounted ducts. Such an arrangement was considered to greatly reduce the risk of crop loss resulting from attack by bacterial blotch, or reduced market values because of scaling.

Loeffen (1991 and 1992) carried out a series of measurements at varying air flow rates to examine uniformity of air delivery at different levels in a shelf-growing room with five levels of cropping surface. He noted a great deal of variation in the cropping surface speeds but did not address improvements in uniformity. He optimised the ducted delivery system and noted that the average cropping surface air speed was directly related to inlet speeds. Leakage rate and its relationship with the rate of flow of re-circulated air was examined in Loeffen (1993) where he found a linear relationship between the two. Loeffen (1994 and 1995) has also re-evaluated the importance of carbon dioxide concentration as a control variable and has shown that minimising carbon dioxide production leads to higher yields and that the rate of carbon dioxide production of fresh compost may be a useful indicator of productivity.

Lomax et al. (1996) examined the effect of air quality on cropping patterns using lightweight flag indicators (Lomax et al., 1995) to observe the airflows. Airflow was constrained to travel along growing surface without mixing with or entraining room air. The changes in quality (temperature, moisture and gaseous components) were recorded and were found to affect the rate of crop growth.

1.4 Relevant work in related areas

The available work on the modelling of mushroom tunnel or growing room air flows appears to be limited to lumped parameter control models where the system is assumed to be perfectly mixed and therefore single point measurements can be taken as representative of the air mass as a whole. Measurement (e.g. Loeffen, 1991 and 1992) has shown that this assumption is not valid and work in Irish mushroom tunnels (Grant, 1995) has shown deficiencies in uniformity of air delivery and temperature gradients vertically and longitudinally in tunnels. There is a need to investigate the details of air flow in growing structures and, in a related area, a greater emphasis on more detailed modelling is to be found in the investigation of agricultural livestock housing.

The heat and mass balance equations for a ventilated room cannot be solved without making some hypothesis concerning the air distribution within the airspace and, as noted above, a common approach is to assume complete mixing which equates the thermodynamic properties of the exhaust air to the average thermodynamic properties of the bulk airspace. A theoretical analysis (Barber and Ogilvie, 1982) suggested that departure from complete mixing may be caused by the formation of multiple flow regions within the airspace, or by short-circuiting of the supply air to the exhaust outlet. The work considered detection and quantification of departure from complete mixing and the consequences of failing to account for incomplete mixing. The theoretical study was complemented by a later paper describing a scale model study of incomplete mixing (Barber and Ogilvie, 1984). Rate-of-decay tracer gas experiments conducted in a one-fifth scale physical model of a slot-ventilated airspace tested the validity of a two-parameter mathematical mixing model. Extensive departure from complete mixing was shown to have occurred in the scale model airspace and invalidated the proposed model. The nature and extent of departure from complete mixing in a particular airspace was shown to be a function of the Archimedes number (ratio of buoyancy to inertial forces), the inlet jet momentum, the inlet/outlet configuration and the geometry of internal obstacles to airflow.

Deurloo et al. (1990) investigated the effects of re-circulated air on the air speed within the animal zone using a delivery duct system related to the sort of design used in

mushroom growing. A duct at a high level on one side of an animal house was used to re-circulate and drive air down to the animals on the floor. The air speed close to the animals was approximately correlated with jet speeds at the re-circulation duct.

Some work on positive ventilation with fans was presented in a series of papers by Randall and Moulsley (1990a, 1990b, 1990c) where they examined fan performance in a test rig. The set of papers aimed at establishing the criteria for selecting propeller fans for use in the ventilation of livestock buildings. Their computer-controlled test rig logged a large number of parameters, monitored criteria such as stability of operation and thereby controlled progress of the tests before producing performance reports.

Empirical studies of airflow have been carried out. An example is the work of De Praetere and Van Der Biest (1990), who examined air flow and ammonia concentration in a piggery with fully slatted floors. The authors stated that previous studies had been limited to the space above the slats as if the slatted floors behaved as a solid surface. They showed that two secondary airflow patterns exist: one that remains above the floor and another that extends into the space below the floor. The flow that extended into the slurry pit was found to affect the slurry temperature and to bring ammonia into the building, thereby creating a spatial variation of the ammonia concentration within the building. This paper showed the potential for modelling of airflows in that the authors expressed their surprise that some results were contrary to expectations. A computational fluid dynamics (CFD) approach to this type of problem could yield an increased understanding of the airflow patterns and inform experimental approaches to such problems.

A mathematical model, based on Darcy's law for fluid flow in a porous medium, was developed (Parsons, 1991) for the flow of gas in a bunker silo after opening the front face. It was assumed that the driving force was the difference in density between the air outside and the carbon dioxide-rich gas inside the clamp. The model was solved using the Phoenics CFD package.

Zhang et al. (1992) developed a heat and mass transfer model for the airspace in a ventilated livestock building. They concluded that, while agreement between the model

and real data was not exact, it was good enough to use as a basis for testing control algorithms. One of the main sources of difficulty related to the airflow within the space. In particular, the response to sudden temperature changes (thermal buoyancy) and a consequent inability to simulate accurately the exhaust airflow delivered by propellor fans. Conditions for stable air flow patterns that avoid draughts at animal level in livestock buildings were examined by Berckmans et al. (1993) for real situations and it was found that the corrected Archimedes number (Randall and Battams, 1979) was a reliable guide to when inlet jets would stay horizontal and when they would fall to the floor.

A three-dimensional turbulence model was used by Hoff et al. (1992) to determine the effects of animal-generated buoyant forces on the airflow patterns and temperature and air speed distributions in a ceiling-slot, ventilated, swine grower facility. The calculations incorporated a low-Reynolds turbulence model and the results were compared with experimental results from a scaled enclosure. There were some discrepancies but these were attributed to errors in the assumptions about the inlet flow development.

Numerical simulation of particle transport (Maghirang et al., 1994) was used to determine the influence of exhaust placement, inlet placement, obstruction and number of inlets on particle transport. In this case the air-flow simulation included the k-ε turbulence model and particle transport was calculated using the equation of motion of the particle. The inlet position influenced particle transport but not the concentration while obstructions produced higher concentrations at animal height.

Baker (1994) constructed 1/8th scale models of poultry transport containers to characterise the effects of external pressure on the flow patterns driving their ventilation. He supported his approach with computational fluid dynamics to predict the effects of changes in vehicle structure and transport practices. Hoff (1995) developed a numerical model to predict flow occurring with opposing pane-wall ceiling jets representative of slot-ventilated livestock facilities. The model was evaluated and found to be successful by comparing its predictions with a low Reynolds' number turbulence model, a laminar model and experimental results from a laboratory-scale test chamber. Airflow patterns

and particle transport through a two-storey stack cage layer facility with mild weather conditions were modelled using the CFD code FLUENT. Several ceiling inlet systems were modelled to identify a system with good particulate removal characteristics. The k-ε turbulence model (see Chapter 3) and the particle motion equations were incorporated to predict air velocities in a 1/5 scale model of the cage layer facility. Calculations agreed well with measurements and the results were used to state that ceiling ventilation may be a viable method for transporting particles away from building occupants.

The effectiveness of a three-dimensional mixed-flow turbulence model describing airspeed and temperature distribution in a laboratory scale slot-ventilated chamber was evaluated by Hoff et al. (1995). Airspeed and temperature profiles were investigated near the inlet jet region and along an axial line extending from the inlet. The observed jet-floor impingement location was accurately predicted, as was average building temperature as a function of the corrected Archimedes number at the inlet. Differences between measured and predicted values were 5% or less. Airspeed and temperature profiles near the inlet and along an axial line from the inlet were compared and large localised errors were found.

Heber et al. (1996) provided experimental data from a climate-controlled, full-scale section of a livestock building with a high air flow rate and simulated animal sensible heat. They measured turbulence time and length scales and provided data for the evaluation of numerical predictions by Harral and Boon (1997) who calculated air flow in the section. They used the PHOENICS code to give mean velocity and turbulence energy distribution under isothermal conditions. They concluded that despite its known shortcomings, use of a standard turbulence model to simulate the turbulence energy generation and dissipation gave reasonable predictions of turbulence energy distribution.

The most typical modelling approach in the research of climate control in livestock buildings has been the use of identification methods operating on experimental data. Recently, Zhang et al. (1996) have used this approach in order to find the criteria that can be used to control the trajectory of the airflow into a ventilated airspace, modelling the drop distance using the Archimedes number. It was shown, using this and the inlet opening and configuration, that the jet drop distance could be accurately modelled using

parameter estimation. In a less detailed approach, a simple but dynamic climate model for pig houses was developed for use in the tuning (gain factors and time constants) of controllers for room temperature, air quality and ventilation (Vanklooster et al., 1995)

CFD is seen as offering a means for simulation of contaminant transport processes, with quantitative richness of detail rarely possible in laboratory testing (e.g. Baker et al., 1994). Such an approach demands an efficient but accurate code and non-specialists must be able to depend on the results from commercial packages. A recent interesting investigation in this regard was the calculation of greenhouse flow fields by Boulard et al. (1997). These workers applied two different commercial codes (PHOENICS and CFD2000) to solve a natural-convection-driven flow in a single span greenhouse. The calculation results disagreed with experimental results and, while this is often attributed to insufficiently accurate boundary information, the packages gave different results, qualitatively as well as quantitatively. The authors allowed the possibility of some experimental error in the data but this does not impinge on the difficulty raised by the disagreements between packages.

1.6 Review of methods

The methods involved in this work are discussed, and the relevant literature is reviewed, as they are introduced in the following chapters. As this thesis draws on a number of areas of work, it proved difficult to assemble all the references into a single coherent review.

Chapter 2

Experimental work

2. 1 Introduction

While a full measurement programme that could be used to validate the output from a mathematical flow model was not implemented in time for use in this work, an experimental investigation of the general characteristics and some important features of the air flow in Irish mushroom-growing tunnels was carried out. An initial investigation (O' Flaherty, 1988) found that the typical speeds used during the cropping phase of the production cycle were of very low magnitude and that this might pose considerable difficulty in studying the characteristics of the airflow at the cropping surface. Because of the reported importance of providing the correct airflow, Tschierpe, H.J. (1973a and 1973b), for the production of a high quality crop, the investigation was continued. All the investigations presented in this Chapter were carried out by the author.

The air delivery system in an Irish mushroom growing tunnel, as described in Section 1.2, has the advantage that it allows high exit speeds (4 to 7 ms^{-1}) at the distribution duct and hence a large volume of air to be supplied while providing the very low air speed that is required for the microclimate at the growing surface. The provision of a speed control for the fan means that the grower has control of the air exchange rates at the crop and a damper system provides a controllable mixture of the fresh and the re-circulated air. The damper is used to provide control of the carbon dioxide concentration in the tunnel and, with suitable outside conditions, can be used to

reduce the internal relative humidity. The crop metabolism produces carbon dioxide and, even though high concentrations are desirable in the early stages of the crop cycle, the demand is always for the reduction of carbon dioxide and, therefore, for the introduction of fresh air. In normal growing practice, there is no demand for the introduction of carbon dioxide to raise the concentration.

2. 2 Methods

The measurements of air speeds and temperatures were made in a full-scale demonstration tunnel at Kinsealy Research Centre and in commercial mushroom growing tunnels.

The cropping surface air speeds were measured with a Dantec Low Velocity Flow Analyser (54N50 Mark II). This is a microprocessor-based anemometer (figure 2.1) for the measurement and the analysis of low air velocities in indoor climate investigations (e.g. Bowman, 1987). It has a fast-reacting omnidirectional velocity sensor which is fully temperature compensated. It works on the basis of the hot-sensor anemometer principle which utilises the relationship between heat transfer and air velocity. The velocity sensor is a Nickel thin-film deposited by sputtering on the surface of a 3 mm glass sphere and is protected against corrosion by means of a 0.5 μm thick quartz coating. The overall accuracy, including the digitising and linearisation errors can be expressed as $\pm$ 1 cm $\pm$ 5% of reading in the 5 to 100 cms^{-1} range. Below 5 cms^{-1} no accuracy is stated as the linearisation table is not based on comparison with other measurements (e.g. laser Doppler anemometry) but on interpolation between 5 cms^{-1}

Figure 2.1. The Dantec Low Velocity Flow Analyzer.

and 0 cms^{-1}. This is not a difficulty in the current work because this is the range of random air movement in mushroom tunnels (section 2.4). The zero flow point, which is a fixed point for the conditioning amplifier is established with a protection tube covering the sensors and closed with a small disc. The instrument has an integration mode that provides the mean velocity over periods of 60, 200 or 400 seconds and instantaneous, linearised values of the velocity are available as analogue voltage signals. The sensor was placed at points on the cropping surface or the tunnel floor, generally halfway between the sidewall and the central distribution ducts.

All the exit speeds from the distribution ducts were measured with a thermal anemometer. The Airflow Developments TA-2-15/3k measures the air speed with automatic temperature compensation and the instrument covers a velocity range of 0 to 15 ms^{-1} with an overall accuracy of ± 2% of the full scale deflection. The average temperature response time was stated as 10 seconds. The air speeds were recorded at the end of the distribution duct distant from the fan because at this point there was no forward component of the exit flow due to the velocity pressure and the readings were taken as representative of the duct exit air speed because the distribution ducts were designed for uniform delivery of air using the method developed by Bailey (1982). When making measurements the sensing element of the anemometer was in the plane of the duct aperture.

The temperatures and the air speeds were recorded on a Grant Instruments Squirrel 1200 data logger. The temperatures were measured with Grant Instruments uu-series thermistor sensors. These had been individually tested against instrumentation that was traceable to US National Bureau of Standards. The maximum deviation of an individual sensor specification from the theoretical characteristic is stated to be 0.1 $^{\circ}$C for the range 0 to 70 $^{\circ}$C. The cable resistance produces an error in the temperature reading and at 25 $^{\circ}$C this is given as 0.0013 $^{\circ}$C per metre of cable. The sensors were supplied with 20 metres of cable giving a maximum error due to the cable resistance of 0.026 $^{\circ}$C. The response time of the sensor in a stainless steel sheath was stated as 6 seconds.

Smoke was used to visualise the flows in the mushroom tunnels. The smoke was white and was produced by burning pellets (ph Smoke Pellets).

The relationship between air speeds leaving the distribution duct and those at the cropping surface was studied for the range of duct exit speeds that are in common use in the industry and under various air conditions. Some investigation of the airflow in multi-level growing systems was carried out.

2. 3 Flow visualisation

Visualising techniques for low velocity airflow can be very useful in gaining an overall understanding of how a particular flow system operates. The technique consists of injecting some visible agent into a flow stream in order to detect the motion of the otherwise-invisible flow. The agent used must be such that it has sufficiently low mass (or neutral buoyancy) so that it does not simply fall through the medium of the flow under examination. For the purposes of these studies, injection of smoke was found to be a convenient and useful means of visualising the fluid flows.

Pellets were burned to produce the smoke and some care in their use was required because of the very low air speed flows that were to be examined. The pellets did not produce sufficient smoke to allow its distribution from the air-handling unit and they had to be introduced at some point in the flow field. Their burning caused a thermal updraught from the pellets and therefore they had to be placed in a position where they did not interfere with the flow to be visualised. Placing them directly below the distribution duct on the floor was the only point at which this could be achieved (figure 2.2) and burning a pellet at any other point across the tunnel did cause a local change in the flow field. The symmetry of the flow field about the centre of the tunnel allowed the effect on the flow pattern to be tested by placing the pellet on the tunnel floor, to one side of the distribution duct, and comparing the flow on the other side to that where the pellet was burned directly under the duct. No differences could be

18

discerned in the main features of the flow and it was therefore concluded that it was acceptable to use the pellets in the centre of the tunnel.

Using the smoke it was established that, while slow-moving, the airflow showed a distinct pattern in its movement and that, at a fixed fan speed, it should be steady on average over a period of time. The air left the polyethylene distribution duct in a series of jets from both sides of the duct. These jets were inclined upwards towards the roof of the tunnel (the apertures were set at approximately 2 and 10 o'clock positions on the circular cross-section duct). This system results in a bilateral symmetry in the airflow that can be appreciated from figure 2.2.

Figure 2.2. Flow visualisation using smoke.

Visualisation showed that the airflow attached to the tunnel wall and maintained the attachment until it reached the crop level. It then flowed across the cropping surface until it met the opposing flow from the other side of the tunnel. The air stream then rose towards the distribution duct again where it was re-entrained by the emerging jets. This is the main flow pattern controlling air speeds at the crop.

19

When the air handling unit was set for full re-circulation within the tunnel there was also a flow along the tunnel towards the re-circulation entry. In the case of the tunnel being fully open, i.e. no re-circulation of air and all the distributed air coming from outside, this flow along the tunnel was in the opposite direction, towards the exhaust outlets that are mounted on the end of the tunnel distant from the fan. In the visualisation these flows were not clearly defined but were observed as a general movement of smoke along the tunnel. Once the smoke was established in the lateral flow, it could be seen that it was moving along the tunnel in both directions but that the movement was faster towards either the re-circulation inlet or exhausts, depending on the damper position.

The smoke testing showed that at the bottom of the tunnel wall there was a region of re-circulation and, as the flow moved towards the centre it began to rise because it was meeting the opposing flow and this resulted in reduced ventilation at the centre of the
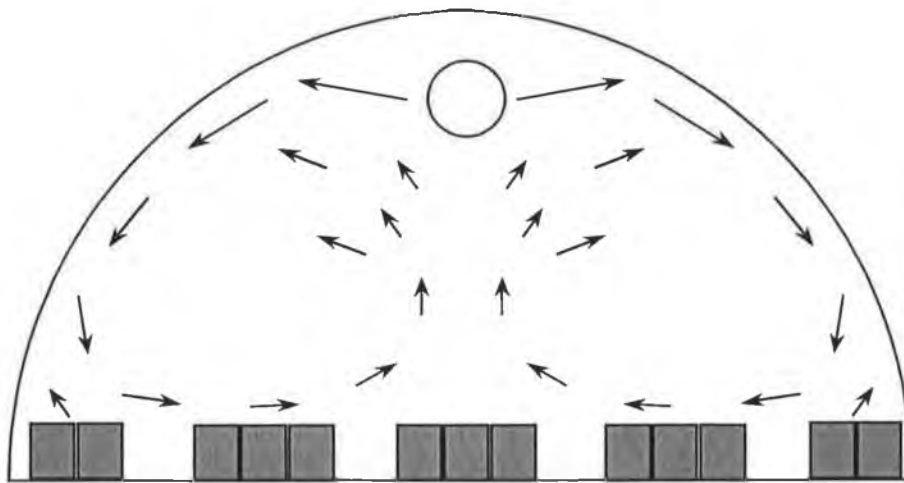

Figure 2.3.   The general air flow pattern in a typical bag and tunnel system of growing.

tunnel, relative to that mid-way between the centre and sidewall. The compost layout is normally in rows that run from front to back of the tunnel with alleyways between

bags for crop management and picker access and, as shown in figure 2.3, this means that some bags at the centre of the tunnel are under-ventilated compared to those that are directly in the airstream. The central area is where some variations in cropping are noted. In particular, the central rows are ready for harvest earlier than others and this is generally due to higher compost temperature. The lack of a positive airflow across them would account for less efficient cooling of bags in the central rows and this is a particular problem in the wider tunnels (8.8 m and 9.2m wide). This has led to the suggestion of alternative layouts for the bags of compost so that there is always a central aisle below the distribution duct.

This general understanding of the airflow pattern indicated that the most useful point at which to measure the air speed at the cropping surface would be halfway between the sidewall and the centre of the tunnel. At this point, or near it, the air speeds measured were expected to give results that could be related to the duct exit air speed and that would represent the flow across the greater proportion of the crop.

The visualisation with smoke was also used as described in section 2.5 to examine the effect of operation of the heating system on airflow at the cropping surface. Before that, the general flow pattern for isothermal conditions was studied to determine the range of speeds that were to be found at the cropping surface and the general relationship with the duct exit speeds.

2. 4   Cropping surface air speeds

A crucial relationship in the air-delivery system is the relationship between the speeds at the apertures on the distribution duct and the corresponding speeds at the cropping surface. At the distribution duct the air speed range (typically 1 ms^{-1} to 8 ms^{-1}) is within the measurement capability of low-cost, widely-available equipment while at the cropping surface the slower-moving flow requires instrumentation that is specialised and expensive and more than a single, rapid measurement is necessary to adequately characterise the flow. It was observed that there were rapid fluctuations in

the speed at the crop and that the averaging of results over a period of time would be required.

Since the air speed at the crop depends partly on the attenuation of the airflow by the tunnel structure and there are many variations in the number and size of holes in the distribution ducts as well as varying models of axial fans with slightly different pressure/output characteristics, it would be anticipated that there would not be a general relationship between the duct exit and the cropping surface air speeds but that the characteristic should be similar in most cases.

Some data from two commercial tunnels are shown in figures 2.4 and 2.5. Both of these tunnels were of the narrower type at 6.7 m wide and 3 m high. Measurements were made halfway between the tunnel sidewall and the central aisle. The air speed sensor was placed on the cropping surface and therefore, due to its height, the speed was monitored at a height of 15 cm above the crop. This was done to avoid any local disturbance of the flow due to the presence of the crop, which would constitute a set of randomly spaced obstructions to the flow near the surface or an overall surface roughness. The smoke testing showed that there was no significant local effect of surface conditions at the measuring height.

The airflow in the tunnels was allowed to establish for a long period at the lowest fan speed setting before beginning the measurements as experience had shown that increasing the air speeds gave the shortest settling time at each point. A number of fan speeds were set and the corresponding cropping surface air speed was measured in each case after allowing a period of time for the flow to become steady at the new setting. This was achieved by monitoring 1 minute averages of the air speed and waiting until the measurements ceased to show an upward trend. The points in the graphs were derived from continuously logged data that were averaged over a 10 minute period.
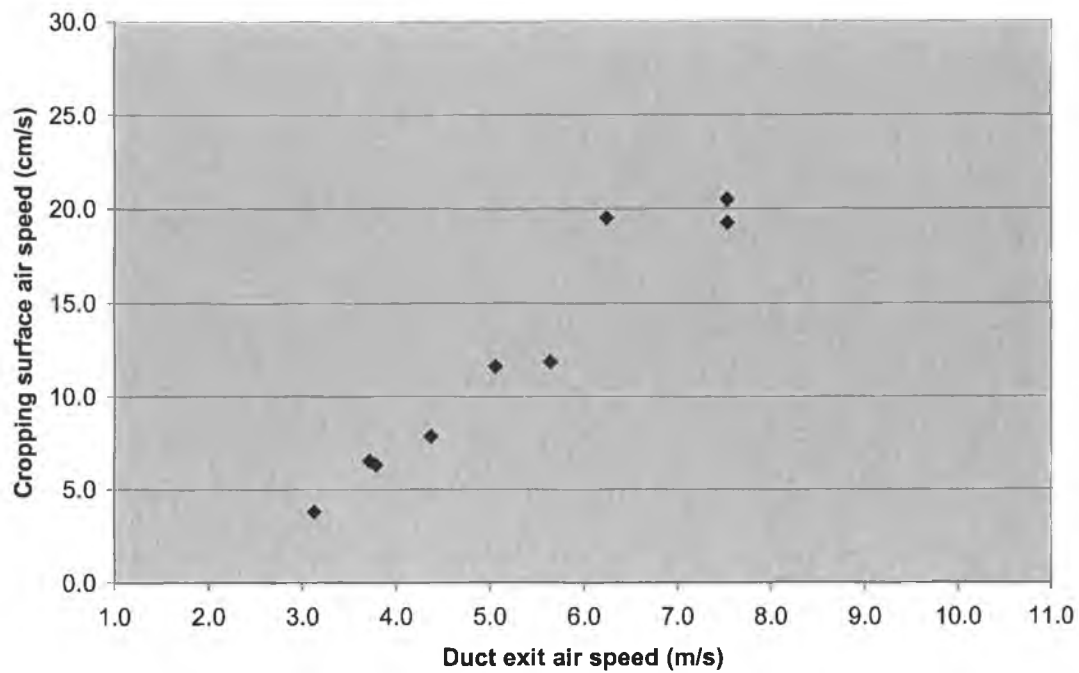
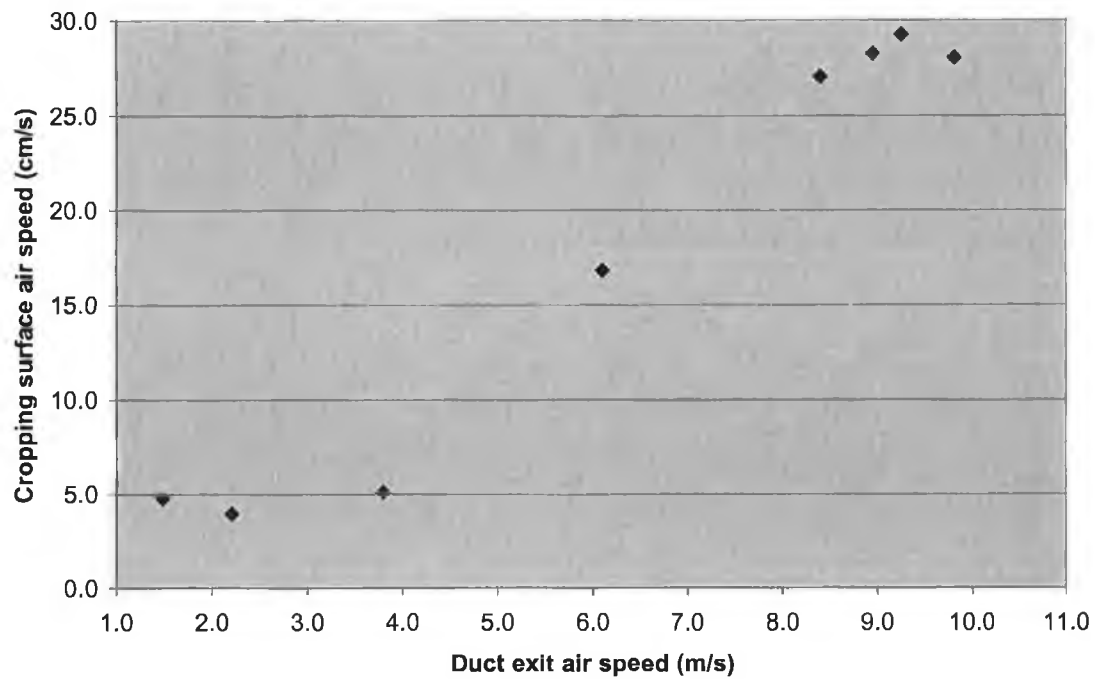**Figure 2.4. Cropping surface air speeds at a number of duct exit air speeds.**

Figure 2.5. Cropping surface air speeds at a number of duct exit air speeds.

The duct exit air speed settings were not uniformly spaced because of the non-linearity of the operating characteristic of the fan speed controller. In practice, it was more convenient to adjust the controller and then measure the resultant air speed rather than adjust the controller to give particular air speeds. The controller for the case in figure 2.4 had a narrower range than the one for the data in figure 2.5 but the latter had an operating characteristic that made it very difficult to resolve the middle of the range. The graphs in figures 2.4 and 2.5 are presented with the same horizontal and vertical scales to allow comparison of the results.

The two graphs show similar overall behaviour of the cropping surface air speed. From figure 2.5, there was a portion of the lower part of the range of the duct exit air speed below 4 ms^{-1} where there was no response in the cropping surface air speed to further reductions in the duct exit air speed. This would be the speed range of random air movement at the cropping surface. Above this there is a response that could be adequately described by a linear relationship. It happens that these two sets of data are not incompatible and could be taken as being from the same system but they cannot be taken as typical or as indicating that general relations can be derived empirically that would apply to all similar structures. However, they do illustrate that it is possible to achieve, with a suitable controller, a wide range of air speeds at the cropping surface.

Adjusting the air speed at the cropping surface can control the crop development and quality but making a change in fan speed also alters the volume of air that passes through the air-handling unit. Adjusting the position of the damper that controls the proportion of fresh air that is introduced can compensate for this if, for example, it is necessary to reduce speed for quality reasons but at the same time maintain a steady ventilation rate to control the carbon dioxide concentration in the tunnel.

To assess the level of random air movement in a mushroom structure a set of measurements were made in an empty tunnel. Four points were examined and logging provided 48 hours of the air speed measurements for each point. The overall average speed over 48 hours was 3.2 cms^{-1}. The highest 10 minute average was 5.1 cms^{-1} and the lowest was 1.5 cms^{-1}. Indications from commercial tunnels during cropping were

that 5 cms^{-1} is typical but, in these situations, it was not possible to allow a long settling time with the fan off.

2. 5   The effects of heating

During the course of measurements and smoke visualisation, a stratification effect was discovered in the tunnels. Smoke was released into the air stream but instead of rising up to the distribution duct and being dispersed by the emerging jets of air it stopped rising and spread laterally at approximately 1 metre above the floor. The smoke finally settled into a horizontal sheet. Any disturbance, such as opening and closing a door, caused waves to move along the surface of the smoke sheet but did not significantly disperse the smoke. The smoke persisted in this sheet for a prolonged period. It was

Figure 2.6. A stratified airflow in a mushroom tunnel.

clear that the main airflow from the distribution duct was not reaching the cropping surface.

Figure 2.6 is a picture of a smoke test that shows a stratified airflow. The sheet of smoke was produced, as described above, by releasing the smoke in the lower stratum. Another approach would be to release the smoke into the air handling unit. In this case, while the small quantity of the smoke would disperse throughout the upper stratum it was possible to see that none of it was entering the lower one and it was completely clear.

This behaviour of the airflow was a possible explanation of the occurrence of some cropping conditions that had been difficult to understand. A condition (bacterial blotching) that is associated with a damp mushroom surface had been observed in several instances where the measured air speeds should have been sufficient to provide the drying required for a high quality crop. Such measurements, if made while the heating was off, could be misleading.

It quickly became clear that the stratification was associated with the action of the heating system and that air speeds at the cropping surface decreased whenever the heating system was switched on. A programme of measurements was undertaken to investigate the effect of the heating on airflow.

In order to avoid the possible confusion of the results all the measurements that are discussed below were made under conditions of low solar radiation. It was known that the outer cover of the insulated structure could be heated to high temperatures under a high solar radiation load. The effect of such high solar radiation was measured by monitoring temperatures within the tunnel structure. The data from such a period is shown in figure 2.7. This set of observation was made in the month May which provided a cold night followed by a bright day and the figure shows temperatures recorded over a 24 hour period from midnight to midnight.
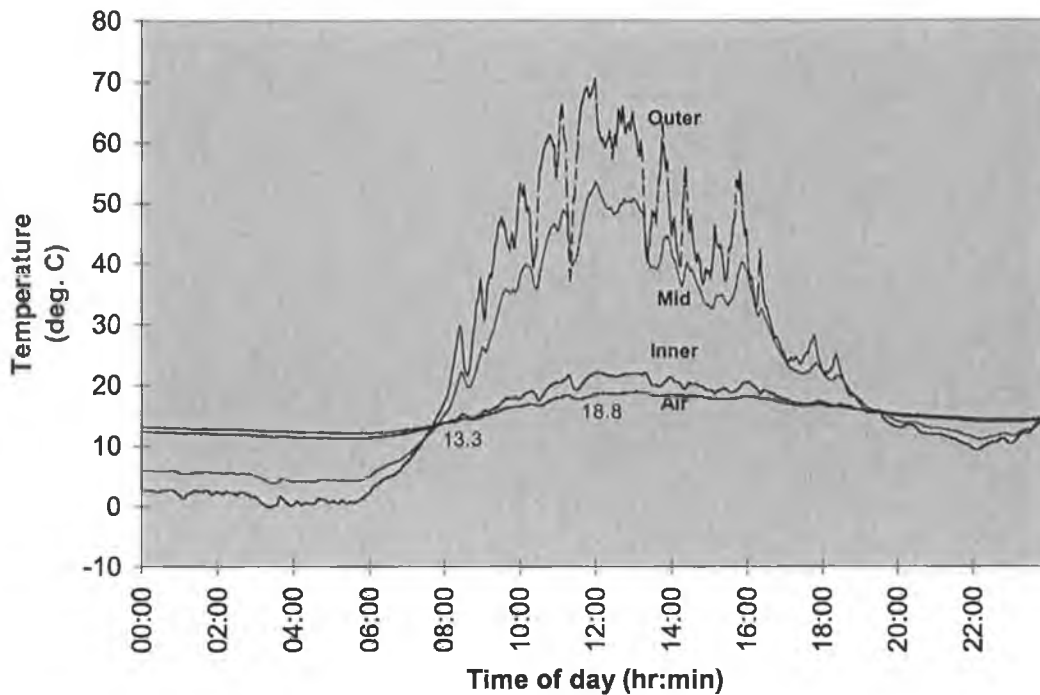
Figure 2.7. Wall and air temperatures measured in a mushroom tunnel under high levels of solar radiation.

Temperature sensors were placed directly under the other cover, in the insulation halfway between the inner and other covers, immediately inside the inner cover and in the tunnel air. The tunnel air-handling unit was set for full re-circulation and maximum airflow. The temperature of the outer skin reached very high values and this caused the other wall temperatures to rise as shown in the traces labelled Mid and Inner. The data in figure 2.7 show that the air temperature in the tunnel rose from 13.3°C to a maximum of 18.8°C and it is clear there is substantial heat entering the system. For the purposes of separating effects in the study of mushroom tunnel airflows, it was therefore important to avoid situations of high solar radiation where the outer covering heated above the ambient temperature. It should be noted that the air speed data presented in Section 2.4 were obtained after these observations and that those measurements were made under conditions of low solar radiation as well as all those that follow.

There appeared to be two effects that contributed to the reduction in air speed. The first was a reduction in the exit speed from the duct and other was the effect of thermal buoyancy in the tunnel air.

In the air handling units, the fan was installed downstream of the other components just before the entry to the distribution duct. This means that whenever the heating was on, the fan was driving warmed, reduced density air and this was reflected in a reduction in the air speed measured at the exit holes in the duct. The thermal expansion of the heat exchanger may also have contributed to the air speed reductions by increasing its resistance to the airflow through it. Figure 2.8 shows measurements made at two different initial duct-exit air speeds. These particular air speeds were important for mushroom growing. The lower one, 4.4 ms^{-1}, was a typical harvesting speed and the upper one, 6.1 ms^{-1}, was at a level recommended for inducing fructification. No heat had been applied for a period before the measurements were made. Air speed was measured at 1 minute intervals.
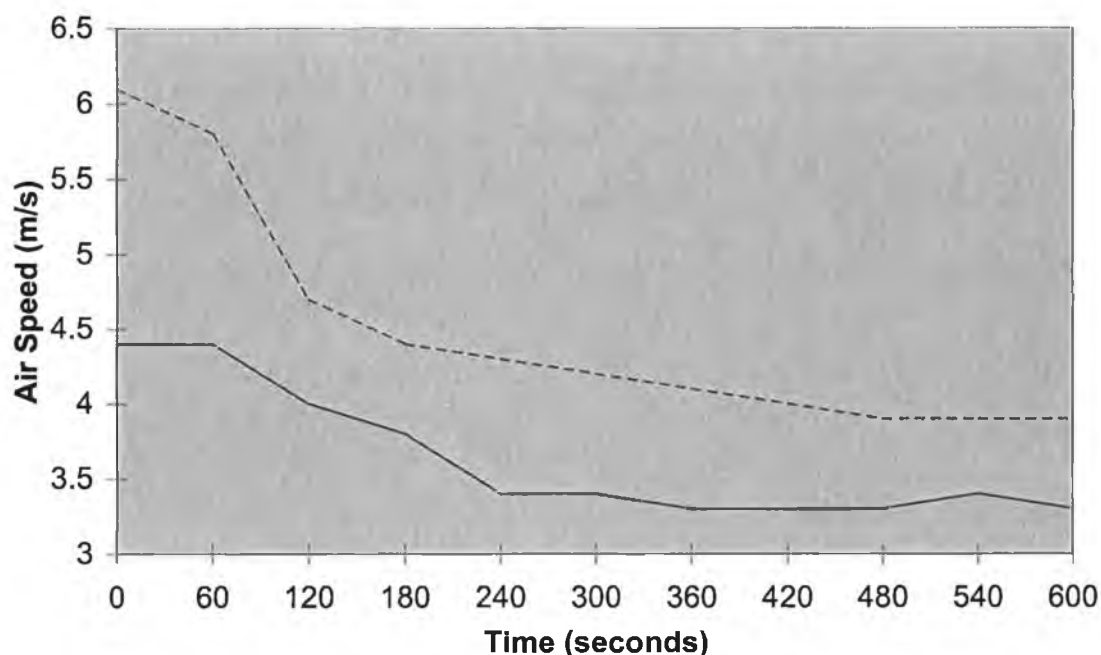
Figure 2.8. The change in duct exit air speeds during heating.

It can be seen in figure 2.8 that the air speed from the duct falls within minutes of the beginning of heating. The lower initial speed falls quickly below the recommended minimum air speed of $4ms^{-1}$, while the higher speed is reduced to one more suitable for harvesting and well below the level required for fructification. Thus, while the reductions in duct exit air speed are 25% and 36% for these two cases, when considered within the operating range of 4 to 6 ms^{-1}, the effects on the cropping process are very much more severe. The time that the heating is on will vary with the demand but would rarely be less than 5 minutes for the standard on-off controllers in use and such reductions would be expected to have an impact on the crop.

The reduction in duct exit air speed is one component of the heating effects but this alone was not enough to account for the very low air speeds at the cropping surface and the stratification. The other is the effect of thermal buoyancy once the warmed air enters the tunnel airflow. The cropping surface speeds are low and easily disturbed and it was found that the combination of reduced air speed and thermal buoyancy was sufficient to cause the airflow at the crop to drop to the level of random air movement.

Figure 2.9. Cropping surface air speed during heating (4 ms⁻¹ at duct exit).

Figure 2.9 shows the effect of heating on the cropping surface air speed with a duct exit speed of 4 ms^{-1}. The lower, rapidly varying line is the cropping surface air speed measured half way between the wall and the distribution duct and the upper traces are the temperature in the distribution duct, the air temperature at a height of 2 metres and the air tempe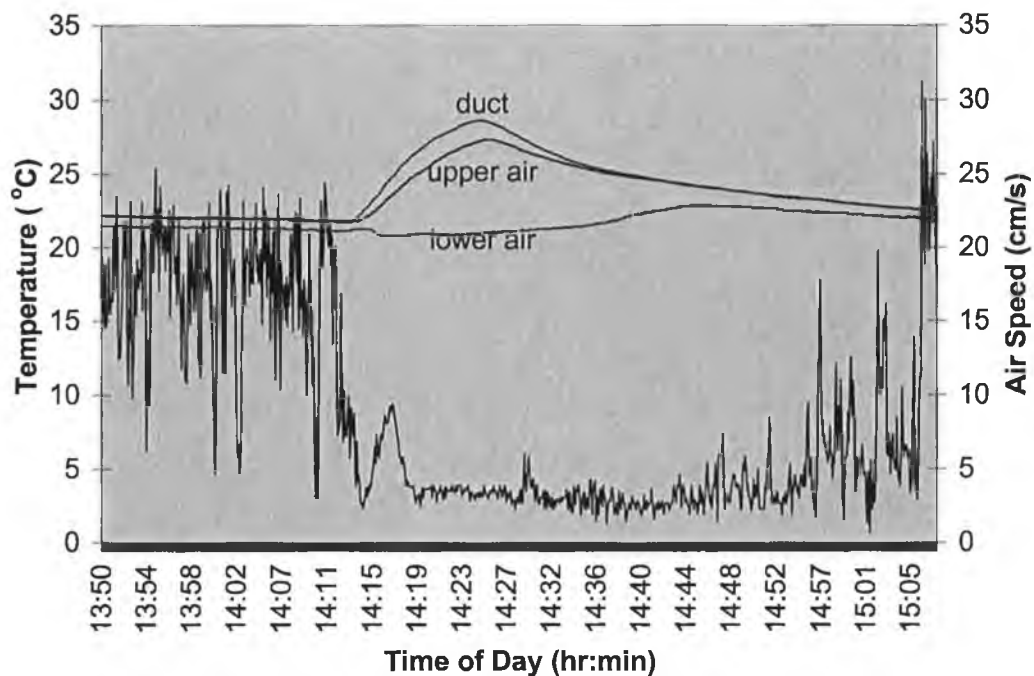rature at the cropping level. The point where heating was applied is seen clearly where the air temperature rises in the distribution duct.

The air speed before the heating was at a mean level of approximately 17 cms^{-1} and the effect of the heating is clearly seen in the figure. When heating was turned on the air speed fell quickly to a very low level and maintained that low level for a period of time. The point where the heat was turned off is where the duct temperature peaks and then begins to fall. The air speed did not recover at this point. The low speed persisted for a period afterwards for a period that was greater that the heat on time.

The upper and lower air temperatures indicate that stratification of the airflow had occurred. The upper air temperature rises with the duct air temperature after a short delay but the lower air temperature remains at approximately its initial level until after the heating is turned off. This reflects the earlier observation that circulation of air continues in the upper part of the tunnel airspace and that heat is distributed efficiently within that region while the air lower section is effectively stagnant and the heat can only reach the cropping surface by diffusion. This was verified with the smoke tests. Injecting smoke into the distribution duct showed that circulation was present in the upper section of the tunnel and that no direct airflow reached the lower part.

The temperature traces also show the re-establishment of good mixing throughout the tunnel. The lower air temperature begins to rise, slowly at first, after heating is turned off and then later rises more quickly. This corresponds to air speed beginning to increase again and indicates improved mixing and that the stratified airflow is returning to a fully mixed pattern. The temperatures do not converge completely until the airflow is fully re-established.

A very similar result was obtained for higher duct exit air speeds. Figure 2.10 shows a result for a duct exit speed of 6 ms^{-1}. The lines in the graph correspond to those in

figure 2.9 and the sequence of events follows the same pattern but in this case the initial air speeds are higher because of the higher duct exit speed. The minimum speed is similar to the case in figure 2.9. It should be noted that there is a difference in the behaviour of the upper air temperature in this case. For the case in figure 2.9 the upper temperature followed the duct temperature more closely. In figure 2.10 it could be that the higher duct exit speed delayed the stratification and thus the upper, warm zone took longer to establish. Once the flow to the floor had stopped, the upper temperature showed a significant increase. As before, the low air speeds persisted for a period after the heat was turned off and equalisation of temperatures (good mixing) was not achieved until the air speeds increased again. The recovery time after heating was shorter for the higher duct exit air speed. Figure 2.10 shows that this time was approximately 10 minutes while the time in figure 2.9 was almost 30 minutes.
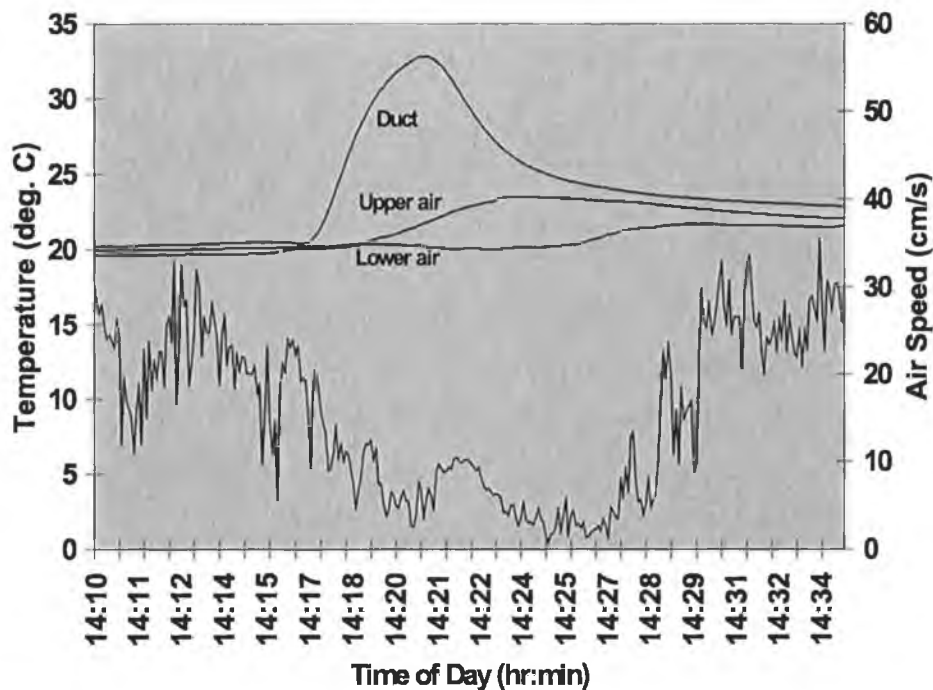
Figure 2.10. Cropping surface air speed during heating (6 ms^{-1} at duct exit).

It might, however, be that the recovery time for the higher speed situations would increase if the stratification was better established. There was only a short period of very low air speed and the upper zone therefore did not heat to the same extent as that in figure 2.9.

There is thus a very pronounced reduction in the air speed whenever the heating is on. The airflow pattern that occurs during the stratification is indicated schematically in figure 2.11. It was common practice in the industry to site the temperature sensor at the cropping surface and this does appear to be the most logical place for it. It can now be seen that this position would produce the worst case for air speed reduction as the sensor would be outside the zone of efficient heat delivery. The heat could reach it by slow processes only.
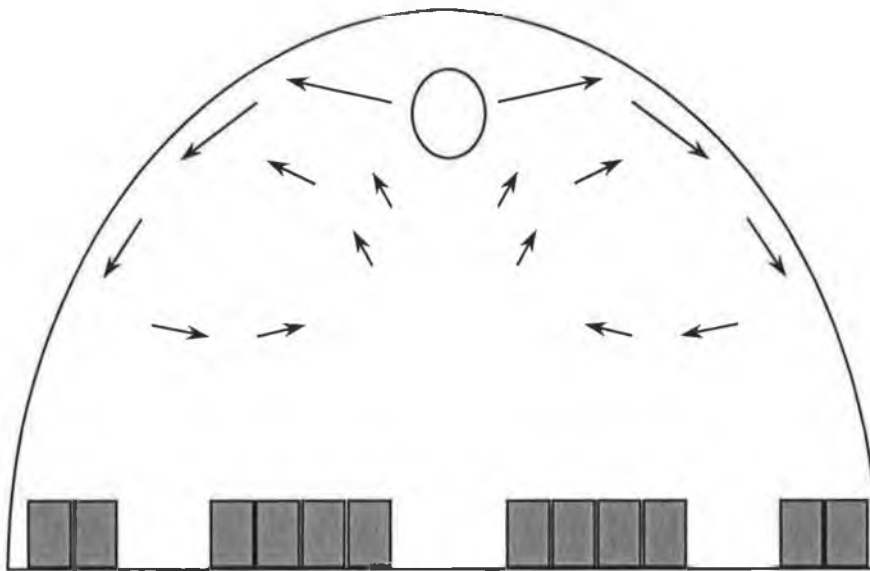
Figure 2.11. The circulation pattern in a stratified flow field.

An important practice for high quality production is the quick drying of the surfaces of the mushrooms after watering. If a mushroom surface remains damp then there is a threshold level of bacterial population that can be exceeded and brown patches or surface pitting can occur (e.g. Tschierpe, H.J., 1973a and 1973b). It was common practice to increase the air temperature after watering in order to stimulate drying. It can now be seen that if this resulted in a prolonged period of heating then the effect

could be to reduce drying for a substantial period of time. Since the air speeds are typically low during the harvesting period in order to avoid excessive drying, it is likely that the recovery from stratification could be prolonged

The results in figures 2.9 and 2.10 would apply to those situations where the heating control is of on-off type. Further observations were made to see if the effect could be avoided by using other control approaches. The use of a proportioning valve with control of duct air temperature and pulsed heating with a solenoid valve were examined.

Figure 2.12 shows a result for a proportioning valve system. In this case, the controller maintained a temperature difference between the air delivered and the air temperature in the tunnel that depended on the control error in the tunnel air temperature, i.e. the difference between the set point and the actual temperature. The figure shows the effect on air speed as the temperature error causes the heating to turn on. The cropping surface air speed is the rapidly fluctuating trace at the bottom of the graph and the temperature in the air distribution duct and in the tunnel air are labelled.
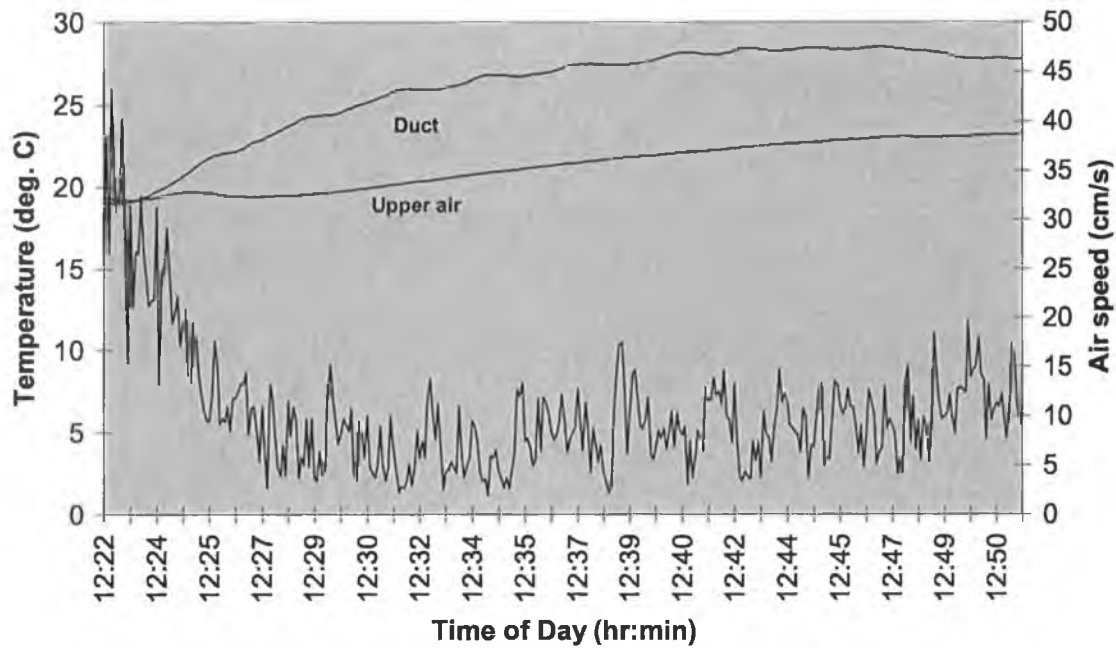
Figure 2.12. Air speed and air temperatures for a proportioning control system.

The initial air speed was approximately 20 cms^{-1} and this was reduced to approximately 5 cms^{-1} a short time after heating began. It appeared that even moderate rises in air temperature could result in a loss of airflow at the cropping surface.

Another alternative heating system is a pulsed system where heat is applied in bursts that are limited in duration but not in frequency. It was proposed that this system should be superior to others in terms of airflow effects because the short on time would not allow the development of a stratified flow field. Such a system was monitored in a commercial mushroom tunnel in normal operation and results are shown in figure 2.13. The lower part of this graph shows cropping surface air speed and the temperature in the distribution duct and a height of 2 m above the ground are also shown.
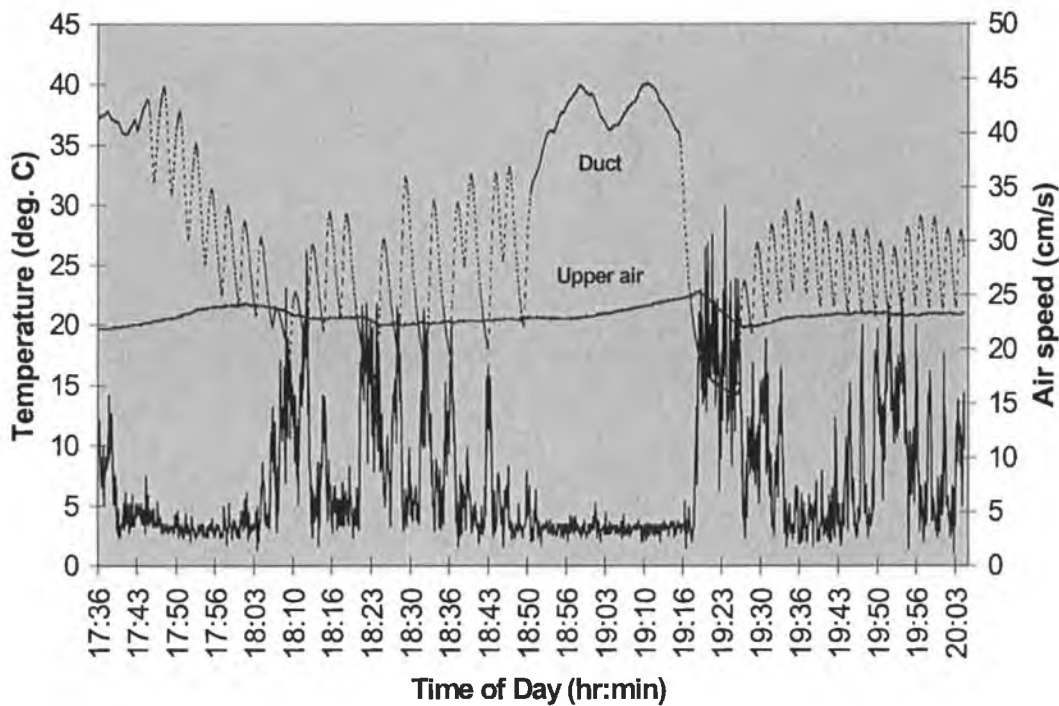
Figure 2.13. Air speed and temperature for a pulsed heating system.

The sequence of measurements starts from a point where the heating load was high and the air temperature in the tunnel was rising to the set point. The air temperature in the distribution duct was more rapidly varying in this case because of the pulsing action. The cropping surface air speed was initially low because of the high duct air temperature and, as the number of pulses decreased and the duct air temperature decreased, the airflow re-established. It can be seen, however, that the first pulse after the air speed had risen was sufficient to cause another decrease. This could have been a normal air speed fluctuation but each succeeding pulse caused another decrease.

After a period of pulsing operation, the heating was turned fully on to deliver sufficient heat to reach the set point. After this the airflow re-established while the heat was off but once pulses began again, the air speed decreased and, as average duct temperature rose, it was severely reduced.

For the type of air-handling unit used in mushroom growing, it appears that there is no simple variation of the heating control that can alleviate the problem of stratification.

The reduction in airflow over a period will depend on the heating load and the effect is therefore seasonal. Figures 2.14 and 2.15 show extremes of the effect. Figure 2.14 shows an evening where the setpoint for temperature was not attained, the heating was on for the entire period shown (and the rest of the night) and airflow was stopped at the crop. Outside temperature during the period fell from 8°C to 4°C and the intake was set for full fresh air. The boiler output temperature was set at a relatively low level of 60°C because it was common practice in the industry to set the boiler temperature relatively low during the summer months. The lowest trace is air speed and the upper traces are the temperatures in the distribution duct and at the cropping surface. The oscillation in the duct air temperature was due to the cycling of the boiler as it maintained its output temperature. The air speed at the beginning of the period shown was well-established but, as heat was applied, it fell to a very low level from which it did not recover until outside temperatures rose the following day.
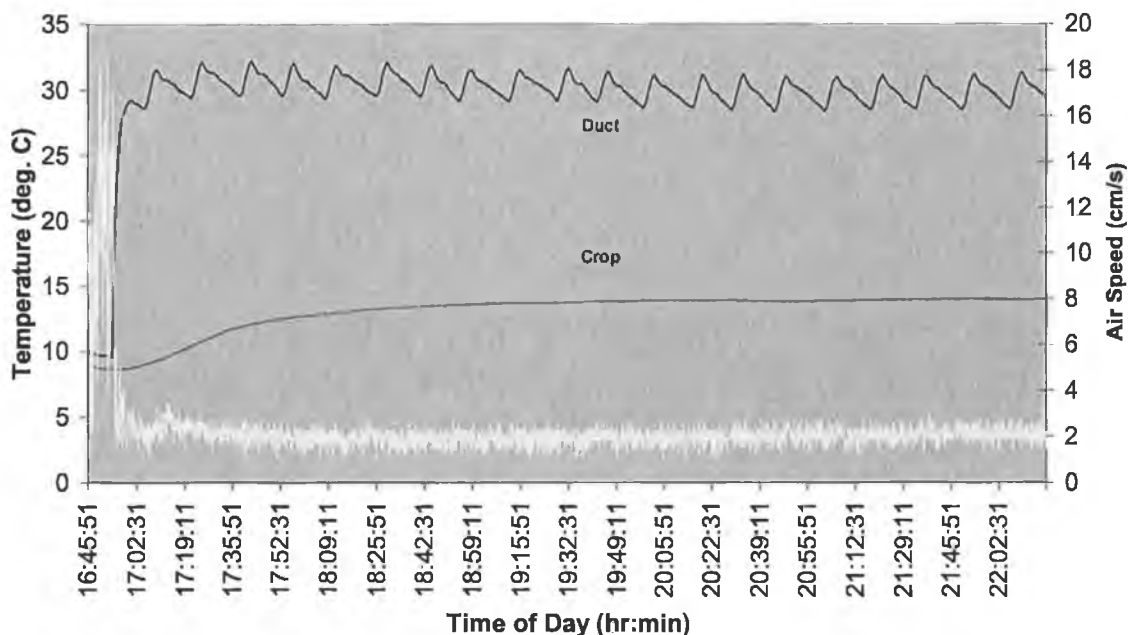
Figure 2.14. Cropping surface air speed and temperatures for a high heating load.

Such situations arise during fructification (the production of the mushroom body) when the crop is producing large quantities of carbon dioxide and ventilation is required to maintain a low concentration in the internal air. Neglecting infrequent extremes, the design criteria for heating systems in horticulture in Ireland assume an outside minimum temperature of -4^OC (O'Flaherty, 1990) and therefore there are likely to be many occasions where the heating load is greater than that during the measurements in figure 2.14. It is clear that very poor airflow due to the operation of the heating system can be a serious problem in mushroom growing.

For the situation in figure 2.14, the heating was on all the time. It was also possible to find some situations where the heating was switching on and off but in which there was insufficient time for the flow to re-establish and the average air speed was very low compared to what it would have been at a given fan speed setting with no heating. Figure 2.15 shows such a situation. Again, the air speed is the lowest trace while the air temperature traces are labelled and these are the temperatures in the distribution duct and at the height of the cropping surface. The intake was set to full fresh air and the outside temperature varied between 3^o C and 6^oC. The boiler temperature was set to 80^oC so that, despite a greater heating load than for the previous case, it can be seen that the temperature setpoint was attained and the heating was switching on and off frequently. The time between heat off and heat on was sufficient to allow some recovery of flow but it only attained the nominal average speed of 15 cms^{-1} for two very brief periods and with an overall average of 3.3 cms^{-1}, there was effectively no significant airflow over the period represented in the graph. In this situation a grower could have no control over growing conditions and could have little expectation of a high quality crop.
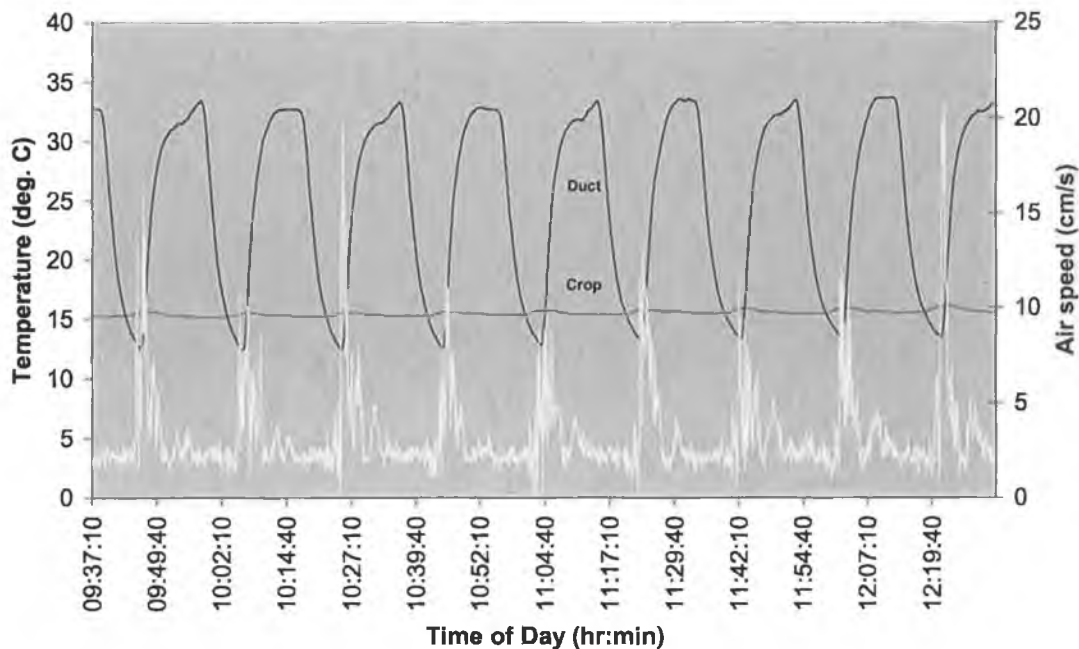
Figure 2.15. Cropping surface air speed and temperatures for a high heating load.

As noted above, positioning the temperature control sensor at the cropping surface places it in the lower layer of air during stratification and this delays the delivery of heat to that area because the direct airflow from the distribution duct cannot reach it. Experiments were conducted to determine whether a higher position for the control sensor could improve the average air speed at the crop. If the heated air reached the sensor more quickly then the heating system would turn off and airflow could re-establish more quickly. The trade-off in such a situation would be that there would be more frequent switching of the heating system and, as well as increased wear and tear on the actuators, there would be more, but briefer, periods of stratification.

It was found that the higher position for the air temperature control sensor resulted in a higher average air speed at the cropping surface in all situations examined except for the case of very high heating load where, despite a small numerical advantage for the lower position the air speeds in both cases were too low to be useful.
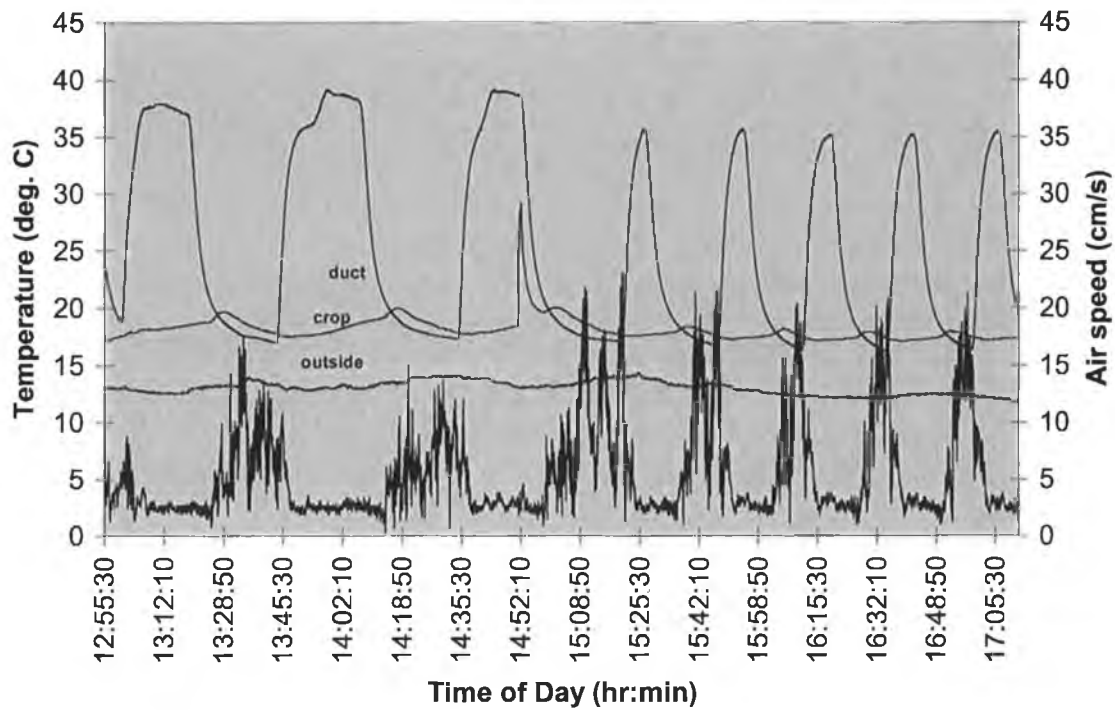
Figure 2.16. Air speed and temperatures for two different positions of the temperature control sensor.

Figure 2.16 shows a comparison of high (2 m above the floor) and low (cropping surface) positions for the temperature control sensor over the course of a day with low solar radiation throughout the period and a low variation in outside air temperature and hence heating load. The higher position was selected to be within the upper, circulating airflow during heating. The height above the ground was set empirically by the need to heat a large proportion of the tunnel air each time thus ensuring sufficient heat to lift the temperature of the entire mass of air after heating. A higher position resulted in the heating system switching too quickly with the consequence of shortening of the useful life of the actuators. A lower position would place the sensor back in the lower layer during stratification.

In the sequence of events shown in figure 2.16, the control sensor was initially in the lower position and after a period of time it was raised to the higher level. The point where it was raised is indicated in the graph by a heating of the cropping surface

temperature sensor to produce a marker spike. Immediately after this there was a transition period and then, once heat was again turned on, the pattern of heat input changes to one of more frequent operation. Taking a number of temperature cycles from before the change and a number afterwards the average air speeds were found to be 4.7 cms^{-1} with the control sensor in the lower position and 6.2 cms^{-1} using the upper position. This might not in itself represent a significant improvement in air flow but figure 2.16 shows that the shorter on time for the heating seems to allow a quicker recovery in the airflow and some brief periods of good mixing ensued. A consequence of this better mixing is seen in that, when the sensor is in the higher position, the swings in the crop level temperature are smaller than when the sensor was lower. The upper layer was overheating and causing an overshoot at the crop level after mixing re-established. More frequent mixing eliminated this effect. The overheating may also have delayed the return to the air speeds that would prevail with no heating.

Experimentation with the different positions for the temperature sensor was difficult because only one structure was available for the tests. Ideally comparison of identical structures running in parallel would be required to eliminate any effects of conditions outside the tunnel but, recognising this, all consecutive tests with similar outside conditions during testing showed an advantage for the upper sensor position.

An ideal process of dehumidification would consist of the removal of moisture from the air without a corresponding change in temperature. In practice, however, this is not possible in a typical mushroom installation. A chilling unit is used to provide a cold surface that is below the dew point of the air stream passing through it. In this situation condensation occurs and moisture is thus removed from the air. This is the primary aim in dehumidifying for mushroom growing but a chilling of the air stream is a consequence of this process and a heat input is required to maintain the air temperature. The air stream in a standard air-handling unit first encounters the chilling coil and then the heating coil is upstream. The purpose of dehumidifying is normally to enhance drying but it was found that the re-heating of the chilled air acted against this intention.

The simplest way to implement the cool/heat type of dehumidification is to allow heating to be controlled independently about a set point. Figure 2.17 shows a sequence of results that illustrates this. Air speed and duct temperature are shown in the graph. The cooling was already on and the duct temperature had reached its minimum value at the beginning of the sequence of events shown. At the same time, the air temperature in the tunnel was dropping and after a short time it reached the set point for heating and this is clearly seen as the point where the duct temperature rose sharply.

When the heating was on the effect on air speed was very similar to the heating situation discussed above. Despite the production of chilled air, the temperature in the duct rose and this caused a similar rapid drop in air speed at the cropping surface.
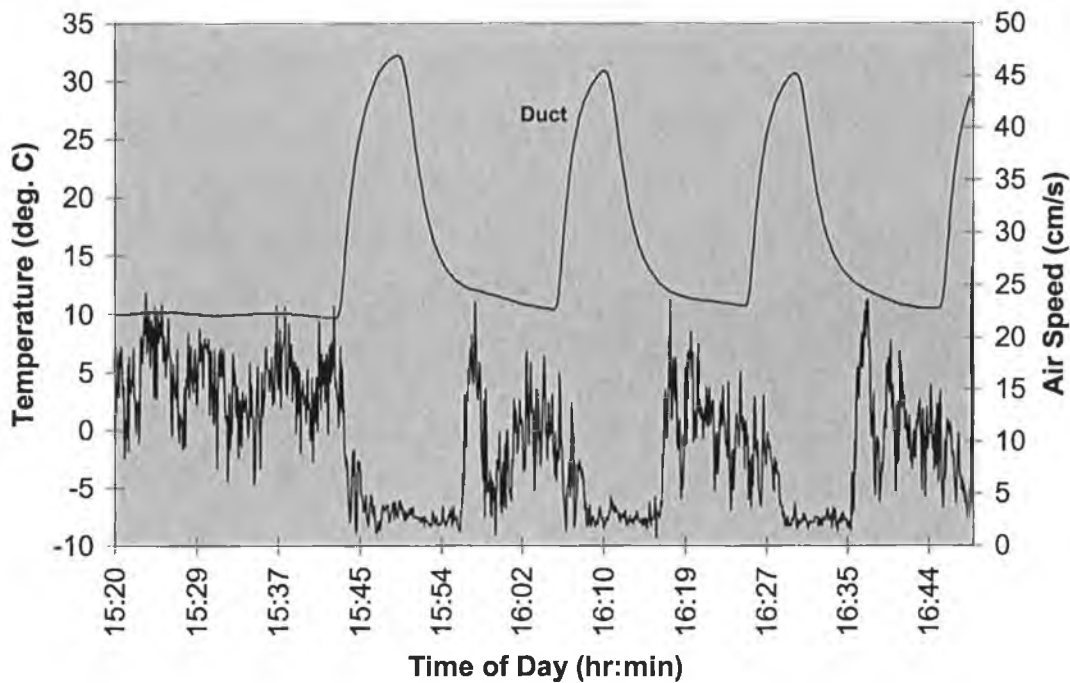
Figure 2.17. Air speed and temperatures during dehumidification.

The average cropping surface air speed was reduced from 15.7 cms^{-1} to approximately 7.1 cms^{-1}. The consequence of this was that the maintenance or improvement of drying that was the object of the exercise was inhibited.

It had been normal in the construction of tunnels to position the exhaust outlet near the roof of the tunnel, above the main doors. It can now be seen that this position is not desirable because, during heating, warmed air would be leaving the tunnel and this would increase the on time for heating. The use of low-level air exhausts would be appear to be preferable in terms of energy conservation and improved airflow.

Strategies such as positioning the temperature sensor at some height in the upper re-circulating layer, using a high boiler water temperature to minimise the on time of the heating system and correct positioning of the exhaust outlets all help to increase the average air speed during cropping but, at some point as the heating load increases, there will be periods of poor airflow and an alternative air distribution system is required to provide a consistent cropping surface air speed. Such a system is discussed in section 2.6.

It may be that recovery from the stratification could be minimised by operating at higher fan speeds but this could damage the crop by causing excessive drying. In principle, a computer-based controller could be programmed with the time delays for the reduction in airflow and the recovery time after heating and then the fan speed could be controlled dynamically. After stratification was present, the fan speed could be increased to the maximum available and then reduced as soon as mixing began again. Ideally, a low-cost sensor suitable for the low speed airflow at the cropping surface could be used to close a control loop with the fan speed controller.

2. 6   An alternative distribution system

While the simple procedures to improve the airflow were easily implemented and produced some improvements, some more radical alteration to the air-handling

arrangements was needed to produce a means of supplying a constant airflow at the crop. One suggestion involved the introduction of a second distribution duct at low level, i.e. near the crop. Following dissemination of the findings on the effect of heating on airflow in mushroom tunnels (Grant, 1995), this had been implemented in a commercial mushroom tunnel by altering the air-handling unit so that conditioned air was delivered to the low-level duct. The original duct was disconnected leaving only a fan to drive air through it and this, therefore, had a re-circulation function only. In the commercial operation a mixing valve controlled the heating and the operator fixed the fan speeds manually. This allowed some measure of compensation for the heating effects and it was found that with the re-circulation duct providing some of the momentum for the airflow that cropping surface air speeds were relatively unaffected for low heat demand but that the flow stratified as the heating load increased. The introduction of a second duct was, however, considered to be useful because there was a reserve of momentum for the flow field that was available from the additional fan and distribution duct and further investigation and development were undertaken.
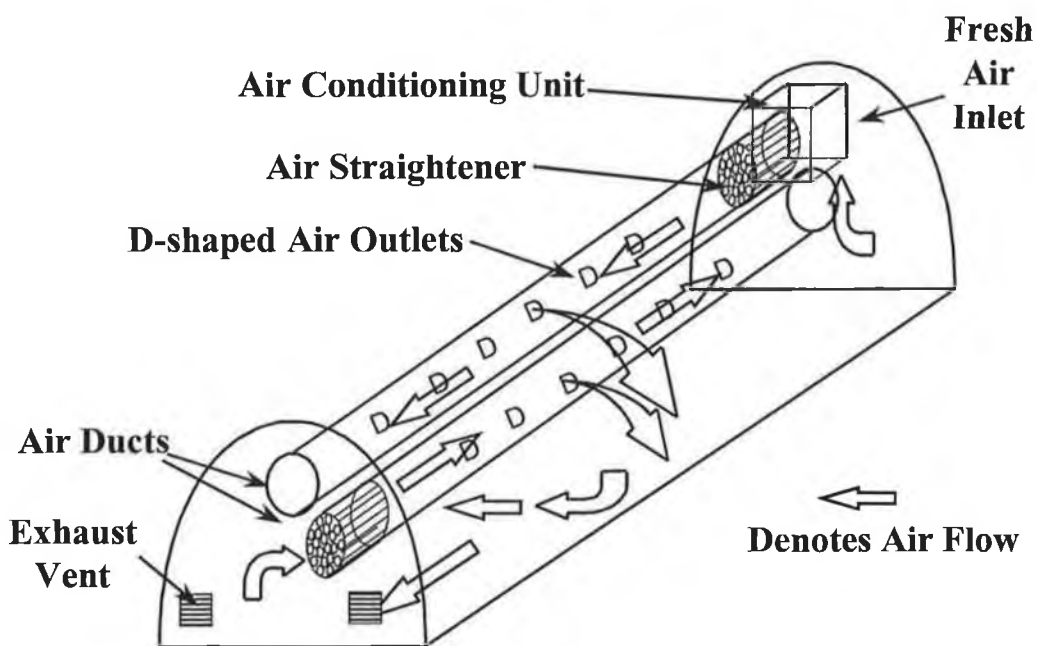
Figure 2.18. Schematic diagram of a two duct air distribution system.

Experimentation showed that it was not necessary to have a low-level duct and that it was not necessary to alter existing air-handling arrangements. It was found that a second duct, for re-circulation only, could be added just below the existing one, i.e. clear of the crop and normal cropping operations (figure 2.18), and this was found to have the potential to counteract the effects of thermal buoyancy on the cropping surface airflow. This approach, then, was entirely a retrofit operation with no changes required in the existing air distribution system.

The new air distribution system for this work consisted of two polyethylene ducts. One of them delivered the conditioned air as usual while the other contributed to the re-circulation and air distribution only. This latter duct was installed so that the fan was at the end of the tunnel opposite to the existing one and this counteracted temperature gradients that could arise when the fresh air intake was closed (Grant, 1995). There were thus two sources of momentum for the flow field. Figure 2.19 shows an elevation of the current arrangement. The figure indicates the general pattern of the flow field and the influence of the two ducts as currently understood. Both sets of air jets drive the air circulation that flows across the cropping surface.
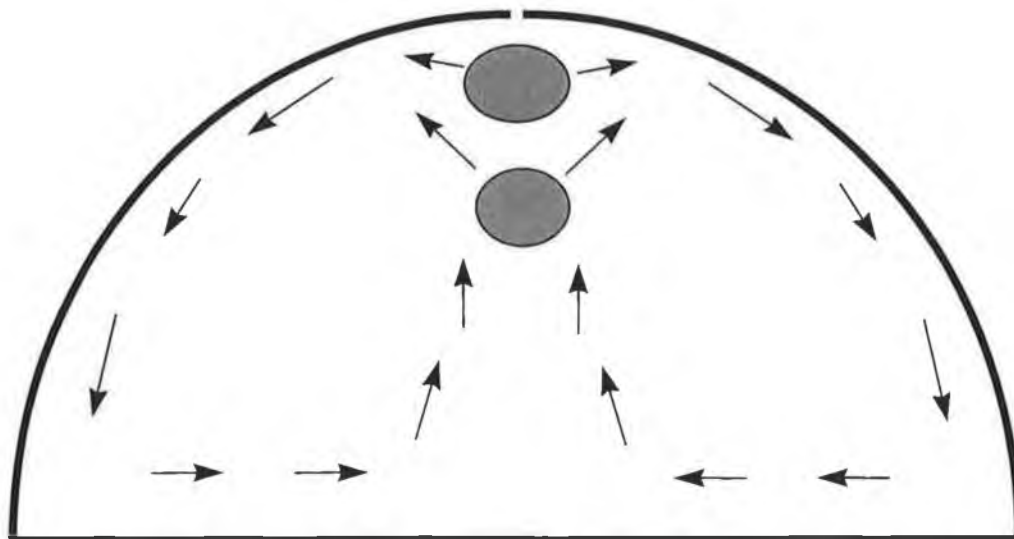
Figure 2.19. The general flow pattern from a two-duct air distribution system.

Measurements were made of exit air speeds at the distribution ducts to relate mean cropping surface speeds to combinations of duct exit speeds. The results given in table

2.1 indicate the main characteristics of speed control in the system for isothermal flow.

The initial settings of 2 ms^{-1} at the conditioned air distribution duct (duct 1) and 3 ms^{-1} at the new re-circulation duct (duct 2) produced a mean cropping surface air speed of 10.7cms^{-1}. The contribution from duct 1 was increased by raising the air speed to 3 ms^{-1} and the cropping surface speed then increased to 17.9 cms^{-1}. The next increase was supplied from duct 2 by increasing its output to 4 ms^{-1} and, again, the air speed at the crop increased. This rise to 25 cms^{-1} was very similar in magnitude to the previous one. These three sets of data demonstrate that either duct can be used to control cropping surface air speed.

A further interesting observation is that if the speeds from ducts 1 and 2 are interchanged, as in the case of the last two rows of table 2.1, then the resultant cropping surface air speeds are very similar. While further definition of the operating characteristics is required, it is at least clear that the flow field can be manipulated by either of ducts 1 and 2 and that their contributions to the momentum of the flow field are of the same order of magnitude. The system output therefore is, to some extent, predictable and controllable.

Table 2.1. Cropping surface air speeds corresponding to a number of exit air speeds at the distribution ducts.

Conditioned air exit speed (ms^{-1})	Re-circulation exit speed (ms^{-1})	Cropping surface air speed (cms^{-1})
2	3	10.7
3	3	17.9
3	4	25.0
4	3	25.7

Tests were carried out where the stratification was induced by starting the heating and then the ability of the system to restore cropping surface air speeds to isothermal levels was examined. Figure 2.20 illustrates the results of such a test. The graph shows the sequence of events as the heating was turned on and the stratification developed. An air speed of 30 cms^{-1} (6.2 ms^{-1} duct exit air velocity) was set initially with a duct temperature of 12.5 $^{\circ}$C and then heat was applied as can be seen by the sharp rise of the air temperature in the distribution duct. Air speeds at the crop fell at this time due to stratification of the flow field. After a period of this minimal flow the fan in the additional, re-circulation duct was turned on to provide an exit speed of 4 ms^{-1} and the effect of this is seen in the immediate rise in the air speed recorded at the crop. After a further period flow from the additional duct was stopped and the cropping surface air speed returned to its initial value and stratification was re-established.
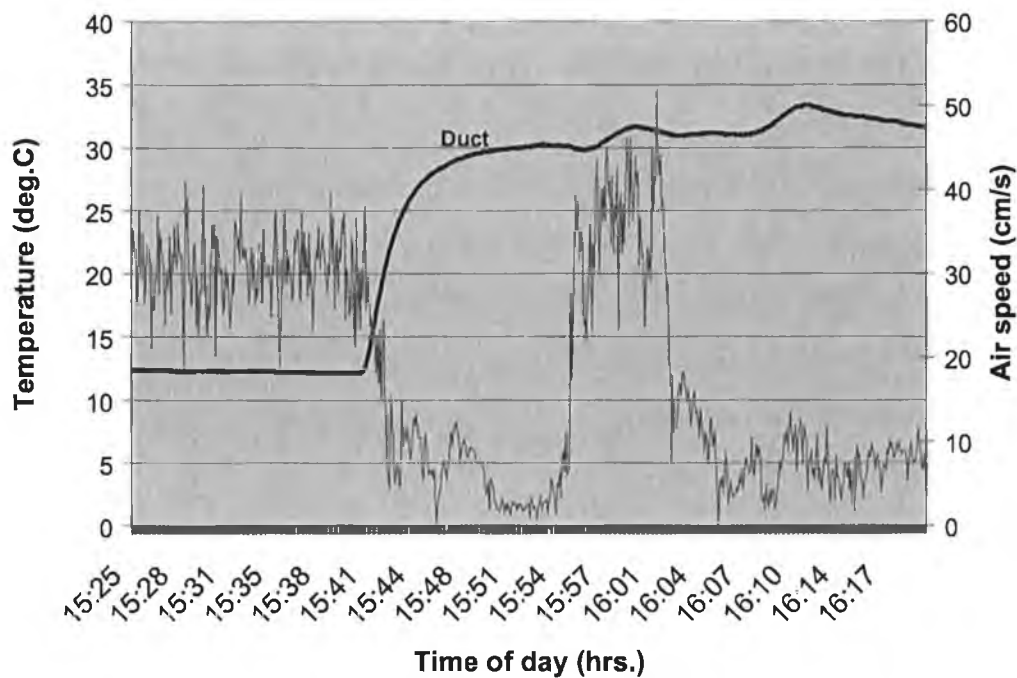
Figure 2.20. Air speed and distribution duct air temperature for a two duct air distribution system

Further tests showed that the restored air speed was not steady but decreased slowly over time and therefore the results could be taken only as initial indications of the potential of the system. More work is needed to establish clearer relationships between heating effects and the ability to restore cropping surface air speed.

The general result is that the new system has the capability to maintain air speeds at desired levels and to avoid the potentially serious consequences of a period of low air movement due to a high heat demand. Taking advantage of the potential of the two duct system depends on having dynamic control of the speed of both fans. For example, in a simple implementation for the widely-used on-off heating control it should be possible to apply a simple increase in the speed of the re-circulation fan whenever the heating is on. Taking full advantage of the additional distribution duct, however, is likely to require a feedback control approach. It was observed that under prolonged heating, a slow reduction occurred in the recovered air speed. Ideally a low-velocity air speed sensor would be used to close a feedback loop and to allow consistent air speeds to be maintained at all times. In a controlled situation the duct exit air speeds would be adjusted to ensure that the stratification would not be allowed to arise.

The addition of a second distribution duct solves, in principle, the problem of inconsistent airflow at the mushroom cropping surface. It also brings other possibilities for the improvement of the ventilation in the tunnels.

i.   The range of speed available is greatly increased because the two ducts drive the air stream together.

ii.   The conflict between the need to be able to reduce cropping surface air speed and yet maintain a high volume delivery for carbon dioxide control could be resolved because it is possible to separate the functions of air conditioning (heating, cooling, fresh air control) and the supply of momentum to the airflow in the tunnel. It is possible to design the conditioned air duct to deliver low speeds with high volume while the circulation duct could be designed to give a

wide range of speeds. The combination could maintain cropping surface air speeds while maximising the volume flowing through the heat exchangers.

iii.   The flow directions in the two ducts can be opposite to counteract temperature and humidity gradients that can arise along a tunnel

iv.   A second duct positioned above or below the first can be extended under or over the existing air-handling unit thereby providing airflow in an otherwise under-ventilated area.

2. 7   Airflow for two rows of three-level shelving

At the time of writing, the mushroom industry in Ireland is changing, with a large number of new growing systems being developed and evaluated. There is great interest in the automation advantages of shelf growing. A crucial question that has to be answered is whether or not the system can produce mushrooms with the quality and shelf life that the retailers require. The microclimate at the cropping surface is of great importance in improving the quality of shelf-grown mushrooms (M. Connolly, 1999) and the main problem to be solved is the production of a uniform and controllable airflow at all points. Loeffen (1992) and Bowman (1991) have studied air flow for shelving systems with 4 or more levels of compost in square-section rooms but no information was found that referred to 2 and 3 level systems in tunnel structures with their curved walls.

The proliferation of new multi-level systems of growing makes the use of purely empirical methods undesirable because of the time required to investigate each system and then to produce a method of providing uniform airflow by trial and error. Two preliminary investigations of methods of altering the airflow pattern in a tunnel are presented here. These methods produce some benefits for the particular growing systems examined but they would have to altered for other systems with, say, different shelf dimensions and different numbers of rows of shelving, and the process of

optimising them would involve a good deal of experimental work in each case. Once validated for a small number of systems, a flow model could allow the testing of options for other systems and would save some of the time required for measurement.

Two methods of correcting air flow to improve the uniformity of airflow for shelf growing have been examined and some encouraging results were obtained. These were as follows:

    i. an alternative configuration of the apertures in the distribution duct and

    ii. the deflection of air flow.

2.7.1 Duct configuration

As described in section 1.2, air normally leaves the distribution duct from two sets of apertures located at 2 and 10 o'clock positions on the distribution duct. An alternative duct configuration to supply air to two rows of shelving was suggested to a grower by members of the IPP Plant (Dutch Mushroom School) and this was implemented for testing. The main flow was directed downwards between the two rows of shelving and a lower volume (fewer and smaller holes) was directed horizontally from both sides of the duct. The general airflow pattern is indicated in figure 2.21. The intention was that the flows, one upward from the floor having passed under and across the lowest shelf and the other travelling down the wall in the conventional manner, would meet at the level of the middle shelf and that this should cause air to move across the middle shelf.
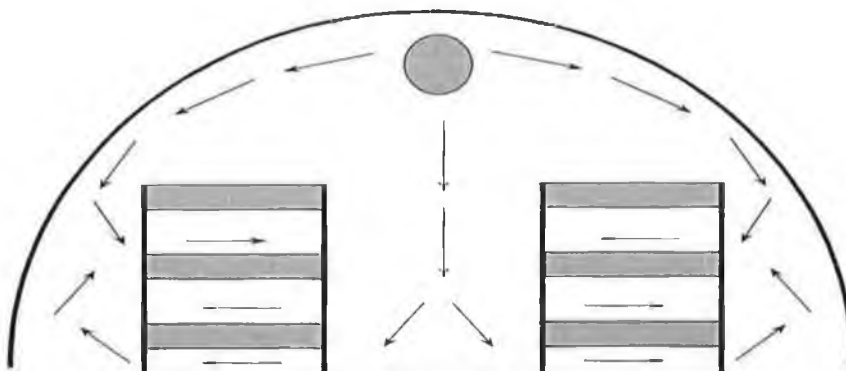

Figure 2.21. Air flow pattern from a modified duct.

49

The general airflow pattern was verified using smoke and measurements of air speeds were made to determine the effectiveness of the approach. It was important that the shelves should be positioned far enough away from the wall to avoid the downward flow striking the edge of the shelf as this would result in excessive speeds across the top shelf. The shelves in this case were 1.4 m wide and there were three levels with their bases at 35 cm, 95 cm and 157 cm above the floor of the tunnel. The shelves had a depth of 21 cm. The horizontal clearance between the top of highest shelf in each row and the sidewalls was 90 cm and the vertical clearance was 88 cm. The tunnel was 7 m wide at the base.

The air speeds were measured in the centre of each level of the shelving and all the measurements were made near the end of the tunnel distant from the fan. The variation in the air speeds across a given shelf would have to be thoroughly evaluated but for the purposes of this experiment a number of the measurements were made at other points across the shelf and these indicated that the speeds recorded at the centre were a reasonable representation of the general airflow pattern. With only one air speed sensor available the work of measuring the air speeds was time consuming because frequent checks had to be made to ensure that the airflow pattern was stable. As before, the measurements were made under conditions of low solar radiation and the heating was off.

The figures in the tables below are the averages over 10 minutes of the data collected. Level 1 refers to the shelf nearest to the ground and level 3 is the topmost one. The results in table 2.2 are for a duct with apertures balanced so that approximately 50% of the total flow was downward with 25% to each side at the maximum air speed from the duct. The speeds at all three levels were at acceptable values for mushroom growing but there was considerable non-uniformity in the results with the speed across the middle shelf being almost double that across the bottom one. The importance of the flow under the shelving was checked by inserting a barrier under the shelves. Air speeds on the two upper shelves dropped to very low levels. Further experiments with the duct configuration altered the distribution of the airflow but failed to improve the uniformity of the air speeds.

50

Table 2.2. Air speeds on shelves using altered duct configuration (7.5ms^{-1} at duct).

Shelf	Air speed (cms^{-1})
Level 3	19.6
Level 2	25.0
Level 1	13.2

Some factors that need to be examined would include the positioning of apertures on the distribution duct, the positioning of the shelves in the tunnel and the spacing between the shelves. Examining all the options experimentally would be very time-consuming and tedious and a model would be very useful for studying the effect of these factors.

2.7.2 Deflectors

The interruption of the airflow pattern using flat deflecting plates fixed to the tunnel wall, as assessed as a means of providing more uniform air speeds across the three levels of shelving in a mushroom tunnel. The measurements were carried out in a 9.2 m wide tunnel equipped with two rows of shelving located symmetrically within the tunnel (figure 2.22). The shelves were 1.4 m wide and there were three levels with their bases at 35 cm, 95 cm and 157 cm above the floor of the tunnel. The depth of the shelves was 21cm. The vertical clearance between the top shelf and the wall/roof of the tunnel was 1.2 m. Horizontal clearance was 1.0 m. The measurement approach used was as outlined in section 2.7.1.

The duct in use was of the conventional type with holes at 2 and 10 o'clock. The air speeds were measured on each level with no flow correction and the results are given in table 2.3. There was extreme non-uniformity as the general pattern of airflow was

similar to that found in single-layer growing where the main flow was across the bottom shelf which corresponded approximately to bag height. The pattern, as indicated by visualisation, is shown in figure 2.22. The average air speed recorded at level 2 was 50% of level 1, while the topmost shelf had an air speed that was only 28% of level 1. This had to be corrected to optimise the production efficiency.

Table 2.3. Air speeds for uncorrected air flow ($6ms^{-1}$ at duct).

Shelf	Air speed (cms^{-1})
Level 3	8.9
Level 2	15.9
Level 1	32.0

The main airflow was passing the upper and middle shelves and methods of redirecting flow to those shelves were examined. A means of interrupting the downward flow was required. Experimentation began by trying to interrupt the airflow at the wall and to force it across level 3. The idea behind deflection is illustrated in figure 2.23. As the flow travelled along the wall a partial block was put in its path to break the attachment of a portion of the airflow.
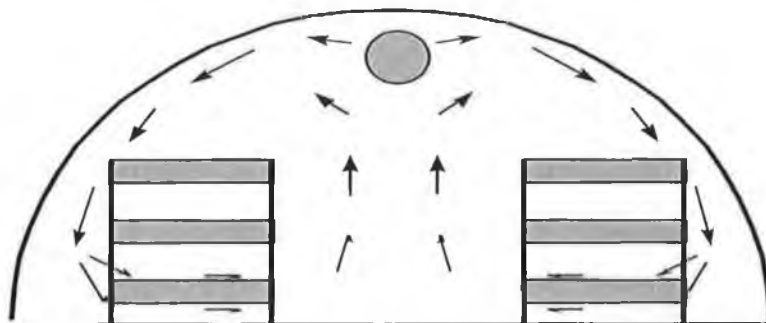
Figure 2.22. The uncorrected air flow pattern
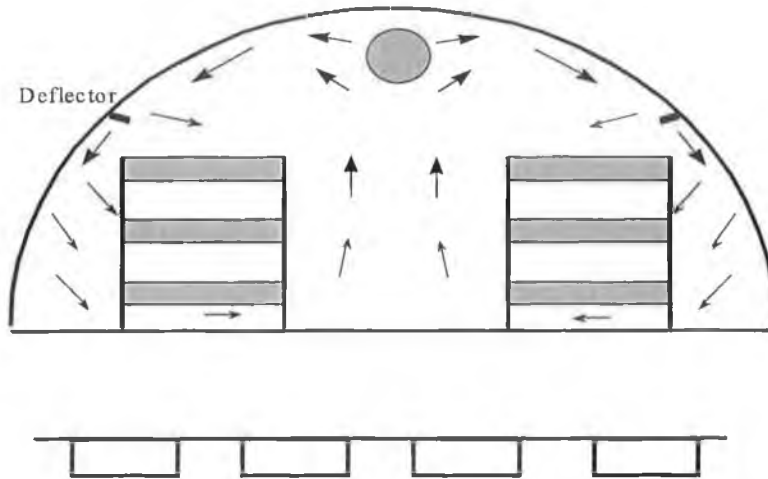
Figure 2.23. A modified air flow pattern.

The blockage or deflector defined a new path for this portion of the air flow and in this case it was 'aimed' at the topmost shelf. In general, a portion of the wall flow was allowed to carry on uninterrupted so that a typical plan view of deflectors is as shown at the bottom of figure 2.23. A series of experiments with the depth, angle and spacing and height of the deflectors showed that it was possible, with the right combination, to control flow across all three shelves.

Tables 2.4 to 2.6 contain the results of some of these experiments. For these sets of results the deflectors were set at an angle of 30° below the horizontal. The angle of the deflector and its height above the shelves was important in producing a suitable airflow across the shelves. A deflector angle and position that allowed the flow to play directly onto the shelves produced excessive air speeds at the near edge of the shelf and the mushrooms were severely damaged. For the results given below the deflector was mounted at a point on the wall that was 0.6 m above the upper shelf. Thus, the purpose of the deflectors was not to direct the airflow onto the shelves but to change the flow pattern in such a way that the new pattern induces a suitable flow across the shelves.

Table 2.4. Air speeds with 24 cm deflectors and 100% blockage.

Shelf	Air speed (cms^{-1})
Level 3	15.4
Level 2	17.8
Level 1	20.0

Table 2.5. Air speeds with 29 cm deflectors and 100% blockage.

Shelf	Air speed (cms^{-1})
Level 3	21.5
Level 2	11.8
Level 1	14.2

Table 2.6. Air speeds with 29 cm deflectors and 75% blockage.

Shelf	Air speed (cms^{-1})
Level 3	18.8
Level 2	18.5
Level 1	18.8

The sequence of actions was as follows. Table 2.4 refers to the measurements made with a continuous deflector with no gaps (100% interruption of the airflow) and a depth of 24 cm. The results in the table indicate that a greater interruption of the wall-attached airflow was required because more flow and a higher air speed was required for level 3. It should be noted that there was a small increase in the airflow at level 2 relative to the results for the uncorrected flow. Increasing the depth of the deflectors to 29cm provided the required boost for level 3, as shown in table 2.5. Unfortunately, this was achieved at the cost of reduced flow at the lower levels. Introducing gaps in the deflectors allowed some of the air to travel uninterrupted along the wall. With the

width of the deflectors set to 30 cm and gaps of 10 cm, i.e. 75% interruption, sufficient flow was restored to the lower levels to produce the uniform flow pattern of table 2.6.

The uniformity of this example demonstrates the control that can be achieved with careful adjustment. Further experiments have shown that, while it may be possible to optimise any given situation, the uniformity of the flow for a given configuration of the deflectors depended on the duct air speed and the distance of the shelves from the wall. In all cases, however, uniformity of air speeds was greatly improved compared to the uncorrected airflow pattern and it could be maintained over a range of air speeds. Very similar results were obtained for the narrower tunnel described in section 2.8.1. In that case the best uniformity was achieved with a the same deflector dimensions as those for table 2.5 but the gaps were set to 30 cm, i.e. 50% interruption and the deflectors were mounted 63 cm above the top shelf at an angle of $15°$ below the horizontal

2.8 Airflow for three-level growing systems in three rows

The wide tunnels (8.8 m and 9.2 m) used in Ireland can accommodate a third row of shelving if it is placed along the centre line of the tunnel. This, unfortunately, is the least efficiently ventilated area of these tunnels using the conventional system of a central duct with exit holes at 2 and 10 o'clock and the air delivery system therefore needs to be altered if positive air flow is to be provided across the extra row. This would involve some changes to the air distribution system but finding solutions for particular cases would be very time-consuming. It is difficult to predict the airflow pattern that will result from a new system and experimentation is essential. This is a problem where a flow model would be very beneficial because each new configuration would require measurement at a number of duct exit speeds and that some adjustment of the positions of the distribution ducts

Two examples of novel air distribution systems are given below. These were installed on commercial holdings and both offer improved airflow across the central row of shelving but with considerable non-uniformity. They are included here as examples of the increasing complexity of the airflow distribution problem for mushroom growing. The task of optimising or re-designing these would be very difficult without a modelling approach.

The first installation is simply the wide (9.2 m) tunnel from section 2.8.2 with an extra row of 1.4 m wide shelving placed in the centre of the tunnel. In an attempt to improve airflow in the centre of the tunnel, the output from the air-handling unit was split and sent to two distribution ducts suspended near the roof but over the aisles between the rows of shelving. Both ducts had the standard layout of air exit apertures at 2 and 10 o'clock positions on the ducts (downward directed flow had been a failure). The flow geometry is shown in figure 2.24.
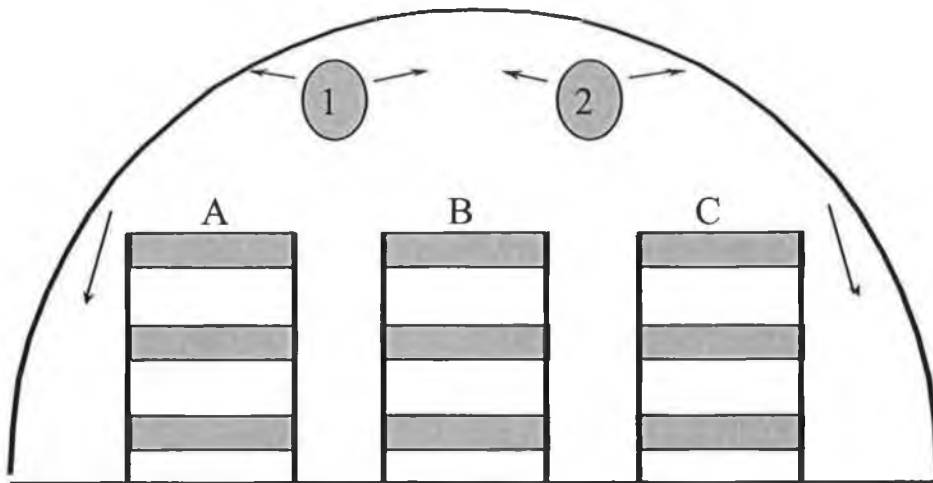

Figure 2.24. Two ducts with three rows of shelving.

The exit speeds at the distribution duct were 6.4 ms^{-1} for duct 1 and 6.6 ms^{-1} for duct 2 (no balancing was possible) and the resultant air speeds at the cropping surface are presented in table 2.7. Air speeds were measured in the centre of the shelves and the shelves were labelled for identification with the tables.

Table 2.7. Air speeds with two conventional ducts.

Shelf	Row A (cms^{-1})	Row B (cms^{-1})	Row C (cms^{-1})
Level 3	8.8	10.6	15.2
Level 2	9.8	18.9	12.7
Level 1	19.6	20.5	26.0

It is clear from the results that air was bypassing the top and middle shelves in much the same way as it would with a conventional duct alone. The air speeds on level 1 were the highest. It is possible that this could be corrected by the use of deflectors to give a uniform airflow over all three levels.

In different growing system, measurements were made in a wide tunnel with three air distribution ducts supplying air to a system comprising three rows of staging with three levels of compost in shallow, plastic-wrapped blocks. The distribution ducts consisted of the original, standard one that supplied conditioned air and two extra ones mounted either side of this. The extra ducts were re-circulating tunnel air. The original distribution duct was in the standard configuration with air outlets at approximately 2 and 10 o'clock positions and the extra ducts supplied air downward towards the floor. The height of the outlets in these ducts was 2.2 m above the floor.

Figure 2.25 shows a schematic of the layout in the tunnel and the labelling for reference to the tabulated results below. The arrows show the directions that the air flows on leaving the ducts. The lowest and highest beds of compost were 1.2 m wide and the middle one was 0.95 m. The depth of each bed was 18 cm and the bases of the beds were at 15 cm, 70 cm and 120 cm above the floor. The horizontal clearance between the top levels and the sidewalls was 1.1 m and the vertical clearance was 1 m.
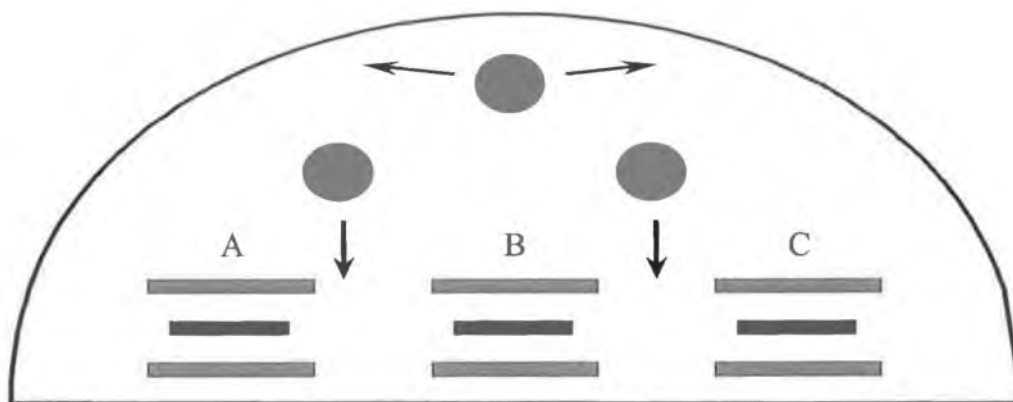
Figure 2.25. Overall system schematic.

Initially, airflow through the extra ducts was turned off and air speeds were measured so that any benefit of the introduction of the extra ducts could be assessed. Level 1 is the lowest bed of compost. Measurements were made at the centre of each level of compost with the duct exit air speed set to the maximum of 6.2 ms^{-1} and results are in table 2.8.

Table 2.8. Air speeds with conventional duct only.

Shelf	Row B (cms^{-1})	Row C (cms^{-1})
Level 3	5.8	16.4
Level 2	12.9	17.1
Level 1	8.0	27.9

Table 2.8 demonstrates that the airflow with the conventional duct follows the usual pattern of travelling down along the tunnel wall and moving across the bottom of the tunnel. The highest speed recorded was for the lowest bed. Also, as in bag growing, there was an area of low speeds near the centre of the tunnel and this was due to the air streams from the two sides of the tunnel meeting in the centre.

The fans in the additional distribution ducts were then turned on and after allowing a long interval for the new flow to become established, the measurements were repeated for the new flow configuration. The exit speeds from the additional distribution ducts were 7.7 ms^{-1}.

Table 2.9. Air speeds with conventional duct and extra ducts.

Shelf	Row B (cms^{-1})	Row C (cms^{-1})
Level 3	22.2	31.9
Level 2	17.9	17.8
Level 1	24.8	17.8

The results are presented in table 2.9. The air speeds can be seen to have improved considerably with flow suitable for mushroom growing (greater than 10cms^{-1}) at all points. There was non-uniformity, however, and the air speed recorded at level 3 in row C, for example, was almost double the air speed at the levels below it. The air speed across level 1 in row C was reduced compared to the results in table 2.8 and this was probably due to the airflows from the conventional duct and the extra ducts opposing each other at this level. There is considerable non-uniformity but limited experimentation with the fan speed control suggested that it might be possible to improve the uniformity by changing the balance between the outputs from new ducts and that of the original one. This configuration should have some of the benefits of the two duct system for overcoming stratification (Section 2.7) in that it provides some sources of flow that are independent of the effects of the heating system.

The number of variations in the cropping systems and the variations in air distribution make the empirical approach to optimising airflow in mushroom tunnels very time-consuming and tedious. There would be little information gained that could be applied to a new situation. The use of flow models and computational fluid dynamics would be a useful part of a problem-solving framework. If a flow model could be validated in a number of situations then the programmes could be used to guide the

experimentation for new situations and could drastically reduce the time required to provide the systems that would deliver a suitable airflow at all times.

Chapter 3

The Mathematical Model

3. 1   The equations of fluid flow

The equations that are used to describe fluid flows for the purposes of this thesis share a common form in that they all obey a generalised conservation principle, i.e. there is a balance between the factors that influence a given dependent variable. These are discussed in the book by Patankar (1980) and, if the dependent variable is ϕ, then

$$\frac{\partial}{\partial t}(\rho\phi) + \nabla \cdot (\rho\tilde{u}\phi) = \nabla \cdot (\varGamma\nabla\phi) + S \tag{3.1}$$

where $\varGamma$ is the diffusion coefficient, S is the source term, t is time, ρ is the density and $\tilde{u}$ is the velocity vector.

Reading equation (3.1) from left to right, the terms are the time dependent term, the convection term, the diffusion term and the source term. The dependent variable ϕ will be the velocity components, temperature and turbulence parameters. The diffusion coefficient $\varGamma$ and the form of the source term S are different for different dependent variables and an appropriate meaning has to be given in each case. Expressing any given equation in the general form by manipulation may result in extra terms that could be included in the source term for the purposes of the solution process.

The other constraint on the flow fields is the continuity or mass conservation equation, given by

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \, \widetilde{u}) = 0. \qquad (3.2)$$

For the isothermal air flows considered in this thesis, density changes are taken as negligible (Patankar, 1980) and the continuity equation reduces to

$$\nabla \cdot \widetilde{u} = 0 \qquad (3.3)$$

3.2 A model for the description of turbulence

Observation has shown that the flow in mushroom growing tunnels is turbulent in nature and that modelling of the airflows involved would require that the effects of turbulence be calculated. There is much general work in the field of turbulent flow calculations and many publications describe work that makes use of techniques similar to those used in the TEACH code. The standard k-ε turbulence model (Launder and Spalding, 1974) which is included in the TEACH package is taken as the means of describing turbulence for the purposes of this thesis. There are a number of modifications that can be made to the details of this modelling (Patel et al., 1985) but until it is possible to carry out some model validation in the mushroom growing room application there is little basis for changing from this formulation of the k-ε equations. For example, Rahnema et al. (1996) implemented a modification in the ε equation for re-circulating flows whilst Liu et al. (1996) compared three k-ε models for predicting the ventilation air jets, both plane-free and plane-wall. They found that the simulations on a number of different grid densities predicted the fluid velocity decay and velocity profile well but over-predicted the jet spread and entrainment ratios by 20 to 40%. A modified turbulence model may ultimately be required for the jets and flow along walls that are part of mushroom tunnel ventilation.

Patankar et al. (1977) presented the numerical prediction of the three-dimensional velocity field of a deflected turbulent jet. This study incorporated the hybrid differencing (see Chapter 4) scheme and a line solution technique making use of a sequential/segregated approach to an iterative scheme. Patankar et al. (1979) modelled the turbulent flow in internally finned tubes and annuli using a mixing-length model of turbulence while Patankar et al. (1978) calculated thermal interactions among the confining walls of a turbulent re-circulating flow. In this latter problem, they used the k-ε model for turbulence and that model is used in the present work.

Rodi (1980) described in detail the various turbulence models and their applications and evaluated the models with regard to their predictive capability and computational effort. Nallasamy (1987) presented a brief account of various turbulence models employed in the calculation of turbulent flows, and evaluated the application of these models to internal flows by examining the predictions of the various turbulence models in selected flow configurations. Among the main conclusions were that the k-ε model is used in the majority of all the 2-d flow calculations reported in the literature and that modified forms of the k-ε model improve the performance for flows with streamwise curvature and heat transfer. Two important factors in the numerical solution of the model equations, namely false diffusion and inlet boundary conditions are discussed (see chapter 4).

Numerical calculations were presented by Hjertager and Magnussen (1981) for two jet-induced three-dimensional flows in rectangular enclosures. The computational method used was very similar to that used in this thesis. Results were compared with available air velocity data and were in general good, although some discrepancies were found. Choi et al. (1988 and 1990) applied the k-ε model in the prediction of the air flow in a slot-ventilated enclosure and for flow around rectangular obstructions. They modified the TEACH code to predict two-dimensional, isothermal air flow patterns and air velocities. Air distribution patterns, velocities, jet growth and attachment, and entrainment predictions were found to agree well with published data. For the obstructions in the flow field, it was found that inaccuracies could arise in the calculation

of the air velocities downstream of the obstructions but it was noted that more grid refinement than that used was likely to be required for accurate predictions.

Maghirang and Manbeck (1993) studied a slot-ventilated airspace that was very similar to the general flow situation in mushroom-growing rooms. They modelled the airflow and the transport of neutrally-buoyant bubbles using the k-ε turbulence model. Comparison between the numerical simulation and experimental results showed good correspondence in the velocity fields and bubble trajectories.

3.3   Approximation of the exact equations

The equations for turbulent flow have their origin in the conservation laws and these give rise to the exact description of a turbulent flow (Tennekes and Lumley, 1972 and Rodi, 1980). The computing resources required for solving the exact equations are large due to the scale of the turbulent elements in a flow relative to the extent of flow domains. For a direct numerical solution (DS), working at scales approximately 10^3 times smaller than the flow domain, at least 10^9 grid points would be required in three dimensions. While this approach provides insight into the nature and mechanisms of turbulence, this sort of computing demand has acted as a driving force in the development of models where a statistical approach is taken to make their solution more practical.

Many of the features of DS can be retained while computing higher Reynolds number flows in the large-eddy simulation (LES) technique. Using a space-filtering procedure, the mean and the large-eddy fields are resolved but the fields smaller than the sub-grid scale have to be mathematically modelled. The instantaneous field ϕ can be decomposed as

$$\phi = \hat{\phi} + \widetilde{\phi} \qquad (3.4)$$

where $\hat{\phi}$ is the filtered (resolved) field, while $\tilde{\phi}$ is the residual (unresolved) fluctuation. Substituting this decomposition into equation (3.1) results in an equation that is formally similar but containing additional terms arising from the filtering procedure. A model is then required to relate the extra terms to the resolved field.

LES, however, is still computationally demanding and the mainstay of engineering applications is a time-averaging technique. Reynolds decomposition is the technique applied and in this the instantaneous values of the dependent variables in the conservation equations (velocity component U_i, pressure P and the scalar quantity ϕ) are separated into mean and fluctuating quantities

$$U_i = \overline{U_i} + u_i \quad , \quad P = \overline{P} + p \quad , \quad \phi = \overline{\phi} + \varphi \tag{3.5}$$

where, for example, $\overline{U_i}$ is a mean whose value is determined over a time that is long compared to the turbulence time scale. The equations and an equation of state (relating, say, the density to temperature) no longer form a closed set because of the introduction of unknown correlations between the fluctuating quantities:

$$\sigma_{ij} = -\rho \overline{u_i u_j}, \qquad q_j = \rho \overline{u_j \phi_j} \tag{3.6}$$

which represent the turbulent transport of momentum and heat or mass. In most flow regions, the turbulent stresses and fluxes are much larger than their laminar counterparts, which are therefore often negligible.

The equations containing the unknown correlations can be solved only if a means of determining these can be constructed. Exact derivations of equations for these terms contain higher-order unknown correlations and therefore models are introduced that approximate the correlations (Tennekes and Lumley, 1972). The model simulates the averaged character of real turbulence and forms a closed set with the mean flow equations. A turbulence model could be defined as a set of equations that determine the turbulent transport terms in the mean-flow equations and thus close the set of equations.

To be useful, the empirical constants in such equations, once defined, should apply over a reasonable range of flow situations.

Turbulence is an eddying motion with a wide spectrum of eddy sizes and a corresponding spectrum of fluctuation frequencies. The largest eddies are determined by boundary conditions in a flow domain and the smallest by viscous forces. It is mainly the large-scale turbulent motion that transports momentum and heat and contributes to the turbulence correlations. The large eddies interact with the mean flow and extract energy which is passed down a cascade of turbulent scales until it dissipated at the level of viscous forces. An important part of practical turbulence modelling is the eddy-viscosity due to Boussinesq (Rodi, 1980) which assumes that eddies are objects that can exchange momentum (as for molecules) and while this assumption breaks down easily because eddies do not retain a discrete identity the idea has been found to work well in practice.

Using this idea, the Reynolds stresses contained in equation 3.6 are assumed to be proportional to the mean velocity gradients. For the general flow situation this can be expressed as

$$-\overline{u_i u_j} = v_t \left(\frac{\partial U_i}{\partial x_j} + \frac{\partial U_j}{\partial x_i} \right) - \frac{2}{3} k \delta_{ij} \tag{3.7}$$

where the Kronecker delta δ_{ij} is necessary to make the expression applicable also to normal stresses but, when making this substitution for the Reynolds stress in the momentum equations, the k term can be absorbed by the pressure-gradient term and the problem is now shifted to the determination of v_t.

In this concept, the eddy viscosity, v_t, is proportional to a velocity scale and a length scale that characterise the large-scale turbulent motion because it is the distribution of these scales that can be approximated in many flows.

3.4 The k-ε model

The details of the determination of v_t give rise to a family of turbulence models of increasing complexity.

(i)   Zero-equation models

These are models that do not involve transport equations for turbulence quantities and are relatively simple descriptions of the eddy viscosity based on empirical formulae or by relating it to the mean-velocity distribution.

The first model to describe the distribution of the eddy viscosity is known as the Prandtl mixing-length hypothesis and, using an analogy from kinetic gas theory, this can be expressed as

$$v_t = l_m^{\,2} \left| \frac{\partial U}{\partial y} \right| \qquad (3.8)$$

which relates the eddy viscosity to the local mean-velocity gradient and involves a single unknown parameter, the mixing length l_m. This concept can be applied successfully for relatively simple flows where l_m can be specified by a simple empirical formula. For more general flows, the hypothesis can be extended from the two-dimensional case but little use has been made of it because it is difficult to specify l_m in flows more complex than shear layers and, while there have been successful flow calculations, they are generally of restricted application.

(ii)   One-equation models

In order to overcome the limitations of the mixing-length hypothesis, models were developed that account for the transport of turbulent quantities by solving a differential transport equation.

If velocity fluctuations are to be characterised by a scale, the physically most meaningful is $\sqrt{k}$ where k is the kinetic energy of the turbulent motion (per unit mass) because k is a direct measure of the intensity of the turbulence fluctuations in three directions. As the energy is contained mainly in the large-scale fluctuations this is a velocity scale for the large-scale turbulent motion and, with the eddy viscosity concept, comes the Kolmogorov-Prandtl expression (Rodi, 1980)

$$v_t = c'_\mu \sqrt{k} L \tag{3.9}$$

(where L is the length scale of the large scale fluctuations) and a transport equation for k where the rate of change of k is balanced by the convective transport due to the mean motion, the diffusive transport due to velocity and pressure fluctuations, the production of k by interaction of stresses and mean-velocity gradients and the dissipation of k by viscous action into heat. Again, the exact form of the equation contains unknown correlations and these are modelled in the same way as above by assuming the diffusion flux proportional to the gradient of k. The dissipation ε is usually modelled as (Rodi, 1980)

$$\varepsilon = c_D \frac{k^{3/2}}{L} \tag{3.10}$$

where c_D is another empirical constant and L is the length scale characterising the large scale motions. This relation follows from dimensional analysis when dissipation rate is determined by large-scale motion characterised by k and L.

The modelled k equation, as used in TEACH and in this thesis, is (from Rodi, (1980) using tensor notation)

$$U_i \frac{\partial k}{\partial x_i} = \frac{\partial}{\partial x_i}\left(\frac{v_t}{\sigma_k}\frac{\partial k}{\partial x_i}\right) + v_t\left(\frac{\partial U_i}{\partial x_j} + \frac{\partial U_j}{\partial x_i}\right)\frac{\partial U_i}{\partial x_j} - c_D \rho \varepsilon \tag{3.11}$$

where the terms, reading from left to right, are the convective transport, diffusive transport, production/generation by shear and viscous dissipation. This form is for

68

high values of the Reynolds number (fully turbulent) and is not applicable to the viscous sub-layer near walls.

The determination of the length scale L distinguishes the various one-equation models that use equation (3.11) and, in many, L is specified by empirical relations but these again will vary with the type of flow and, for complex flows, there is little information available on the length-scale distribution. Formulae have been proposed but these are often complex and expensive to compute.

(iii)   Two-equation models

The difficulty in finding a widely valid specification for L stimulated the development of a transport equation for determining the length-scale. The length-scale characterises the size of the large, energy-containing eddies and is subject to transport processes in a similar manner to the energy k.   Such a transport equation could be used with equation (3.11) to form a two-equation model. The length-scale equation need not have L as a dependent variable; any combination of the form

$$Z = k^m L^n$$ (3.12)

could be used because k is known from solving the k equation.

The length scale L can be characterised by the dissipation rate ε using equation 3.10. Distribution of L can be calculated by constructing a transport equation for ε and using it as part of the two-equation turbulence model.

The exact equation in this case requires drastic assumptions for complex correlations. Diffusion, generation and destruction terms require modelling and the outcome as applied in this work is the following (Rodi, 1980):

$$U_i \frac{\partial \varepsilon}{\partial x_i} = \frac{\partial}{\partial x_i}\left(\frac{v_t}{\sigma_\varepsilon} \frac{\partial \varepsilon}{\partial x_i} \right) + c_{1\varepsilon} \frac{\varepsilon}{k} G - c_{2\varepsilon} \frac{\varepsilon^2}{k}$$ (3.13)

69

where the terms, reading from left to right, represent the convection, diffusion and with the last two for generation/destruction. G is the term for the generation of turbulent kinetic energy from equation (3.11). Along with the relation, from equations (3.9) and (3.10)

$$v_t = c_\mu \frac{k^2}{\varepsilon}$$ (3.14)

equations (3.11) and (3.13) form the k-ε model.

The k-ε model has become popular and, as a consequence, is one of the most widely tested models. The constants, for high Reynolds number flows, are determined by examining a number of special cases of the modelled k-ε equations and a widely used set is due to Launder and Spalding (1974). These values of the constants are $c_\mu = 0.09$, $c_D = 1$, $c_{1\varepsilon} = 1.44$, $c_{2\varepsilon} = 1.92$, $\sigma_k = 1$ and $\sigma_\varepsilon = 1.3$ and they are the values used in all the work presented in this thesis. Rodi (1980) notes that the k-ε model with these constants has been applied successfully to a number of flow types, including re-circulating flows. Patel et al. (1985) noted that, for flows with high Reynolds number, the model behaviour depends only on the values of the five constants and the results can be sensitive to the precise values of $c_{1\varepsilon}$ and $c_{2\varepsilon}$. While both Hjertager and Magnussen (1981) and Choi et al. (1988) adjusted these constants and found improved agreement with their data as a result, such modifications are unlikely to possess the generality required of a standard model.

Chapter 4

Solution of the Mathematical Model

4.1 Discretisation of the equations

Major challenges to computational fluid dynamics come from

(i)   the physical complexity of the fluid flow, such as the phenomenon of turbulence,

(ii)   the difficulty of balancing the requirements of accuracy, stability, and economy for numerical procedures, and

(iii) the topological complexity of flow fields.

It is a formidable job to take of all these into account, which must be done when simulating real-life flow situations. Only those schemes which have a low level of complexity can find widespread use in practical engineering calculations. More and more numerical practices make it seem likely that, for the moment, no single scheme can simultaneously fulfil the requirements of accuracy, stability, boundedness and algorithmic simplicity.

The finite volume (control volume) method is one practice used to discretise the equations. This technique is readily understood by non-specialists and lends itself to direct physical interpretation. The calculation domain is divided into a set of non-overlapping volumes such that one surrounds each grid-point and the differential

equations are integrated over each control volume. Piecewise profiles expressing the variation of the dependent variable ϕ are used to evaluate the required integrals.

The derivation of discretisation equations for the finite volume method are well-known and there is a lucid presentation in Patankar (1980). The finite volume method continues to be popular despite the development of the finite volume element method (Patankar, 1980 and McCormick, 1989) which promises more effective treatment of boundary conditions, non-linearities and irregular grids.

The discretisation equation

$$A_p \phi_p = \Sigma A_{nb} \phi_{nb} + b \qquad (4.1)$$

where the general grid point variable ϕ_p is described in terms of the grid point coefficient A_p the neighbour coefficients A_{nb} and a source term b and the equation expresses the conservation principle for the finite/control volume, just as a differential equation expresses it for an infinitesimal control volume. An attractive feature of the finite volume approach is that exact satisfaction of integral conservation of quantities, such as mass and momentum, is expressed in the solution over the whole calculation domain. Even coarse-grid solutions exhibit this integral balance.

In the approach used in this thesis, and detailed by Patankar (1980), the solution is sought in the form of grid point values only. The interpolation formulae or profiles are regarded as a part of the derivation of the discretisation equations only and, once these are derived, the profile assumption can be forgotten. Different profile assumptions could be made for integrating different terms in the differential equations for velocity, pressure and turbulence.

The discretised equations were developed to obey four basic rules to ensure physical realism and overall balance.

Rule 1. Consistency at control-volume faces. Fluxes across faces common to adjacent control volumes must be represented by the same expression in the discretization equations for the two volumes.

Rule 2. All coefficients (centre-point coefficient A_p and neighbour coefficients A_{nb}) must always be positive. The resulting formulation is intended to avoid physically unrealistic solutions.

Rule 3. Negative-slope linearisation of the source term. Source terms are linearised $(S_c + S_p\phi)$ and this precaution avoids the possibility of the centre-point coefficient becoming negative via the slope of the linearised source term, i.e. the S_p term.

Rule 4. $A_P = \Sigma A_{nb}$. This reflects the fact that a given solution will also satisfy the differential equation if it is shifted by an arbitrary constant.

In this thesis, the starting point for the solution of the discretised equations was the use of a Fortran code called TEACH (Gosman and Ideriah, 1976) applies a standard tri-diagonal solver (see following section) in an iterative scheme. Non-linearities were accounted for by calculating a set of equation coefficients with the best available estimate of the dependent variables. This interim linear set is solved sufficiently to produce an updated set of variables. The coefficients are then recalculated and the process continues until some desired degree of convergence is attained. At this point the variables have effectively ceased to change due to the recalculation of coefficients.

The measure of convergence is a given reduction of the residuals for all of the equations to be solved. For equation (4.1), the residual is defined as

$$R = \Sigma A_{nb}\phi_{nb} + b - A_p\phi_p. \tag{4.2}$$

When the residual is zero then the discretisation equation is satisfied but in practice a suitable convergence criterion is that the value of the sum of the residuals for all equations to be solved is less than a selected small number.

A number of the features of the method that are of particular relevance to the work presented in this thesis are outlined in the following sections.

4.2   Solution of the linear algebraic equations

As noted above, the discretised algebraic equations are solved by a very convenient algorithm known as the Thomas or TDMA (tri-diagonal matrix algorithm). The presentation here follows that by Patankar (1980). The designation of the method refers to the required layout of the non-zero coefficients along three diagonals of the coefficient matrix. This layout arises here because when implementing a line solver all the off-line terms are added to the source term as a constant for that line and only coefficients referring to the two neighbouring cells on the line are considered.

Writing the discretised equation as

$$a_i \phi_i = b_i \phi_{i+1} + c_i \phi_{i-1} + d_i \tag{4.3}$$

for $i=1,2,3,\ldots,N$ which relates the dependent variable ϕ_i to the neighbouring points ϕ_{i+1} and ϕ_{i-1}. For the boundary condition equations set $c_i = 0$ and $b_N = 0$ so that values outside the calculation domain play no role. From this, ϕ_1 is known in terms of ϕ_2, and since ϕ_3 is known in terms of ϕ_2 and ϕ_1 then ϕ_3 can be expressed in terms of ϕ_2 alone. This process can be continued until ϕ_N is expressed in terms of ϕ_{N+1}. Because ϕ_{N+1} has no meaning, the numerical value of ϕ_N is determined at this stage.

Back-substitution can then begin which allows all the other values of ϕ to be calculated. If, in the forward substitution process, a relation is sought such that

$$\phi_i = P_i \phi_{i+1} + Q_i \tag{4.4}$$

and the previous equation was

$$\phi_{i-1} = P_{i-1}\phi_i + Q_{i-1} \tag{4.5}$$

then substitution of equation (4.5) into equation (4.3) leads to the recurrence relations

$$P_i = \frac{b_i}{a_i - c_i P_{i-1}} \tag{4.6}$$

$$Q_i = \frac{d_i + c_i Q_{i-1}}{a_i - c_i P_{i-1}}. \tag{4.7}$$

To start the recurrence process, it is noted that

$$P_1 = \frac{b_1}{a_1} \tag{4.8}$$

and

$$Q_1 = \frac{d_1}{a_1}. \tag{4.9}$$

For ϕ_N, note that $P_N = 0$ because $b_N = 0$ and therefore

$$\phi_N = Q_N \tag{4.10}$$

The sequence of operation for the algorithm is as follows:

(i) Calculate P_1 and Q_1 from equations (4.8) and (4.9).

(ii) Use the recurrence relations (4.6) and (4.7) to obtain all P_i and Q_i.

(iii) Set $\phi_N = Q_N$.

(iv) Use equation (4.4) for $i = N-1, N-2, \ldots, 3, 2, 1$ to obtain $\phi_{N-1}, \phi_{N-2}, \ldots \phi_3, \phi_2, \phi_1$.

4.3 The staggered grid

Harlow and Welch (1965) described a new technique for the numerical investigation of the time-dependent flow of an incompressible fluid. The full Navier-Stokes equations were written in finite-difference form and the solution was accomplished by finite-time-step advancement. The primary variables were the pressure and the velocity components which was a departure from the usual stream function and vorticity methods that were predominant at that time. A staggered grid equivalent was used where velocities were defined at cell boundaries and pressures at cell centres. The TEACH code uses a staggered grid. A co-located grid is one where all variables to be calculated are coincident on the grid points. A staggered grid is one where some variables (velocities in this case) are displaced from the positions of others (pressure in this case).

A staggered grid is used because a particular difficulty arises in the grid representation of pressure gradient terms in the velocity equations. The pressure drop is across the control volume is the quantity of interest but, with a piecewise-linear profile for pressure, the discretised momentum equations contain the pressure difference between alternate grid points and not adjacent ones.

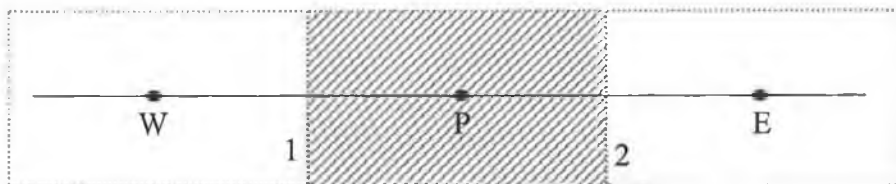Consider the representation of the term $-dp/dx$ over the control volume shown in figure 4.1.

Figure 4.1. Control volume for the pressure gradient calculation.

This is calculated from the pressure difference between the faces 1 and 2 in figure 4.1. To derive the face values a linear profile may be assumed between the grid points W and P and P and E. If the faces lie mid-way between the grid points then

$$p_1 - p_2 = \frac{p_w - p_P}{2} - \frac{p_P + p_E}{2} = \frac{p_W - p_E}{2} \qquad (4.11)$$

Thus the dependence on alternate grid points is demonstrated. Because of this a zig-zag pressure field, however unrealistic, such as that shown in figure 4.2 could be calculated as a uniform pressure field because the pressure difference between two alternate points is zero. Should such pressure fields arise during an iterative solution, they would be preserved to convergence because the momentum equations would not calculate their presence. Note that there is also a loss of accuracy because the pressure gradients are derived from a grid that is coarser than the one prescribed for the problem.

100	300	100	300	100	300
5	27	5	27	5	27
100	300	100	300	100	300
5	27	5	27	5	27
100	300	100	300	100	300
5	27	5	27	5	27

Figure 4.2. A zigzag pressure field that could appear uniform to the momentum equations.

The under-lying problem here is that the pressures used to calculate the gradient term are not on grid points and must be interpolated. One solution to this problem is the use of a staggered grid.

A staggered grid means that each dependent variable is regarded as having its own grid attached to it. The usual practice is that the grid layouts are identical but some grids can be displaced relative to others in order to gain advantage in the type of calculation above. For example, figure 4.3 shows two grids, one displaced from the other. The solid gridlines show the first set of grid points which could be used for scalar quantities. The dashed lines represent the grid for the u velocity component. This is displaced so that the velocity component points lie midway between the scalar grid points. The grid points for the scalar quantities are represented by solid symbols and the displaced points are represented by circles.
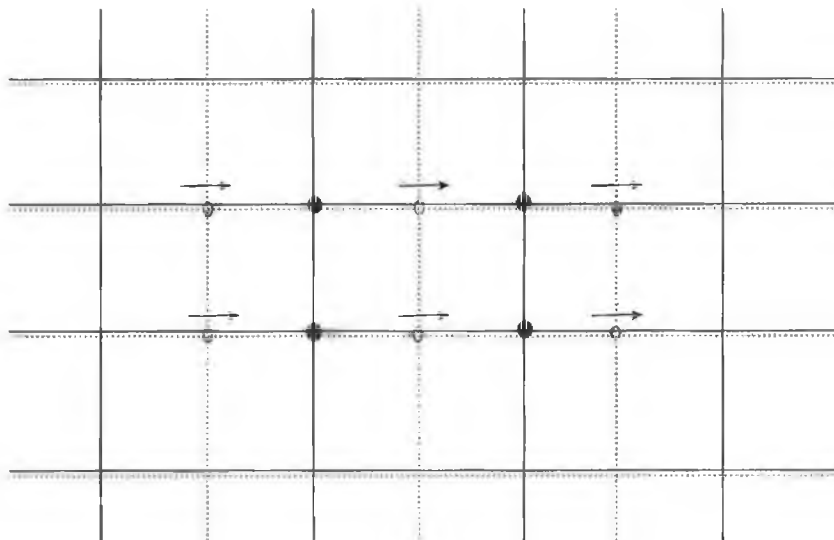

Figure 4.3. Displacement of the u velocities relative to the scalar grid.

Displacing the grids for the velocity components and the pressures allows the pressure gradient term to use adjacent grid pressures because their difference is now the natural driving force for the momentum equations. The other velocity components are similarly

displaced to achieve the same effect. The v component would be displaced vertically in figure 4.3 and the z component would move off the plane of the diagram in a three-dimensional case.

A similar problem arises for the discretised continuity equation. A wavy velocity field can satisfy this because it demands equality of velocities at alternate grid points and not at adjacent ones. In the staggered grid only reasonable velocity fields would satisfy the continuity equation because adjacent points are used.

Also, mass flow rates across the control volume faces can be calculated without any interpolation for the relevant velocity component.

Peric et al. (1988) compared finite-volume numerical methods with staggered and co-located grids in two dimensions. The results demonstrated that convergence rates, dependence on under-relaxation parameters, computational effort and accuracy were almost identical for both solution methods. They considered that co-located grids offered some advantages for multi-grid techniques and for non-orthogonal grids.

4.4   Location of control volume faces

There are two options for the placement of control volume faces relative to grid points, namely,

(i)   Faces located midway between the grid points.

(ii)   Grid points at the centres of the control volumes.

Note that these options are identical for a uniform grid. They are illustrated for a non-uniform grid in figures 4.4 and 4.5. In both of these, and in figure 4.6, the solid line intersections define the grid points and the dotted lines indicate the positions of interfaces between the cells around each grid point. Shaded areas denote typical cells.

Note that method 1 usually arises from selecting the grid points first and then constructing the control volumes.

Method 2 results from first defining the control volumes. Figure 4.4 shows the volumes that would be constructed by method 1 to fit the grid shown. Figure 4.5 is for the same control volumes but it is assumed that these are already in place and the grid points are then laid out at their centres.
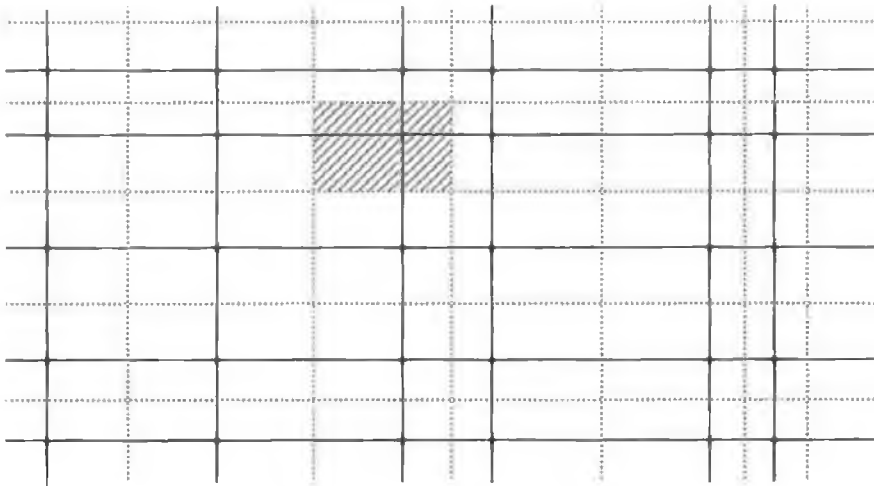

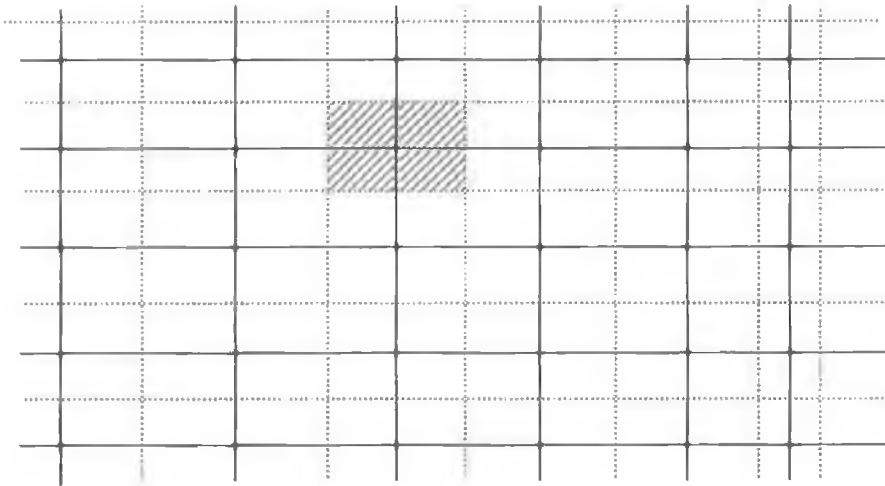**Figure 4.4. Control volume faces located by method 1.**

Figure 4.5. Grid points located by method 2, i.e. at control volume centres.

The midway faces provide greater accuracy in calculating fluxes across a face because the slope of a piecewise linear profile is the same as the slope of any parabolic profile evaluated midway between the grid points, i.e. the result of a linear profile corresponds to a less crude parabolic one.

Unfortunately, the point between two grid points, where the flux is calculated, can be well off the centre of the control volume face and the assumption that this is representative can lead to inaccuracy. Further, if the grid point is not set at the centre of a volume its usefulness as a representative value for that the volume will vary with its position in the control volume and, for a non-uniform grid, will vary through the calculation domain.

Method 2 should not suffer to the same extent from these sources of inaccuracy because the grid point is located in the centre of the control volume but it does not have the profile advantage. It does, however, offer more convenience in setting out a calculation domain because the control volume is the basic unit of discretisation. Boundary discontinuities, e.g. material properties, can be conveniently handled.

81

Control volumes can be designed to avoid discontinuities occurring within faces. Setting out the grid points in method 1 and then constructing control volumes makes it more difficult to set them in desirable locations.

In figure 4.6, where the shaded areas represent a solid boundary outside a region of flow, it can be seen that, by defining individual control volumes as elements of either the boundary or flow domain, a typical volume face is located on the boundary and the available boundary condition data can be used directly.
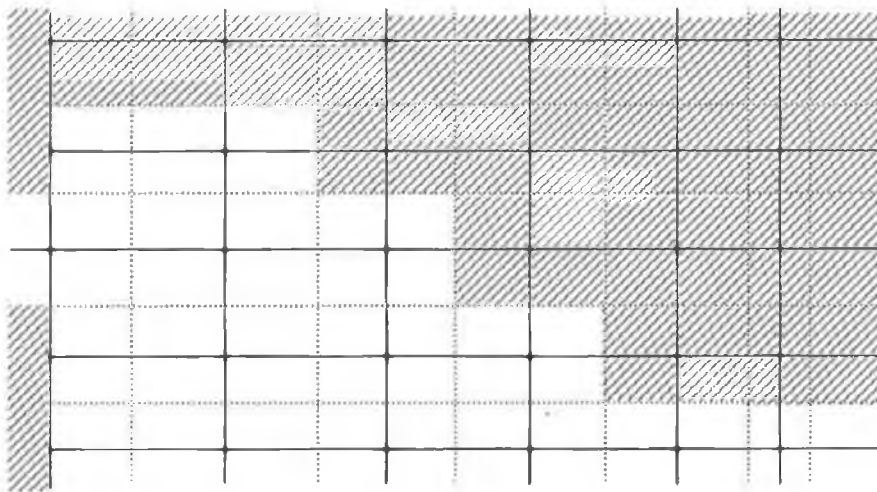

Figure 4.6. A method 2 domain showing convenience in setting boundary conditions.

Method 2 also avoids complications in the construction of coarser grids based on the one initially constructed. In this work these coarse grids are constructed from groups of finer grid control volumes and the faces are coincident. The coarse grids are to be used in the evaluation of a multigrid acceleration technique for the new solution methods developed.

For the rest of this thesis attention is devoted to grids in the Cartesian co-ordinate system in order to avoid complications and potential for coding errors in the specification of grid geometry and control volume dimensions. However, the methods used are not in principle limited to the Cartesian system and they can be used with a grid in any orthogonal co-ordinate system, e.g. polar co-ordinates. The features introduced by such a system are mainly geometric and provided the lengths, face areas and volumes are correctly calculated no new principles should be required. Orthogonality allows profiles defined by just two grid points to be used.

The grids generated for this work were simple in structure and no automatic generation method was required. The TEACH code used as the basis for this work did not have any grid generation capability and did not allow the use of body-fitted co-ordinates.

4.5   Source term linearisation

Often the source term is a function of the dependent variable and it is often difficult, although desirable, to account for this in the construction of the discretisation equation. Formally it is only possible to account for a linear dependence and an average value can be discussed and expressed as

$$S = S_C + S_P \phi_P \qquad (4.12)$$

where S_C is the constant part of S and S_P is the coefficient of the dependent variable. The incorporation of a linear dependence is preferable to treating the source term as a constant. The appearance of ϕ_P shows the presumption of the usefulness of an average value that applies over the entire control volume, i.e. a step-wise profile. (It should be noted that the use of a stepwise profile in the source term does not limit the profile assumptions for the other terms.)

When S is a nonlinear function of the dependent variable it must be linearised, i.e. specify the values of S_C and S_P that may themselves depend on the variable value.

During the iteration process, the source terms are updated as the variable changes. S should be a good representation of the relationship and the basic rule about non-positive values for S_P must be obeyed.

The linearised source term can be used as a means of setting the desired values for variables at a grid point. To understand this, note the role of the source terms in the discretized two-dimensional equation

$$a_P\phi_P = a_E\phi_E + a_W\phi_W + a_N\phi_N + a_S\phi_S + b \qquad (4.13)$$

where

$$b = Sc\Delta V \qquad (4.14)$$

and

$$a_P = a_E + a_W + a_N + a_S - S_p\Delta V . \qquad (4.15)$$

ΔV is the volume of the finite control volume, ϕ is the dependent variable and a represents the convection/diffusion coefficients for the grid point and interfaces with the neighbouring values.

By setting S_C and S_P to very large values then other terms become negligible. Thus, using an arbitrarily chosen, very large number,

$$S_C = 10^{30}\phi_{P,desired} \qquad (4.16)$$

and

$$S_P = -10^{30}. \qquad (4.17)$$

The consequence is that

$$S_C + S_P\phi_P \approx 0 \qquad (4.18)$$

84

and

$$\phi_P = -\frac{S_c}{S_P} = \phi_{P,desired} \qquad (4.19)$$

This property was used in the approximation of curved walls within the rectangular grid calculation domain and in the insertion of blocks to represent shelving. Air velocities for the volumes outside the flow can be conveniently set to zero

4.6   Under-relaxation of the equations

In the iterative solution procedure of the algebraic equations, under-relaxation is used as a method of slowing down changes in the values of the dependent variables from one coefficient update to another. It is a common technique for avoiding divergence in the iterative solution of strongly non-linear or coupled non-linear equations.

Equation (1.1) can be written as

$$\phi_p = \frac{\Sigma A_{nb}\phi_{nb} + b}{A_p} \qquad (4.20)$$

If we add ϕ_p^* to the right-hand side and subtract it, we have

$$\phi_p = \phi_p^* + \left(\frac{\Sigma A_{nb}\phi_{nb} + b}{A_p} - \phi_p^* \right). \qquad (4.21)$$

A relaxation factor α can be introduced to control the term in brackets, which represents the change in ϕ_p produced by the current update cycle, so that

$$\phi_p = \phi_p^* + \alpha \left(\frac{\Sigma A_{nb}\phi_{nb} + b}{A_p} - \phi_p^* \right) \tag{4.22}$$

or

$$\frac{A_p}{\alpha}\phi_p = \Sigma A_{nb}\phi_{nb} + b + (1-\alpha)\frac{A_p}{\alpha}\phi_p^*. \tag{4.23}$$

Values of α between 0 and 1 produce under-relaxation and the values of ϕ_p stay closer to ϕ_p^*. When the iterations converge then ϕ_p becomes equal to ϕ_p^*, which is a necessary condition because the result of a solution scheme must satisfy the original equation and be free of any effect of any arbitrary constant which is introduced.

4.7 Differencing scheme

The choice of differencing scheme for the convection term is a very active field of research and has lead to a number of alternative proposals. In integrating over the control volume (shaded area) shown in figure 4.1, values for the dependent variable are required at the boundaries of the volume. It is natural to set the value at the interface between the points labelled P and E as a linear interpolation between the values at the two points, i.e.

$$\phi_e = \frac{1}{2}(\phi_E + \phi_P) \tag{4.24}$$

This gives rise to the central difference scheme but this scheme fails if the cell Peclet number, P, given by

$$P = \frac{\rho u L}{\Gamma} \tag{4.25}$$

(where ρ is the density, Γ the diffusion coefficient, u is a velocity component and L is a length scale (in this case the grid cell dimension)), is greater than 2 (for a uniform grid) (Patankar, 1980). It was found that numerical instability could be overcome by using one-sided, upwind, first-order finite-difference approximations to represent the derivatives in the convection terms (Courant et al., 1952) and solutions for high Reynolds number flows were published by, for example, Bozeman and Dalton (1973). (The Reynolds number is defined as the ratio of representative inertial and viscous forces, $Re = \rho u L / \mu$ where ρ is the density, u the characteristic air speed, L the characteristic length and μ is the viscosity.) The upwind scheme calculates the convection terms in the discretised equation from the assumption that the value of the dependent variable at an interface is equal to its value at the grid point on the upwind side of the face.

It was recognised that first-order upwind difference approximations can generate significant truncation errors, the most serious of which produces a diffusive effect which augments the effects of viscosity (de Vahl Davis and Mallinson, 1976 and Raithby, 1976). As a consequence, considerable controversy ensued over the significance of this false diffusion, especially in the case of high Reynolds number flow when the physical diffusion is weak. Solutions obtained at that time demonstrated that using upwind differencing in solutions predicted the gross features of the flows adequately but there was insufficient experimental evidence to permit comparisons sufficiently accurate to allow quantitative estimates of the effects of truncation errors on the solutions. De Vahl Davis and Mallinson (1976) stated that if accuracy was required then high-order approximations were necessary. It was also stated that false diffusion from first-order approximations could introduce significant errors when the direction of flow was inclined to the axes of the finite difference mesh. Bozeman and Dalton (1973) tested several methods for approximating convection terms in the governing equations. They showed that some of the upwind difference approximations used previously produced unsatisfactory solutions at high Reynolds numbers. However, they were forced to use an upwind difference approximation to obtain high Reynolds number solutions.

A hybrid scheme was developed by Spalding (1972) which addressed difficulties with the earlier upwind differencing scheme. The name hybrid indicates that the scheme can essentially be seen as a combination of central differencing and the upwind scheme with upwind differencing at high Peclet numbers and central differencing at low Péclet numbers (less than 2 on a uniform grid). The combination produces unconditionally positive coefficients in the discretised equations.

By invoking an analogy between false diffusion and the conduction of heat in an isotropic medium, de Vahl Davis and Mallinson (1976) derived theoretical expressions for estimating the false diffusion coefficients in steady 2-dimensional flows. The coefficients were found to be proportional to the magnitudes of the fluid velocity vector, the grid/mesh size and dependent on the angle between the fluid velocity vector and the grid lines. Expressions were derived for three dimensions because the need for such expressions is even greater in 3-dimensional flow calculations since storage limitations necessitate the use of coarser grid distributions, resulting in larger truncation errors. Raithby (1976) presented a detailed analysis of the magnitudes and practical effects of the false diffusion under various one-dimensional flow situations. Although formally only first-order accurate, Raithby's skew differencing scheme yielded significant reductions in numerical diffusion by taking the direction of the fluid velocity vector into account (Raithby, 1976).

More and more complex fluid flows are being computed numerically and it has often been impossible to separate numerical errors from deficiencies of the turbulence (or other) models. A number of upstream weighted difference schemes of higher order accuracy and satisfactory numerical stability have been proposed but these were slow to find wide applicability because of their complexity and lack of boundedness under certain flow conditions. For example, non-physical overshoots and undershoots are sometimes produced in flow regions with very steep property gradients (e.g. Leschziner, 1980 and Han et al., 1981). Gresho and Lee (1981) suggested that the presence of wiggles in a solution should perhaps be taken as an indication of inadequacies in the computational grid employed.

A desire to use high-order approximations and small mesh intervals to ensure accuracy has to be balanced against limitations imposed by the complexity of the problem being

solved, the available computing equipment and the stability requirements of the solution algorithm. Determining the effects of a given compromise between these aims is not always easy and determining confidence limits for a particular solution, as opposed to convergence of a particular scheme can be very difficult. Early attempts to apply second order approximations to the terms in the governing equations representing convection (e.g. Burggraf, 1966) failed to produce stable solutions for high values of the relevant Reynolds number, values for which experiments (Pan and Acrivos, 1967) showed steady laminar flow.

Methods were introduced that allowed greater flexibility in setting variable profiles between grid points by using more than two neighbouring grid points in interpolating variable values. An upwind, second-order differencing scheme was proposed by Warming and Beam (1976) but one such method that has attracted a good deal of attention in recent years is the QUICK (Quadratic Upstream Interpolation for Convective Kinematics) method and its variants, e.g. QUICKER. The method was originally proposed by Leonard (1979) and it allows greater flexibility in setting variable profiles between grid points by using more than two neighbouring grid points in interpolating variable values. Leonard proposed that it was necessary to develop methods that avoided the stability problems of central differencing while remaining free of the inaccuracies of numerical diffusion associated with upstream differencing (Patankar, 1980). While it was possible to eliminate the latter problem by sufficiently refining the grid, this could impose a requirement of considerable extra computational effort (Raithby and Torrance, 1974). Leonard states that numerical experience of skew differencing (Raithby, 1976) indicated that the remaining truncation error was of the same order as the skewness error which had been eliminated.

Leschziner and Rodi (1981) examined the performance of three discretisation schemes for convection and three turbulence-model variations when used to simulate the re-circulating flow in an annular and a plane twin-parallel jet in still air. The schemes considered were the hybrid scheme, the skew-upwind differencing scheme and the quadratic, upstream weighted differencing scheme of Leonard (1979). It was shown that the hybrid scheme resulted in severe solution errors and in a generally anomalous response to changes in the modelled viscosity field. In contrast, the others yielded, in all

cases, similar results and responded in an expected manner to the modifications of the turbulence models.

Stubley et al. (1980) examined the connections between approximations made in the discretisation process and final solution errors, particularly when convection was a dominant factor. The discrete equations were reformulated and two different schemes were tested. While the schemes were not proposed as practical for numerical prediction (e.g. they were not conservative) they were seen as providing a basis for such and thought to hold considerable promise for future development. Also, Stubley et al. (1982), considered the error in a finite-difference solution as arising from two sources: the profile error component arose because the assumed profile used in deriving the scheme did not match the exact solution and the operator error component which is due to the failure of the finite-difference operator to accurately simulate the convection-diffusion process. They found profile error was relatively easy to quantify but to evaluate the performance of differencing schemes in convection dominated problems required a test for operator error. They proposed such a test and examined how errors due to sharp changes in velocity and difference coefficients over the solution domain could be reduced.

Numerical calculations were presented by Hjertager and Magnussen (1981) for two jet-induced three-dimensional flows in rectangular enclosures. The computational method used the three velocity components and pressure as the primary variables. It was found that the numerical smearing introduced by upwind schemes did not upset the solutions. Some explanations for this phenomenon are offered by McGuirk and Rodi (1978) who have shown that the question of first-, second- or higher-order accurate finite-difference schemes are not always the proper indication of accuracy. They point out and show that it is more appropriate to look at the relative magnitude of the different terms occurring in the differential equations at different locations in the solution domain to assess accuracy. Cheng and Shubin (1978) also concluded, after a study of the one-dimensional, steady state Burger's equation, that the formal order of accuracy of a difference algorithm may not reflect the magnitude of the computational errors. A similar view on numerical diffusions is expressed in the book by Patankar (1980).

However, development of the applications of QUICK continued and Leszchiner (1980) published an evaluation of hybrid differencing, skew-upwind differencing and the QUICK scheme. Computations were presented for a number of test configurations, both linear and non-linear. It was shown that, in all cases, the skew-upwind and QUICK differencing schemes yielded similar results that were superior to those obtained with the hybrid scheme, although the former schemes were found to suffer to a limited extent from boundedness problems. The schemes were applied to two confined, laminar, re-circulating flows for which experimental data were available. The predictions of all three schemes were similar for the flows examined (backward-facing step and sudden pipe expansion) and the author argued that, for the low Reynolds number cases that were examined, numerical diffusion gave rise to small errors.

Some answers may lie in the use of higher-order techniques with filtering techniques to remove spurious oscillations. In some such methods (e.g. Chapman, 1981), artificial diffusivities are introduced in regions where the solution with a purely higher-order scheme becomes unbounded. The major problem appears to be that of determining the bounds for each of the dependent variables in such flows with complex interactions.

Pollard and Siu (1982) evaluated upwind, hybrid and quadratic differencing schemes for a number of laminar, two-dimensional flows. They introduced two new forms of the quadratic formulation to overcome stability problems with the original one and found them to be stable although one of them seemed to be sensitive to the size of the grid. The accuracy of the solutions obtained with the new schemes was equal to or better than those from the upwind and the hybrid schemes. They noted the difficulties encountered by other authors (e.g., Han et al., 1981) that required the introduction of pseudo-source term in the QUICK scheme to achieve stability and provided a formulation that ensured that the equation coefficients were always positive.

Demuren (1985) presented the derivation of the three principal components of the false or numerical diffusivities which result from the use of first-order upwind difference schemes in the calculation of steady three-dimensional flows. The effects of the false diffusion were illustrated by comparison of the first-order calculations with those with a higher-order scheme and experimental data, for the flow of a row of turbulent jets

issuing at right angles to a nearly-uniform cross-stream. The adverse effect of the false diffusivities in suppressing the physics of the flow was shown.

Patel and Markatos (1986) made a comparative study of eight discretisation schemes, including the previously developed variants of QUICK (referred to as Quadratic Upstream Differencing Scheme (QUDS)). The latter were found to give the best performance in terms of accuracy for the two test cases (a sudden expansion pipe flow and a lid-driven cavity flow) although the QUICK schemes were found to require 30% to 60% more computational effort in the tests than others. While generally concluding that QUICK schemes offer the best compromise of accuracy and efficiency it was noted that divergence of the solution procedures was encountered when the authors applied them to some more complex engineering situations and they expressed the concern that their range of application might be restricted. In a later study Patel et al. (1987) applied a number of discretisation schemes to supersonic jet flows and in this case failed to achieve a convergent solution for the QUICK schemes that were employed.

An alternative method for the reduction of numerical diffusion was proposed by Huh et al. (1986). A corrective scheme was developed that compensated for the effect of cross-flow diffusion by reducing the effective, anisotropic, cross-flow diffusion coefficient in the diffusion portion of the simulation. However, for certain cases characterised by high Peclet numbers and substantial curvature of the velocity field over a single mesh spacing, unphysical oscillations were found to arise which could degrade the solution accuracy very rapidly.

Freitas et al. (1985) simulated flow in a cavity using a finite-difference code with the incorporation of a modified QUICK scheme and claimed the first successful calculation of a three-dimensional laminar lid-driven cavity flow. The code used was essentially the TEACH programme of Gosman and Ideriah (1976) and with the inclusion of a QUICK scheme they were able to numerically reproduce the experimentally-observed vortices and other general flow structures. Four different convection-diffusion schemes were used by Huang et al. (1985) who found that QUICK gave very accurate results.

An interesting example of a body fitted coordinates application was the work of Shyy et al. (1985). They presented a finite-difference algorithm for re-circulating flow that incorporated a staggered grid. Implementation in curvilinear co-ordinates led to interesting results in the application of three differencing schemes. The authors found results that were quite different to those in cartesian coordinates in accuracy and efficiency aspects. They stated that the choice of grid distribution system was at least as important as the choice of difference scheme. In particular, when nodes were more concentrated near boundaries it was found that the QUICK became unstable and convergence became difficult even at very low under-relaxation factors.

Leonard (1988), in developing a method for highly convective simulation of step profiles, notes that QUICK produced unphysical overshoots and oscillations. He also noted that convergence difficulties with the TEACH code were sometimes due to using a single sweep direction in the solution procedure. Chen and Falconer (1992) applied QUICK schemes to advection-diffusion modelling for coastal and estuarine flows and found that QUICK-based schemes were attractive compared with other comparable second-order schemes.

Patankar (1988) presented a summary of the state-of-the-art in methods rather than applications. At the time he considered that attempts to improve differencing schemes had come full circle and left the hybrid scheme among the most stable but subject to numerical/false diffusion while the QUICK scheme, among others, was subject to 'wiggles' and convergence difficulties. Gaskell et al. (1988), for example, note that the characteristic lack of boundedness of such schemes have ensured the continued use of the hybrid scheme for problems that contain steep gradients in one or more of the dependent variables. Patankar also noted that for more complex flows it becomes difficult to distinguish between these unphysical results and the detailed behaviour of the flow. For example, Demirdzic et al. (1987) provided a general finite-volume approach to the calculation of turbulent flow in geometrically complex domains and the general result was that although numerically accurate (i.e. grid independent) results could be achieved there were appreciable discrepancies between the predictions and the experimental data which the authors attributed to turbulence modelling deficiencies.

A paper by Zhu and Rodi (1991) addressed convection modelling and flow field zoning procedures destined for finite-volume methods for incompressible steady-state flows with irregular boundaries. A composite oscillation damping algorithm, which is capable of yielding bounded and low diffusive solutions, was used to approximate the convection terms of transport equations. To deal with complex geometries, a zonal procedure was introduced into an advanced finite-volume formulation that used general non-orthogonal and non-staggered grids. The solution domain was divided into simple sub-regions each covered by a separate mesh and this enabled the FV method to handle those domains which were difficult to cover with a single grid. Information transfer between zones was accomplished by overlapping of the grids at the interfaces.

Zhu and Rodi used a procedure combining second-order upwind central differencing and first-order upwind schemes with the switch from one scheme to the other being automatically controlled by a convection boundedness criterion. Test results showed that it strictly preserved the boundedness of solutions, while maintaining low numerical smearing of steep gradients. Comparison with the QUICK scheme showed that it retained the non-diffusive nature of QUICK but without the unboundedness in the vicinity of sharp corners. It could be seen that the results were nearly as good as those obtained by QUICK with respect to resolving the steep gradients. Comparisons between the new scheme and the hybrid central/upwind scheme showed non-negligible discrepancies between the hybrid scheme and the experimental data for coarse grids (22×10) where errors were clear and good agreement for the new scheme. Moving to a finer grid (32×20) removed the discrepancies. It was stated that all results showed the new scheme was more accurate on coarse grids and that differences decrease with grid refinement.

Ikahogi et al. (1992) noted that in the course of calculations for a 3D duct flow that numerical instability had occurred using the QUICK scheme.

A new formulation developed by Hayase et al. (1992) used the four rules for physically realistic solutions having overall balance as stated by Patankar (1980) and given above in Section 4.1. Testing showed that this formulation was more stable and converged faster than any of the previous formulations. The testing included the relative evaluation of the boundary conditions approximated by second- and third-

order finite-difference schemes as well as calculations performed at higher Reynolds number (greater than 10,000) than previously reported.

Barton (1995) presented numerical solutions for laminar flow over a backward-facing step with five difference schemes, including QUICK and a second-order upwind scheme (Warming and Beam, 1976) as well a novel variant of the latter. QUICK and the second-order upwind schemes gave similar results while the hybrid scheme suffered from excessive numerical diffusion except for very fine grids and central-differencing converged only for the finest grids. Evaluation was made by comparing results with experimental data.

Li and Baldacchino (1995) supplied a generalised formulation to implement the QUICK scheme, the second-order upwind scheme and the second-order hybrid scheme (Li and Rudman, 1995) on non-uniform grids. The implementation method was straightforward in that it was formulated as a first-order upwind scheme with extra terms for the other schemes included as a 'deferred correction' in the source terms. This was following the method of Khosla and Rubin (1974), who described an unconditionally stable, second-order accurate, implicit, finite-difference method where the coefficient matrix was tri-diagonal and always diagonally dominant. The accuracy of the results was evaluated by comparison with known exact solutions for a number of problems and they found that the order of accuracy of the examined schemes was preserved if the distribution of the non-uniform grid points was properly chosen. A reduction in accuracy was noted for a two-dimensional flow calculation but the treatment of cells at boundaries was not consistent in that the problem of references to values lying outside the calculation domain was solved by reducing to the first-order upwind scheme for those cells lying next to the boundary. A more consistent approach would be desirable.

4.8 The SIMPLE algorithm

The solution procedure which was taken as the starting point for the calculation of flows in this thesis, and which has been used as a standard against which to judge the performance of the new algorithms, is the well-known Semi-Implicit Method for Pressure Linked Equations (SIMPLE). The procedure is described in Patankar and Spalding (1972), Caretto, Gosman, Patankar and Spalding (1972) and Pantankar (1975 and 1980). It involves the sequential solution of the set of equations, i.e. each equation in turn, for velocities, a pressure correction equation and those for other variables such as temperature or turbulence parameters. In the notes below u, v and w are the velocity components and p is the pressure.

In order to solve the momentum equations, a pressure field has to be specified (estimated) and for the solution to proceed a means of improving the initial specification has to exist. Patankar (1980) wrote the desired pressure and velocity fields in terms of the estimate and a correction, i.e.

$$p = p^* + p', \tag{4.26}$$

$$u = u^* + u' \tag{4.27}$$

and similarly for the v and w components of the fluid velocity, and then derived velocity correction formulae from the momentum equations. These have the form, for example,

$$u_e = u^*_e + \frac{A_e}{a_e}(p'_P - p'_E) \tag{4.28}$$

where u_e is the velocity at the east face of a grid cell, a_e is the coefficient of this velocity, A_e is the area of the interface and p_P and p_e at the pressures at the cell centre qrid and at the centre of the eastern neighbour cell. Substituting these corrections into the discretised continuity equation yields a discretisation equation for values of p' which is

$$a_P p'_P = a_E p'_E + a_W p'_W + a_N p'_N + a_S p'_S + a_F p'_F + a_B p'_B + b \qquad (4.29)$$

where

$$a_E = \rho_e \frac{A_e}{a_e} \Delta y \Delta z, \qquad (4.30)$$

$$a_W = \rho_w \frac{A_w}{a_w} \Delta y \Delta z \qquad (4.31)$$

and the other coefficients have the same form except

$$a_P = a_E + a_W + a_N + a_S + a_F + a_B \qquad (4.32)$$

and

$$b = [(\rho u^*)_w - (\rho u^*)_e]\Delta y \Delta z + [(\rho v^*)_s - (\rho v^*)_n]\Delta z \Delta x$$
$$+ [(\rho w^*)_b - (\rho w^*)_f]\Delta x \Delta y \qquad (4.33)$$

This b term represents a mass source which must be removed by the pressure corrections through their associated fluid velocity corrections.

The key feature of the SIMPLE procedure, which has motivated a number of alternative formulations such as SIMPLER (Patankar,1980) and SIMPLEC (Van Doormaal and Raithby, 1984), is the application of corrections to the pressures and fluid velocities. There are approximations made in the construction of the pressure correction equation which, while not influencing the converged solution, do influence the course of the solution procedure. The term semi-implicit is used in the algorithm title to acknowledge that a term is omitted which has an indirect or implicit influence on the variable estimates as the solution proceeds.

With the pressure correction equation thus developed the sequence of operations for the SIMPLE algorithm can now be specified. The sequence is

(i) Guess a pressure field p^*.

(ii) Solve the momentum equations to obtain u^*, v^* and w^*.

(iii) Solve the pressure correction equation.

(iv) Calculate p by adding the pressure correction to p^*.

(v) Calculate u, v and w from u^*, v^* and w^* using the fluid velocity correction equations.

(vi) Calculate any other variables that influence the flow field.

(vii) If the solution has not converged, then treat the corrected pressures as a new guess to restart the cycle at step (i).

The proposal underlying the work of Raithby and Schneider (1979) was that coupling between the momentum and mass conservation equations for incompressible flows was often the major cause of the slow convergence of iterative solution techniques. Several methods of handling this coupling were examined and the results of their application to a test problem were compared. The application was made in a manner that completely isolated the effect of the coupling and led to a clearer understanding of how the methods performed. This was accomplished by solving for an 'exact' solution. The coefficients and exact values of the fluid velocities and pressures were stored and these coefficients used to solve from an initial guess of zero values. They noted that methods which eliminate the pressure correction steps needed further study.

SIMPLEC is the more successful of the alternative formulations. Savings in computation times relative to SIMPLE vary from 50% to approximately 70% in the report by Van Doormaal and Raithby (1984). They introduced a "consistent" approximation in the construction of the pressure correction equation (as opposed to

the almost arbitrary removal of a term in the derivations for the SIMPLE method). The computational improvements appear to be largely due to the new approximation becoming exact when diffusion and advection of momentum are negligible compared to the pressure gradient and source terms.

Armfield (1991) developed a non-staggered mesh scheme and compared it to SIMPLE on a staggered mesh. The result was equivalent in accuracy for the two schemes and improved the efficiency for the new scheme on non-uniform grids. The pressure correction technique for solving the Navier-Stokes and Euler equations was studied by Connell and Stow (1986). A general pressure correction analysis was presented using the full form of the linearised momentum equations. It was shown how the SIMPLE procedure resulted from adopting the most approximate form of the momentum equations. Two extended pressure correction techniques were derived using higher-order approximations and shown to be up to 50% faster than SIMPLE on the test example.

Pressure correction methods continue to be developed. Davidson (1996) presented an application of SIMPLEC on a cell-centred, collocated finite volume method for general unstructured meshes. In this work cells could have an arbitrary number of grid points (cell vertices) and the applications considered varied between three and six cell faces. The discretized equations were solved using either a symmetric Gauss-Seidel solver or a conjugate gradient solver with a preconditioner. The method was applied to two-dimensional test cases (incompressible, laminar flows) and it was stated that extension to three dimensions and compressible flows was straightforward.

An alternative method (PISO - Pressure Implicit with Splitting of Operators) for handling the coupling of the implicitly-discretised time-dependent fluid flow equations was described by Issa (1985). The main feature of the technique was the splitting of the solution process into a series of steps whereby operations on pressure were decoupled from those on velocity at each step, with the split sets of equations being amenable to solution by standard techniques. Patankar (1988) interprets this method as a version of SIMPLER.

Properties of the PISO (Pressure Implicit with Splitting of Operators) and SIMPLE algorithms for the solution of the momentum equations were examined by comparison of

the computational effort required for reaching the same convergence criterion in two test problems by Wanik and Schnell (1989). Swirling and non-swirling flows were considered in axisymmetric geometries and in the computations both iterative and time-marching versions of the methods were considered. Optimal values of the numerical parameters, on which the effectiveness of both algorithms depended for various different grids, were found and the results for the swirling flow problems were verified on the basis of the available experimental data. These workers found that PISO required approximately 50% of the effort required for SIMPLE but Wanik and Schnell found that the standard k-ε model was inadequate for swirling flows.

SIMPLE was compared to other two other methods of handling pressure-velocity coupling by McGuirk and Palma (1993) for incompressible turbulent flow problems. In this study SIMPLE did not provide the fastest rate of convergence but it was competitive with the best technique (PISO) and had advantages in terms of storage requirements. The study did not lend any support to a move away from SIMPLE as a popular technique.

The use of SIMPLE with the QUICK scheme is reported in the work of Shyy et al. (1985) and Barton (1995) and these were discussed in Section 4.7.

4. 9 QUICK Formulation

Leonard (1979) achieved a conservative formulation with stable convective sensitivity by using a three point upstream-weighted quadratic interpolation scheme for each wall of a control volume individually. Using a constant grid spacing for convenience the formula used for interpolation of the value of, say, ϕ_r, when u_r is positive to the right as illustrated in figure 4.7, was

$$\phi_r = \frac{1}{2}(\phi_C + \phi_R) - \frac{1}{8}(\phi_L + \phi_R - 2\phi_C) \qquad (4.34)$$

which may be interpreted as a linear interpolation corrected by a term proportional to the upstream-weighted curvature. Leonard noted that because of the consistent quadratic interpolation used in modelling both convective and diffusive terms, the overall truncation error in a solution obtained by the QUICK algorithm is third-order in spatial grid size and that highly accurate solutions could be obtained using practical grid spacing.

In all the formulations given below, the flow rate F and the conductance D are given by, for the east faces of the finite volume, for example,

$$F_e = (\rho u)_e A_e \tag{4.35}$$

$$D_e = \frac{\Gamma_e A_e}{(\delta x)_e} \tag{4.36}$$

where A_e is the face area of the finite volume between the grid point and its eastern neighbour and δx_e is the distance between those two points.
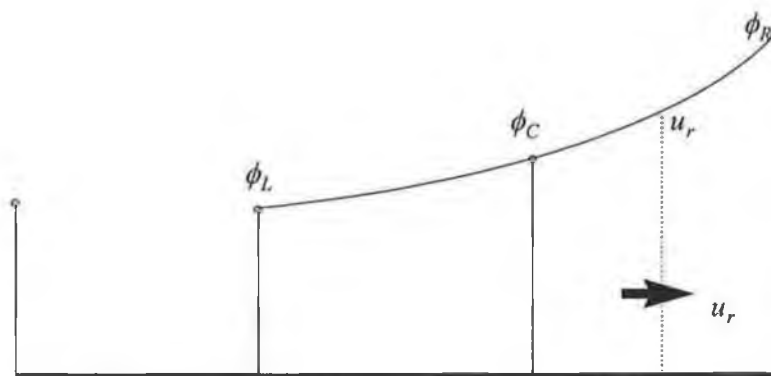
Figure 4.7. Quadratic upstream interpolation for ϕ_r.

The first formulation of Leonard's scheme used in this thesis is that due to Pollard and Siu (1982). The equations for two dimensions are stated here in full for Cartesian co-ordinates and the formulation was extended to three dimensions in the calculations for this thesis. The superscript ϕ indicates that the equations are applied to each of the dependent variables.

$$a_E^\phi = \left\| (D_e^\phi + \frac{3}{4}F_e^\phi), (D_e^\phi - \frac{3}{4}F_e^\phi - \frac{1}{8}F_w^\phi), (D_e^\phi - \frac{3}{4}F_e^\phi), (D_e^\phi + \frac{3}{4}F_e^\phi - \frac{1}{8}F_w^\phi) \right\| \quad (4.37)$$

$$a_W^\phi = \left\| (D_w^\phi - \frac{3}{4}F_w^\phi), (D_w^\phi + \frac{3}{4}F_w^\phi + \frac{1}{8}F_e^\phi), (D_w^\phi + \frac{3}{4}F_w^\phi), (D_w^\phi - \frac{3}{4}F_w^\phi + \frac{1}{8}F_e^\phi) \right\| \quad (4.38)$$

$$a_N^\phi = \left\| (D_n^\phi + \frac{3}{4}F_n^\phi), (D_n^\phi - \frac{3}{4}F_n^\phi - \frac{1}{8}F_s^\phi), (D_n^\phi - \frac{3}{4}F_n^\phi), (D_n^\phi + \frac{3}{4}F_n^\phi - \frac{1}{8}F_s^\phi) \right\| \quad (4.39)$$

$$a_S^\phi = \left\| (D_s^\phi - \frac{3}{4}F_s^\phi), (D_s^\phi + \frac{3}{4}F_s^\phi + \frac{1}{8}F_n^\phi), (D_s^\phi + \frac{3}{4}F_s^\phi), (D_s^\phi - \frac{3}{4}F_s^\phi + \frac{1}{8}F_n^\phi) \right\| \quad (4.40)$$

$$a_P^\phi = a_E^\phi + a_W^\phi + a_N^\phi + a_S^\phi \quad (4.41)$$

$$S_a^u = C^u(p_P - p_E) - \frac{1}{8}(M_w^+ F_w^u u_{WW} + M_s^+ F_s^u u_{SS} - M_e^- F_e^u u_{EE} - M_n^- F_n^u u_{NN})$$

$$- \frac{9}{8}(M_e^+ F_e^u u_E + M_n^+ F_n^u u_N - M_w^- F_w^u u_W - M_s^- F_s^u u_S) \quad (4.42)$$

$$S_a^v = C^v(p_P - p_N) - \frac{1}{8}(M_w^+ F_w^v v_{WW} + M_s^+ F_s^v v_{SS} - M_e^- F_e^v v_{EE} - M_n^- F_n^v v_{NN})$$

$$-\frac{9}{8}(M_e^+ F_e^v v_E + M_n^+ F_n^v v_N - M_w^- F_w^v v_W - M_s^- F_s^v v_S) \qquad (4.43)$$

$$S_b^u = \frac{1}{8}(M_e^+ F_u^u + M_n^+ F_n^u - M_w^- F_w^u - M_s^- F_s^u)$$

$$+\frac{9}{8}(M_s^+ F_s^u + M_w^+ F_w^u - M_e^- F_e^u - M_n^- F_n^u) \qquad (4.44)$$

$$S_b^v = \frac{1}{8}(M_e^+ F_u^v + M_n^+ F_n^v - M_w^- F_w^v - M_s^- F_s^v)$$

$$+\frac{9}{8}(M_s^+ F_s^v + M_w^+ F_w^v - M_e^- F_a^v - M_n^- F_n^v) \qquad (4.45)$$

where

$$M_i^+ = \frac{F_i^\phi + \left| F_i^\phi \right|}{2F_i^\phi}, \quad i = n, e, w, s \qquad (4.46)$$

and

$$M_i^- = \frac{F_i^\phi - \left| F_i^\phi \right|}{2F_i^\phi}, \quad i = n, e, w, s \qquad (4.47)$$

This set of coefficients gives positive values, regardless of the magnitude and direction of the convective effect. They constitute the QUICK extended form or QUICKE.

The source terms S_b^ϕ can assume positive values and may render a negative value to the central coefficient of the discretised equation. To avoid such an occurrence, the source terms are revised as follows:

$$S_a^{\phi^*} = S_a^\phi + S_b^\phi \phi_P^* \qquad (4.48)$$

103

and

$$S_b^{\phi*} = 0 \qquad (4.49)$$

where * denotes the current value of the variable. The method, with this last revision, was referred to as the QUICK Extended and Revised form or QUICKER. In testing this form, Pollard and Siu (1982) found that, while the method was stable, solution times were greater for those cases where the source term revision was not made. They reported that in flow situations with regions of re-circulating fluid flow that the number of iterations required by the QUICKER method increased drastically over those needed to secure convergence using the hybrid method. In general, their results showed that the QUICKER method gave solutions that were equal to or more accurate than the hybrid scheme, especially in those situations where the flow was highly re-circulatory.

The uniform grid formulation of the Li and Baldacchino (1995) scheme can be stated as follows in the form used in this thesis. The coefficients of the discretised equation are formulated as for the first-order upwind scheme (Patankar, 1980) and the deferred correction terms for inclusion in the source terms are calculated as follows:

$$
F_E^{dc} = (0.5\phi_E - 0.5\phi_P - 1/8(\phi_E - 2\phi_P + \phi_W))F_E^+
$$
$$
+ (0.5\phi_P - 0.5\phi_E - 1/8(\phi_P - 2\phi_E + \phi_{EE}))F_E^- \qquad (4.50)
$$

$$
F_W^{dc} = (0.5\phi_W - 0.5\phi_P - 1/8(\phi_W - 2\phi_P + \phi_E))F_W^-
$$
$$
+ (0.5\phi_P - 0.5\phi_W - 1/8(\phi_P - 2\phi_W + \phi_{WW}))F_W^+ \qquad (4.51)
$$

$$
F_N^{dc} = (0.5\phi_N - 0.5\phi_P - 1/8(\phi_N - 2\phi_P + \phi_S))F_N^+
$$
$$
+ (0.5\phi_P - 0.5\phi_N - 1/8(\phi_P - 2\phi_N + \phi_{NN}))F_N^- \qquad (4.52)
$$

$$
F_S^{dc} = (0.5\phi_S - 0.5\phi_P - 1/8(\phi_S - 2\phi_P + \phi_N))F_S^-
$$

$$+ (0.5\phi_P - 0.5\phi_S - 1/8(\phi_P - 2\phi_S + \phi_{SS}))F_S^+ \tag{4.53}$$

$$F_F^{dc} = (0.5\phi_F - 0.5\phi_P - 1/8(\phi_F - 2\phi_P + \phi_B))F_F^+$$
$$+ (0.5\phi_P - 0.5\phi_F - 1/8(\phi_P - 2\phi_F + \phi_{FF}))F_F^- \tag{4.54}$$

$$F_B^{dc} = (0.5\phi_B - 0.5\phi_P - 1/8(\phi_B - 2\phi_P + \phi_F))F_B^-$$
$$+ (0.5\phi_P - 0.5\phi_B - 1/8(\phi_P - 2\phi_B + \phi_{BB}))F_B^+ \tag{4.55}$$

where

$$F_i^+ = \frac{F_i^\phi + \left| F_i^\phi \right|}{2}, \quad i = n, e, w, s \tag{4.56}$$

and

$$F_i^- = \frac{F_i^\phi - \left| F_i^\phi \right|}{2}, \quad i = n, e, w, s \tag{4.57}$$

The deferred correction source term is given by

$$c^{dc} = -F_E^{dc} + F_W^{dc} - F_N^{dc} + F_S^{dc} - F_F^{dc} + F_B^{dc}. \tag{4.58}$$

The formulation of Hayase et al. (1992), which was published as a two-dimensional scheme and is extended here to three dimensions, is given by

$$a_E = -F_e^- + D_e \tag{4.59}$$

$$a_W = F_w^+ + D_w \tag{4.60}$$

$$a_N = -F_n^- + D_n \tag{4.61}$$

$$a_S = F_s^+ + D_s \tag{4.62}$$

$$a_F = -F_f^- + D_f \tag{4.63}$$

$$a_B = F_b^+ + D_b \tag{4.64}$$

$$a_P^\phi = a_E^\phi + a_W^\phi + a_N^\phi + a_S^\phi + a_F^\phi + a_B^\phi \tag{4.65}$$

where F_i^+ and F_i^- (i = n,e,w,s,f,b) are as defined in equations (4.48) and (4.49).

The source term is given by

$$S = -S_e^+ F_e^+ - S_e^- F_e^- + S_w^+ F_w^+ + S_w^- F_w^- - S_n^+ F_n^+ - S_n^- F_n^- + S_s^+ F_s^+ + S_s^- F_s^-$$
$$- S_f^+ F_f^+ - S_f^- F_f^- + S_b^+ F_b^+ + S_b^- F_b^- \tag{4.66}$$

where

$$S_e^+ = -\tfrac{1}{8}\phi_W - \tfrac{1}{4}\phi_P + \tfrac{3}{4}\phi_E \tag{4.67}$$

$$S_e^- = -\tfrac{1}{8}\phi_{EE} - \tfrac{1}{4}\phi_E + \tfrac{3}{8}\phi_P \tag{4.68}$$

$$S_w^+ = -\tfrac{1}{8}\phi_{WW} - \tfrac{1}{4}\phi_W + \tfrac{3}{4}\phi_P \tag{4.69}$$

$$S_w^- = -\tfrac{1}{8}\phi_E - \tfrac{1}{4}\phi_P + \tfrac{3}{8}\phi_W \tag{4.70}$$

$$S_n^+ = -\tfrac{1}{8}\phi_S - \tfrac{1}{4}\phi_P + \tfrac{3}{4}\phi_N \tag{4.71}$$

$$S_n^- = -\tfrac{1}{8}\phi_{NN} - \tfrac{1}{4}\phi_N + \tfrac{3}{8}\phi_P \tag{4.72}$$

$$S_s^+ = -\tfrac{1}{8}\phi_{SS} - \tfrac{1}{4}\phi_S + \tfrac{3}{4}\phi_P \qquad (4.73)$$

$$S_s^- = -\tfrac{1}{8}\phi_N - \tfrac{1}{4}\phi_P + \tfrac{3}{8}\phi_S \qquad (4.74)$$

and four further terms, for the front and back faces of the three-dimensional cell, follow the same pattern of subscripts.

4.10 Boundary Conditions

At the boundaries of the computational domain, the general discretised equation does not apply and, hence, special measures are required for the cells next to the boundaries. Boundary conditions are implemented in this work by substituting known information about the boundary node into the discretised equation for the cell adjacent to the boundary. The grid arrangement is such that the boundaries coincide with the cell walls and the boundary information is handled by setting the boundary coefficient to zero and incorporating the information into the central coefficient and source terms of the adjacent cell. Fixed value (Dirichlet) conditions are used, e.g. setting an inlet plug velocity by specifying the actual boundary value required.

In order to ensure consistent treatment of the boundaries in a QUICK scheme the interpolation of the boundary values in the formulation of equation coefficients requires special attention because the scheme may require a value of the dependent variable outside the calculation domain. The most consistent is the third order treatment (Leonard, 1988 and Hayase et al., 1992) where the point that is referenced outside the calculation domain is estimated by, for example,

$$\phi_o = 3\phi_b - 3\phi_1 + 4\phi_2 \qquad (4.75)$$

where the subscripts are o for the value to be estimated, b for the boundary value, 1 and 2 for the first and second points inside the calculation domain. The use of these interpolations has been noted by Hayase et al. (1992) as giving rise to some instability in the solution of the flow equations. The corresponding second order treatment would estimate the same point as

$$\phi_o = 2\phi_b - \phi_1 \qquad (4.76)$$

Hayase et al. (1992) demonstrated the superior accuracy of the third-order treatment. It is incorporated in the source terms of the QUICK formulations used in this thesis.

The k-ε model as outlined in Chapter 3 is for high Reynolds number flows and requires some additional work to adapt it for more accurate description near boundaries. The wall function method as used in TEACH was adopted for this study for convenience and as a reasonable starting point for use in the testing of the new three-dimensional versions of the CELS solver covered in a later chapter.

The wall function method is due to Launder and Spalding (1974) and has been applied in a application similar to the mushroom one with some success by Choi et al. (1988) and Choi et al. (1990) to animal housing structures.

In the wall function method the dimensionless distance (local Reynolds number) from the wall to the first grid point is evaluated. If this distance y^+ is less than 11.63 then it is assumed to be purely viscous and if greater than 11.63 then it is treated as purely turbulent.
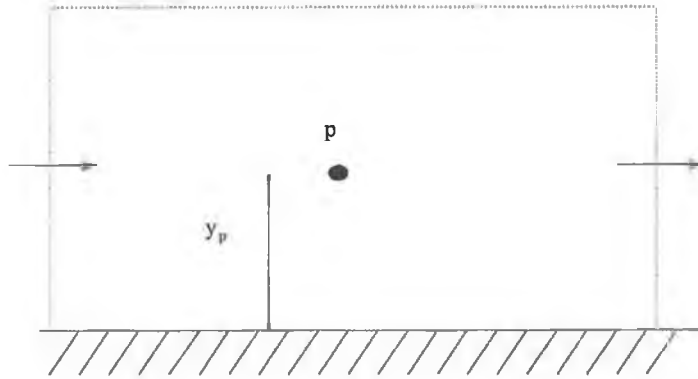
Figure 4.8. Conditions for a near wall grid point p.

The momentum equations take the following form in these two situations

$$\text{For } y^+ \leq 11.63 : \quad u^+ = y^+ \tag{4.77}$$

$$\text{For } y^+ > 11.63 : \quad u^+ = \ln\frac{Ey^+}{\kappa} \tag{4.78}$$

where

$$y^+ = \frac{yu}{\nu}, \tag{4.79}$$

$$u^+ = \frac{U_p}{u_*}, \tag{4.80}$$

and

$$u_* = \left(\frac{\tau_w}{\rho}\right)^{1/2}. \tag{4.81}$$

The values $\kappa = 0.4$ and $E = 5.1$ were used with success by Choi et al. (1988 and 1990) for their slot flows which were similar to the type of flow encountered in mushroom growing. For this reason these values are also adopted here.

In handling the boundary conditions for the k equation the coefficient for the wall face of the boundary control volume is set to zero. This is because the turbulence energy goes to zero at the wall and there is therefore no flux contributed.

The source $G - c_D \rho \varepsilon$ is incorporated in the linearised source term. The full generation term from the k equation is modified using the wall shear stress

$$\text{For } y^+ \leq 11.63 : \quad \tau_w = \mu \frac{u_p}{y_p} \tag{4.82}$$

$$\text{For } y^+ > 11.63 : \quad \tau_w = \rho \kappa c_\mu^{1/4} k_p^{1/2} \frac{u_p}{\ln(Ey^+)} . \tag{4.83}$$

To complete, the source term ε is calculated from

$$\text{For } y^+ \leq 11.63 : \quad \varepsilon = \frac{c_\mu^{3/4} k^{3/2}}{y_p} y^+ \tag{4.84}$$

and

$$\text{For } y^+ > 11.63 : \quad \varepsilon = \frac{c_\mu^{3/4} k^{3/2}}{\kappa y_p} \ln(Ey^+) . \tag{4.85}$$

In wall flows, dissipation rate ε reaches its highest value at the wall. In order to overcome difficulties in modifying the coefficient for the control volume face coincident with the wall, a fixed value of ε is set by modifying the source coefficients to insert the value. Local equilibrium is used to estimate this as

$$\varepsilon = c_\mu^{3/4} \frac{k_p^{3/2}}{\kappa y_p} \qquad (4.86)$$

where k_p and y_p are the turbulence kinetic energy and distance from the wall at the grid point for the boundary control volume.

In this work the inlet treatment for the turbulence parameter equations as used in the TEACH package was retained. Choi (1991) used this prescription for his calculations and the similarity of the situations modelled and the results would indicate that this choice is reasonable for the current work. The inlet values are estimated as

$$k_{in} = 0.04u_{in}^2 \qquad (4.87)$$

and

$$\varepsilon_{in} = \frac{k_{in}^{3/2}}{0.005L} \qquad (4.88)$$

The review by Nallasamy (1987) comments on the lack of a standard approach to inlet boundary descriptions for the turbulence parameters in the k-ε model for the calculation of turbulence parameters. He notes that, in particular, the inlet profiles of k and ε can have a significant effect on the flow downstream and can lead to wrong conclusions about the performance of a turbulence model. Leschziner and Rodi (1984) note, in a study of axisymmetric free jets, that inlet conditions for k and ε play as crucial a role in achieving predictive accuracy as turbulence modelling details.

Chapter 5

The Coupled Equation Line Solver (CELS) Technique

5.1 Selection of the CELS method

The selection of a numerical method to replace SIMPLE as the solver for the discretised equations reflects both an initial aim of this work to apply the multigrid acceleration technique (see Chapter 6) and the work that was already done in extending the TEACH code to three dimensions to allow modelling of mushroom growing rooms. This latter part of the work is outlined in Chapter 8 where attention is drawn to long convergence times using the SIMPLE and the role of the pressure/velocity coupling in giving rise to these.

The structure of the SIMPLE method was an important factor in determining its efficiency. SIMPLE can be described as a sequential or segregated solver because each of the equations for the dependent variables are solved in turn. Alternative segregated methods to SIMPLE, such as SIMPLEC (Van Doormaal et al., 1984) and SIMPLER (Patankar, 1980), improve the pressure-velocity coupling and solver efficiency (see Section 4.9) but seemed unlikely to allow successful application of multigrid methods to the solution of the linear coefficient sets within each coefficient update cycle because the computational effort per coefficient update was still small compared to the total required to achieve convergence. Miller and Schmidt (1988) showed that the benefits of multigrid applied to inner iterations alone could be limited depending on the inner solver used and on pressure-velocity coupling.

Simultaneous methods (velocities and pressure solved together) give rise to a reduced number of coefficient updates and greatly increased effort in solving the linear sets of equations provided by the updates. This would allow more scope for improvements in the convergence time by multilevel acceleration. Smith (1990) has reported on the successful application of a point-by-point solver and multigrid method with local refinement to reduce convergence time for the fluid flow equations treated in this thesis. The same equation set and the hybrid scheme were used but calculations were carried out in a domain where the grid points for all the variables were coincident (co-located grid). Some of these methods were developed from the well-known work of Caretto et al. (1972).

The requirement to preserve as much as possible of the work that had been done in the extension of the TEACH code to three dimensions placed a considerable constraint on the alternative solution methods that could be considered. Only methods based on the staggered grid were examined. Galpin et al. (1985) presented a solver based on a staggered grid with the same formulation of the fluid flow equations. This Coupled Equation Line Solver (CELS) produced significant improvements in convergence time relative to SIMPLE and was competitive with SIMPLEC. Convergence time improvements relative to SIMPLEC overall were problem dependent and grid density dependent but an important feature of the method was the reduction in the number of coefficient updates. Updates were reduced from several hundred to less than fifty. For broadly the same computational effort, this means that the method allows greater effort per coefficient update and more potential for improving the performance of the inner solver for the linear sets. Multigrid methods have also been applied to this method for laminar flows and significant acceleration was achieved by the Additive Correction Multigrid (ACM) approach (Hutchinson et al., 1988). This technique is discussed in chapter 6. The work published on this method was restricted to two dimensions and laminar flows.

5.2   Derivation of the equations for the CELS method

The direction of the staggering relative to the scalar grid (figure 5.1) was opposite to that used by Galpin et al. (1985) but as this was simply a 180 degree rotation of the grid it was not expected to cause any difficulty. The solved equations were to be reformulated for this arrangement. In view of the amount of work done that had been done in extending TEACH, this was considerably more convenient than altering the original code.
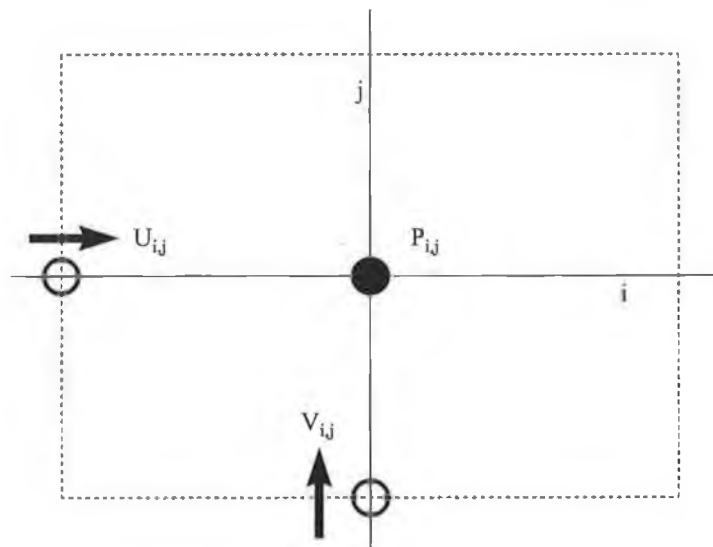

Figure 5.1.   Staggering of grid for derivation of CELS equations.

The use of under-relaxation factors as set out in chapter 3 and the hybrid and QUICK differencing schemes were employed for this work whereas Galpin et al. (1985) used upstream-weighted differencing (Raithby and Torrance, 1974) and a different, although algebraically equivalent, formulation for the under-relaxation factor.

Thus, the equations that formed the starting point were the two-dimensional fluid flow equations as used in the development of the TEACH code. These are

$$A_P^u u_{ij} = A_E^u u_{i+1j} + A_W^u u_{i-1j} + A_N^u u_{ij+1} + A_S^u u_{ij-1} - C^u(p_{ij} - p_{i-1j}) + b^u, \qquad (5.1)$$

$$A_P^v v_{ij} = A_E^v v_{i+1j} + A_W^v v_{i-1j} + A_N^v v_{ij+1} + A_S^v v_{ij-1} - C^v(p_{ij} - p_{ij-1}) + b^v \qquad (5.2)$$

and

$$0 = A_E^c u_{i+1j} + A_w^c u_{ij} + A_N^c v_{ij+1} + A_S^c v_{ij} \qquad (5.3)$$

where the superscript c refers to the continuity equation. At constant j, i.e. solving along a line of constant j, the off-line terms can be gathered into a modified source term so that

$$A_P^u u_{ij} = A_E^u u_{i+1j} + A_W^u u_{i-1j} + C^u(p_{i-1j} - p_{ij}) + B^u, \qquad (5.4)$$

$$A_P^v v_{ij} = A_E^v v_{i+1,j} + A_W^v v_{i-1j} - C^v p_{ij} + B^v, \qquad (5.5)$$

and

$$0 = A_E^c u_{i+1j} + A_w^c u_{ij} + A_S^c + B^c, \qquad (5.6)$$

where the source terms are given by

$$B^u = b^u + A_N^u u_{ij+1} + A_S^u u_{ij-1}, \qquad (5.7)$$

$$B^v = b^v + A_N^v v_{ij+1} + A_S^v v_{ij-1} + C^v p_{ij-1}, \qquad (5.8)$$

and

$$B^c = A_N^c v_{ij+1}. \qquad (5.9)$$

The continuity equation (5.3) can be used to derive expressions for $v_{i,j}$, $v_{i+1,j}$ and $v_{i-1,j}$. These are then used to eliminate v in equation (5.2), leaving the pressure expressed in terms of u. Manipulation of the resulting equation yields the convenient form

$$A_P^p p_{ij} = A_{EE}^p u_{i+2j} + A_E^p u_{i+1j} + A_W^p u_{ij} + A_{WW}^p u_{i-1j} + B^p \qquad (5.10)$$

where the coefficients without subscripts refer to the cell labelled ij and

$$A_P^p = -C^v, \qquad (5.11)$$

$$A_{EE}^p = \frac{A_E^v A_{Ei+1j}^c}{A_{Si+1j}^c}, \qquad (5.12)$$

$$A_E^p = \frac{A_E^v A_{Wi+1j}^c}{A_{Si+1j}^c} - \frac{A_P^v A_E^c}{A_S^c}, \qquad (5.13)$$

$$A_W^p = \frac{A_W^v A_{Ei-1j}^c}{A_{Si-1j}^c} - \frac{A_P^v A_W^c}{A_S^c}, \qquad (5.14)$$

$$A_{WW}^p = \frac{A_W^v A_{Wi-1j}^c}{A_{Si-1j}^c}, \qquad (5.15)$$

and

$$B^p = \frac{A_E^v}{A_{Si+1j}^c} B_{i+1j}^c - \frac{A_P^v}{A_S^c} B^c + \frac{A_W^v}{A_{Si-1}^c} B_{i-1}^c - B^v. \qquad (5.16)$$

Equation (5.10) can now be used to eliminate p in equation (5.1), yielding an expression for u in terms of its neighbouring cells, namely,

$$a_P^u u_{ij} = a_{EE}^u u_{i+2j} + a_E^u u_{i+1j} + a_W^u u_{i-1j} + a_{WW}^u u_{i-2j} + b^{*u} \qquad (5.17)$$

where

$$a_P^u = A_P^u + C^u \frac{A_W^p}{A_P^p} - C^u \frac{A_{Ei-1j}^p}{A_{Pi-1j}^p}, \qquad (5.18)$$

$$a_{EE}^u = -C^u \frac{A_{EE}^p}{A_P^p}, \qquad (5.19)$$

$$a_E^u = A_E^u + C^u \frac{A_{EEi-1i}^p}{A_{Pi-1j}^p} - C^u \frac{A_E^p}{A_P^p}, \qquad (5.20)$$

$$a_W^u = A_W^u + C^u \frac{A_{Wi-1j}^p}{A_{Pi-1j}^p} - C^u \frac{A_{WW}^p}{A_P^p}, \qquad (5.21)$$

$$a_{WW}^u = -C^u \frac{A_{WWi-1j}^p}{A_{Pi-1j}^p}, \qquad (5.22)$$

and

$$b^{\circ u} = B^u - C^u \frac{B^p}{A_P^p} + C^u \frac{B_{i-1j}^p}{A_{Pi-1j}^p}. \qquad (5.23)$$

The penta-diagonal form of equation (5.17) lends itself to an efficient solution by the use of a standard technique. These equations show the expected symmetry with those derived by Galpin et al. (1985). As in their work, a penta-diagonal solver of the form presented in the following section was used to solve the u equation.

5.3   Solution Procedure

A general penta-diagonal system of equations can be written as follows:

117

$$a_P \phi_i = a_{EE} \phi_{i+2} + a_E \phi_{i+1} + a_W \phi_{i-1} + a_{WW} \phi_{i-2} + b_p \qquad (5.24)$$

and it is solved by manipulating it into the upper diagonal (recursive) form

$$\phi_i = A_i \phi_{i+2} + B_i \phi_{i+1} + C_i \qquad (5.25)$$

where

$$A_i = a_{EE} / F \qquad (5.26)$$

$$B_i = (a_E + DA_{i-1}) / F \qquad (5.27)$$

$$C_i = (b_P + DC_{i-1} + a_{WW} C_{i-2}) / F \qquad (5.28)$$

$$D = a_W + a_{WW} B_{i-2} \qquad (5.29)$$

$$F = a_P - a_{WW} A_{i-2} - DB_{i-1}. \qquad (5.30)$$

With the values of u derived along the line of constant j, values of v can be calculated from the continuity equation (5.3) and the pressure value from equation (5.10).

The method consists of line-by-line sweeps in the solution domain. Each line solution for u using the penta-diagonal solver provides information for an improved estimate of the v velocity component and the pressure for use in the calculation of the following line in the sweep. These updated values enter the calculations in the modified source terms. The sweep proceeds from the line corresponding to the highest value of j inside the calculation domain and proceeds downwards to the line $j = 3$.

On the last line of the sweep through the plane ($j = 2$), a special procedure is required because the v velocities are defined by boundary conditions (see figure 5.2) and not, as for the others, by a discretised equation. The values for v are known from the

118

boundary conditions and then the u values are determined by mass conservation. The pressure then needs a special procedure and it is calculated from equation (5.2) for the v velocity component.
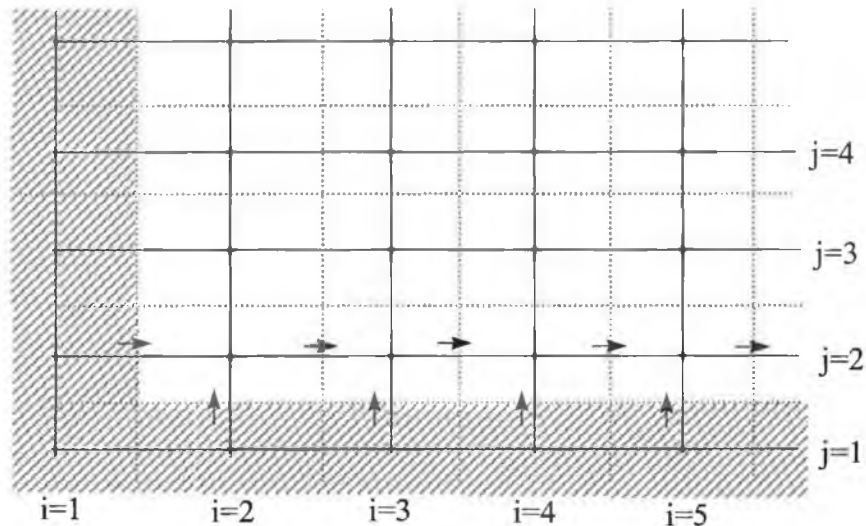
Figure 5.2. Velocity components for the last line solution of a sweep through horizontal lines of grid points.

A similar derivation along the lines of constant i yields penta-diagonal equations for the v velocity components. Values for the u velocity components and the pressure can be calculated from the continuity equation and the pressure equation in terms of v. These equations enable the calculation domain to be swept in two directions.

CELS is operated as an iterative procedure, where a set of coefficients is generated for the discretised equations, starting with initial values of zero for the velocity components and the pressure, and these are then solved by sweeping the calculation domain using the equations just derived and updating the dependent variables line by line. In the calculations, a residual reduction factor less than unity (α) is set such that the sweeping of the calculation domain, after each coefficient update, continues until the initial residual is reduced by that factor. The coefficients of the discretised

119

equations are then updated. The CELS equations are under-relaxed within the CELS procedure in the same way as the discretised equations are in the outer, coefficient update procedures (Section 4.6).

Chapter 6

Additive Correction Multigrid for the Flow Equations

6.1   Fundamental concepts of multigrid methods

Multigrid methods arise essentially from the spectral or Fourier mode picture of relaxation procedures. From this viewpoint, errors in interim solutions produced by iterative methods can be regarded as consisting of a range of frequency components. High frequency errors are associated with the solution between adjacent grid points while the low frequency components are associated with the errors across sections of or the entire calculation domain.

Many standard iterative schemes possess what is called the smoothing property. This makes these methods very effective at reducing the high-frequency or oscillatory components of the error, while leaving the low frequency or smooth components relatively unchanged.

The smoothing property is reflected in the classic path to a converged solution. Many methods converge quickly at first but then the rate of convergence slows. In the work presented here this is seen as an initially rapid drop in the magnitude of the residuals that are used to monitor convergence and then the slowing down of residual reduction. More effort is expended in producing a given reduction in residuals as the solution procedure progresses. The initial rapid drop is due to the efficient removal of oscillatory modes of the error. Once these are removed the relaxation of the smooth components is far less effective.

The smoothing property is a serious limitation of these methods but associating the frequency of the error components for which the methods are efficient with the grid spacing in the numerical method suggests the remedy that leads to the multigrid methodologies. The problem becomes one of finding a way of using variable grid spacing so that the relaxation methods can be adapted to work on all the error modes or frequencies with the same efficiency.

6.2   Removal of smooth error components

With the idea that a given method is effective on error components related to the grid spacing it seems that operating on a larger grid spacing (a coarser grid) would move the smooth components into the range of effectiveness of the solver. For example, a good way to improve the early stages of a relaxation process is to start with a good initial guess. This often comes from a calculation on a coarser grid. The results of this are then interpolated onto the fine grid.

Assume that a particular method has been applied until only the smooth components remain. How do these components look on a coarser grid? Figure 6.1 shows such a situation (adapted from an excellent introductory text by Briggs (1987)).

A smooth wave on the fine grid is projected directly onto a coarser one with twice the grid spacing. It can be seen that using the solver on the coarser grid is now effectively working with a more oscillatory component. In principle the error can now be made smooth quickly by removing this oscillatory component on the coarse grid. Further, when the solver runs into difficulty with smooth components on this coarser grid, it will be possible to project on down to another grid with a greater spacing.
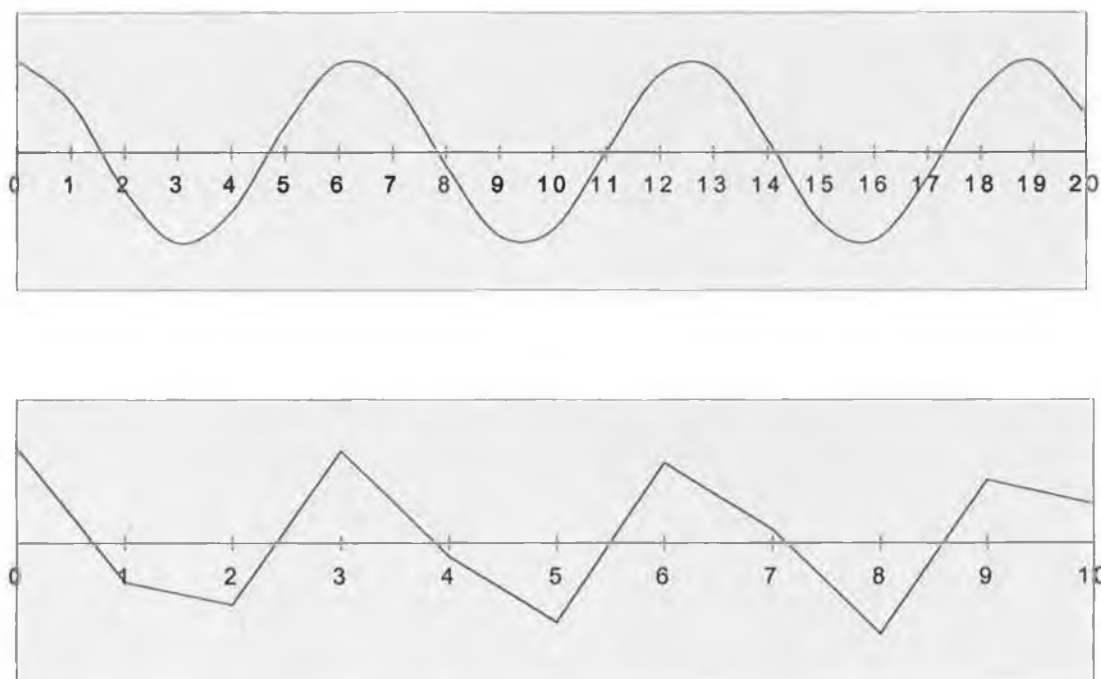
Figure 6.1. Projection of a smooth error component onto a coarser grid.

These deceptively simple observations supply the general principle of multigrid or multilevel procedures. The question to be answered is what is the procedure for transferring information between grids.

6.3   The residual equation

Information on the error is required before we can proceed. But if we have an expression for the error then the problem is solved. Attention in these methods is directed to what is called the residual equation because of its incorporation of the error.

Denote a system of linear equations by

$$\mathbf{Au} = \mathbf{f} \tag{6.1}$$

Let **u** denote the exact solution and **v** is an approximation to this which is generated by the solution method. The algebraic error is a measure of the usefulness of **v** as an approximation. This is given by

$$e = u - v \tag{6.2}$$

which in itself contains no useful information but it is related to the residual, given by

$$r = f - Av \tag{6.3}$$

which is a computable measure of the amount by which v fails to satisfy the original equation for **u**. By uniqueness, $r = 0$ if and only if $e = 0$ but the norm of these quantities are not necessarily related. Manipulation of these relationships yields an important relation between the error and the residual. It is

$$Ae = r \tag{6.4}$$

and this is called the residual equation. This says that when **f** is replaced by the residual then **e** plays the same role as **u**.

Another important principal motivates the use of the residual equation. This is that operating on the original equation with an initial guess **v** is equivalent to using the residual equation with an initial guess $e = 0$. Approximations for **e** can be obtained by this method and these errors can be used as corrections for **v**. They should combine with the calculated **v** to improve its approximation to **u**.

6.4　Multigrid solution procedures

A principal strategy incorporates the use of the residual equation to relax on the error. The method employed in this thesis uses a very-closely related technique. The procedure is generally composed of cycles or combinations of cycles that may be described as follows:

(i)　Relax on $A\mathbf{u} = \mathbf{f}$ on the fine grid to obtain an approximation $\mathbf{v}$.

(ii)　Compute the residual $\mathbf{r} = \mathbf{f} - A\mathbf{v}$.

(iii)　Transfer A and $\mathbf{r}$ to a coarser grid.

(iv)　Relax on the residual equation $A\mathbf{e} = \mathbf{r}$ on the coarse grid.

(v)　Use the coarse grid error estimates to correct $\mathbf{v}$ on the fine grid.

(vi)　Return to step (i) unless a convergence criterion is satisfied.

This procedure is the basis of coarse grid correction. Relaxation on the fine grid (level 0) is continued until it ceases to be effective, i.e. the oscillatory components of the error are removed. The problem is then transferred to the coarse grid (level 1) and solved sufficiently to generate useful corrections for the fine grid approximation.

The error calculated on the coarse grid can become smooth enough to slow convergence and at this stage the procedure above can be nested. In this situation transfer to a still coarser grid (level 2) is implemented and corrections for $\mathbf{e}$ on level 1 are generated by a residual equation on level 2. Nesting can continue until a level is reached where smoothing does not significantly slow convergence.

Nesting gives rise to a number of solution strategies such as the V and W cycle illustrated in figures 6.2 and 6.3.
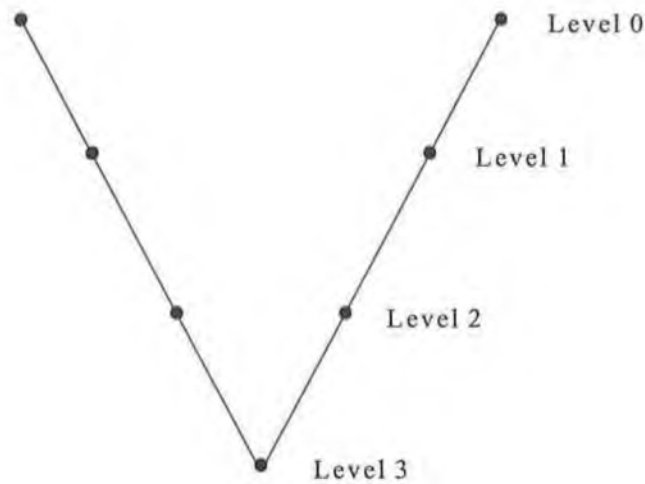
Figure 6.2. Multigrid V cycle.

The V cycle shows a simple strategy that begins on the finest grid (level 0) on the top left of the figure and moves down after smoothing to the next coarsest grid which is level 1. When this in turn becomes smooth, the problem is transferred to level 2 and finally, in this example to level 3. At the coarsest level, consisting of a small number of grid points, a direct solver is often applied. The corrections derived are now passed up a level, correcting the results calculated for that level and usually a number of iterations are carried out to take advantage of the temporary improvement in efficiency that is provided by the introduction of corrections. Corrections are passed up each level in turn until the finest grid is corrected and an improved solution is produced on this one. If convergence is not achieved at the finest level then the V-cycle can be repeated.
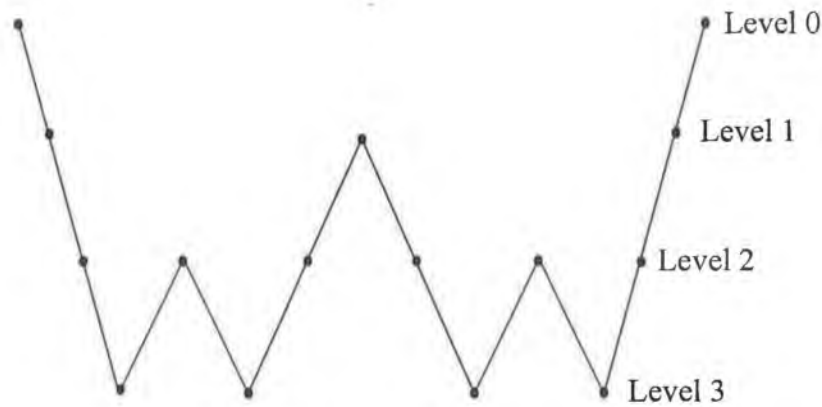
Figure 6.3. Multigrid W cycle.

The W-cycle shows the sort of action that is used in this thesis. The problem begins as in the V-cycle. As each level becomes smoother, the problem is transferred down to the next, more efficient operating, grid. In this case there is a requirement to produce a certain degree of convergence at each level. This could be specified as a reduction of some proportion of the initial residual measure.

When the procedure returns the first time to level 2, the iterative process continues to reduce the residuals on this level. If the errors again become smooth (residual reduction efficiency decreases) then a transfer to level 3 is required again. On the next transfer to level 2 the residual reduction criterion is met and corrections are applied to level 1. On level 1 the residual reduction criterion is not met before there is a drop in efficiency and so there is a requirement to move down to levels 2 and 3 again. After some time repeating this sort of action, convergence criteria are met on each level in turn, after correction, and activity on the finest grid is resumed. As with the V cycle, the entire procedure may need to be repeated to achieve convergence on the finest level.

What is required now is a method for transferring the original problem down to a coarser grid and for moving corrections up to finer grids.

6.5 Prolongation and restriction operators

First consider restriction operators. The process is generally called interpolation or prolongation. Only the simplest case of linear interpolation is shown here but there is a general opinion in the multigrid literature that more complex prolongation is rarely justified. For grids in which the coarse one has twice the spacing of the fine one, we can define a linear interpolation operator to move from a coarse (c) to a fine (f) grid by

$$I_c^f \mathbf{v}^c = \mathbf{v}^f , \tag{6.5}$$

where

$$\mathbf{v}_{2j}^f = \mathbf{v}_j^c \tag{6.6}$$

$$\mathbf{v}_{2j+1}^f = \frac{1}{2}(\mathbf{v}_j^c + \mathbf{v}_{j+1}^c) . \tag{6.7}$$

Under this operator, values for even-numbered fine grid points are transferred directly from the coarse grid. Odd-numbered fine grid points use the average of the adjacent coarse grid values. Interpolation will work well if the real error on the fine grid is smooth. In this case interpolation of a good coarse grid approximation will yield reasonable results. If the real error, however, is oscillatory then the interpolant from a good coarse grid result may not be accurate. In the light of this, sufficient smoothing has to be achieved initially before the multigrid method can work effectively.

Moving from a fine to a coarse grid is accomplished by what is generally called a restriction operator. The simplest restriction operator is injection is given by

$$I_f^c \mathbf{v}^f = \mathbf{v}^c \tag{6.8}$$

$$\mathbf{v}_j^c = \mathbf{v}_{2j}^f. \tag{6.9}$$

The coarse grid simply takes its value directly from the corresponding fine grid point. An alternative and popular prescription is call the full weighting operator and it is given by

$$\mathbf{v}_j^c = \frac{1}{4}(\mathbf{v}_{2j-1}^f + 2\mathbf{v}_{2j}^f + \mathbf{v}_{2j+1}^f). \tag{6.10}$$

Coarse grid values are given by weighted averages of the values at neighbouring fine grid points. It is interesting and instructive to note that smoothing is required for accurate transfers in both directions.

6.6 Multigrid schemes

Coarse grid equations obtained using these techniques in terms of a correction to the existing fine grid solution give rise to what is known as the Correction Storage (CS) scheme. In this type of scheme the coarse grid equations are determined by the current estimates for coefficients and variables on the fine grid and corrections only are generated on the coarse grids. The CS scheme is suitable for linear problems or those non-linear, iterative methods where there is a large computational effort per coefficient update. The Full Approximation Storage (FAS) scheme was developed for the general non-linear situation. In this method the coarse grid coefficients are updated as the solution proceeds on the coarse grid. The solution matrix must be available on all grids in order to accomplish this and the definition of the coarse grid variable is the fundamental difference between the schemes. In a Full Multigrid (FMG) scheme, the initial operation is a sequence of calculations which solves the equations, beginning on the coarsest grid and moving up to the finest one.

As the method to be applied in this thesis is a variant of the CS scheme, no more attention will be given to the FAS or other schemes. There is now a vast literature on the subject of multigrid methods and a good proportion of it is devoted to flow problems. Good surveys which could now be viewed as introductions to the general application of the multigrid method are to be found in publications by Brandt (e.g. 1984 and 1987), one of the originators of the entire field.

In Ghia et al. (1982), the vorticity-stream function formulation of the two-dimensional incompressible Navier-Stokes equations is used to study the effectiveness of the coupled strongly-implicit multigrid (CSI-MG) method in the determination of high Reynolds number, fine mesh flow solutions. The driven flow in a square cavity is used as the model problem. Solutions are obtained for configurations with Re as high as 10,000 and meshes of as many as 257×257 points. The FAS scheme was employed and a four-fold increase in efficiency was found.

A multigrid method for finding stationary solutions of the Euler equations was described by Mulder (1985) and tested with spatial discretization obtained by upwind differencing and implicit time discretization. The Correction Storage scheme was employed and gains in efficiency depended on the grid density.

The accuracy of many turbulent flow and convective heat transfer problems may be affected by the presence of regions containing large gradients of the dependent variables. Philips and Schmidt (1984) successfully applied a combination of local grid refinement and multigrid methods. The paper presented investigations of interpolation methods between differing grid densities and accuracy and computational requirements. In Philips and Schmidt, 1985a, a multigrid concept of local grid refinement was applied as a method for improving the accuracy of the calculations. In combination with the multilevel acceleration techniques, an accurate and efficient computational procedure was developed. In addition, a robust implementation of the QUICK finite differencing scheme was described. Calculations of a test problem were presented to quantitatively demonstrate the advantages of the multilevel-multigrid method. The QUICK scheme was implemented on the finest level only, the argument being that the function of the

finer meshes was error smoothing for the accurate finest grid and that finest level source terms were contained in residuals on a coarser one. The QUICK implementation was found to be more accurate than the hybrid scheme and nodes adjacent to the boundary were handled by applying the hybrid scheme to these (having found this to be superior to the use of the upwind scheme).

A combination of local grid refinement with multilevel operations and the QUICK differencing scheme with the PISO (pressure implicit split operator scheme) showed significant improvement (25%) in total CPU timings for a lid-driven cavity (Philips and Schmidt, 1985b). A buoyancy-driven flow was solved with just 15% of the effort required for PISO alone. The authors noted that agreement with benchmark solutions was not as good as between the improved methods tested but did not comment on their acceptability.

A procedure applicable to the finite-difference formulations using staggered grids was presented by Vanka (1986a). A smoothing technique called symmetrical coupled Gauss-Seidel was proposed and empirically observed to give good smoothing rates. FAS-FMG multigrid was used and, while detailed comparisons with other methods were not made, Vanka observed that at Re=1000 and a grid of 81×81 nodes the method was fifty times faster than SIMPLE. An efficient finite-difference calculation procedure for three-dimensional re-circulating flows was presented by Vanka (1986b). The algorithm was based on Vanka (1986a) and the symmetrical coupled Gauss-Seidel technique was observed to provide good rates of smoothing. Calculations were made of the fluid motion in a three-dimensional cubic cavity with a moving top wall. The efficiency of the method was demonstrated by performing calculations at different Reynolds numbers with finite difference grids as large as $66 \times 66 \times 66$ nodes. The CPU times and storage requirements for these calculations were observed to be very modest.

Barcus et al. (1987) applied multilevel methods to a SIMPLE-based solution method. FMG was employed and the novelties of the approach were a co-located variable arrangement and the use of a Correction Scheme instead of the more often used Full Approximation Scheme. The SIMPLE algorithm was employed on all grid levels for the pressure velocity coupling although operating on residual equations rather than the

131

primitive variables. Test cases showed considerable savings. Interestingly, the authors showed the contribution that the initial coarse grid estimates (prolongation scheme) for the first fine calculations made to the savings by completing the solutions on the fine grid with and without the telescoping multigrid steps.

Miller and Schmidt (1988) showed that the benefits of multigrid applied to the inner iterations in an iterative procedure for the flow equations could be limited, depending on the inner solver used and on pressure-velocity coupling.

Shyy and Sun (1993) developed a successful FAS scheme application for staggered-grid based solvers for incompressible re-circulating flows. The algorithm was implemented in conjunction with a pressure correction method using curvilinear co-ordinates. Grid-generation and adaptive grids were also used. The multigrid technique produced impressive results in all problems examined.

6.7 The development of additive correction multigrid (ACM)

In the classifications of multigrid methods currently in use, Additive Correction Multigrid (ACM) can be viewed as a special case of the Correction Storage system. It has been developed for and is used as an inner solver, i.e. it operates on the solution of a given set of equations after the coefficients have been updated and the main gain from its use is in the reduction of the effort per coefficient update. It can also reduce the number of coefficient update cycles by accelerating the solution of low-frequency error components.

A general formulation of additive correction methods was presented by Settari and Aziz (1973). The methods were applied to the solution of finite difference equations resulting from elliptic and parabolic partial differential equations. It was found to be competitive with other methods and they noted the correction approach could be easily applied to any iterative method. The methods as proposed by Settari and Aziz (1973) were

132

applied by Hutchinson and Raithby (1986), interpreted as a multigrid method (ACM) and shown to be effective in reducing the cost of solving an algebraic equation set for a single variable. It was then successfully applied to the coupled fluid flow equations in two dimensions, as used in CELS, by Hutchinson et al. (1988). In the two sample problems presented by Hutchinson et al., the ACM method proved very attractive.

The first problem was a two-dimensional driven cavity flow where a laminar flow within a square enclosure was driven by the lid (upper wall) sliding across the cavity at a constant speed. The solution times were reduced by a factor of 5 for a 64×64 grid and by a factor of 8.7 for 96×96 grid when compared with the performance of CELS alone. The relationship between solution time t and the number of grid points (N) was of the form $t \propto N^{\beta}$ where β was approximately 2 for CELS and 1.4 for CELS with ACM.

When ACM was applied to two-dimensional, natural convection in a cylindrical polar annulus, the performance improvements were less marked relative to CELS. On a 32×32 grid the improvement factor was 2.3 and for a 64×64 grid the result was 2.1. For this problem β was approximately 1.2 for CELS and 0.6 for CELS with ACM. On both of these grids CELS out-performed SIMPLEC.

It appeared that the less marked improvement due to ACM in the natural convection case was due to coupling between the equations at the outer solver level (the temperature-velocity coupling). More coefficient updates were required to take account of this coupling than would have been the case without it. A consequence of the outer solver coupling is that the effort required per iteration of the inner solver is reduced. A large input of computational effort in the solution of a given set of coefficients (by setting a low target figure for the residual reduction factor) is not justified when they are subject to the outer coupling. An earlier study of a CELS-based solver was an examination of the treatment of the temperature-velocity coupling by Galpin and Raithby (1986). The technique developed showed reductions in computational effort for modelling of natural convection but it seemed that much of the saving could be made by application of ACM to the unmodified CELS method applied to the same problem (Hutchinson et al., 1988).

The question to be addressed by the application of CELS in two dimensions in this thesis is what improvements can be expected relative to SIMPLE when the outer coupling is due to the action of the turbulence equations in updating the effective viscosity which is then used in the u,v,p inner solver.

6.8   Formulation of equations

For the application of the ACM method, the coarse-grid equations are obtained directly from the fine-grid ones. This is attractive for two reasons:

(i)   Discretisation is required on the fine grid only. This can be an advantage for some users as it eliminates the complexity of grid generation at coarse levels. It also eliminates the possibility of inconsistencies in discretisation at different levels.

(ii)   The important conservation property of the fine grid equations is maintained on coarser grids by this method because the coarse grid equations are combinations of the fine-grid ones. Solution of coarse grid equations yields conservative fields.

In the method as previously applied, combinations of fine grid cells are made by gathering them in 2×2 blocks as shown in figure 6.4. Note that the direction of staggering here differs from that in the references as explained in the discussion of the CELS method.
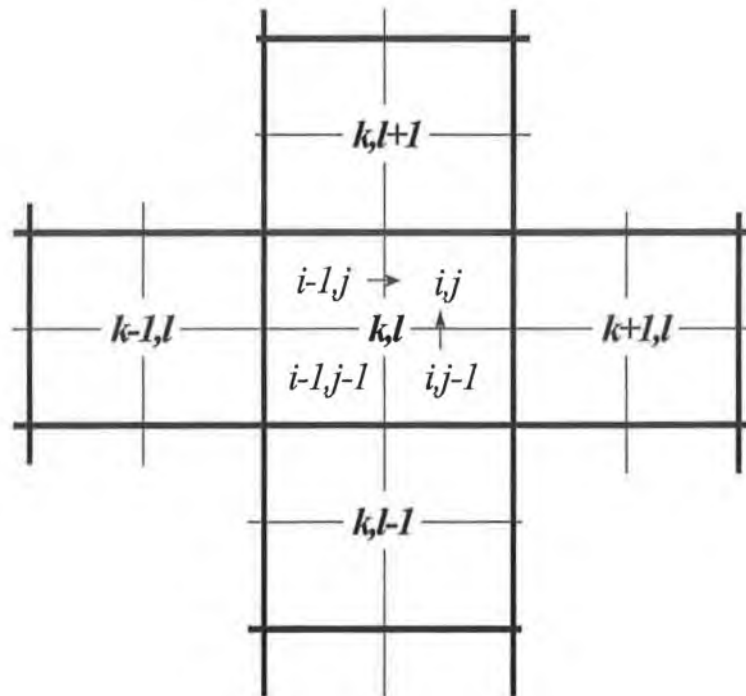
Figure 6.4. Construction of a coarse grid from 2x2 combinations of fine grid control volumes.

The figure shows that four adjacent fine grid control volumes are combined into one for the construction of the coarse grid. The i,j volume on the fine grid is shown with its associated, staggered variables. These are the velocities that share the i,j designation. The heavier lines show the coarse grid cells and k,l is adopted to denote coarse grid locations.

The simplest case to consider is the continuity equation. For this presentation the formalism presented by Hutchinson et al. (1988) is dispensed with as, while being concise, it doesn't appear to offer any practical advantage.

The continuity equations for the four fine grid cells in the figure are combined and then a substitution containing a correction is made.

The continuity equations for the fine grid cells are given by

$$i,j \text{ cell:} \quad 0 = A^c_{Ei,j}u_{i+1,j} + A^c_{Wi,j}u_{i,j} + A^c_{Ni,j}v_{i,j+1} + A^c_{Si,j}v_{i,j} \qquad (6.11)$$

$$i\text{-}1,j \text{ cell:} \quad 0 = A^c_{Ei-1,j}u_{i,j} + A^c_{Wi-1,j}u_{i-1,j} + A^c_{Ni-1,j}v_{i-1,j+1} + A^c_{Si-1,j}v_{i-1,j} \qquad (6.12)$$

$$i,j\text{-}1 \text{ cell:} \quad 0 = A^c_{Ei,j-1}u_{i+1,j-1} + A^c_{Wi,j-1}u_{i,j-1} + A^c_{Ni,j-1}v_{i,j} + A^c_{Si,j-1}v_{i,j-1} \qquad (6.13)$$

$$i\text{-}1,j\text{-}1 \text{ cell:} \quad 0 = A^c_{Ei-1,j-1}u_{i,j-1} + A^c_{Wi-1,j-1}u_{i-1,j-1} + A^c_{Ni-1,j-1}v_{i-1,j} + A^c_{Si-1,j-1}v_{i-1,j-1} \qquad (6.14)$$

where it is necessary to retain the cell identification for the coefficients. These four equations are added and then a substitution is made for each fine grid variable by expressing it as its current value added to a coarse grid correction, e.g. if the best available estimate of the fluid velocity on the fine grid is $\overline{u}$ then this velocity is to be corrected by a uniform adjustment within each block where

$$u_{ij} = \overline{u}_{ij} + u_{kl} \qquad (6.15)$$

where the corrections u_{kl} satisfy a continuity relation on the coarse grid. The variables u, v and p carrying the subscripts for each of the four cells indicated in figure 6.4 are corrected by u_{kl}, v_{kl} and p_{kl}. The value u_{i+1j}, for example, 'belongs' to the $k+1,l$ block in the diagram and it would be corrected by $u_{k+1,l}$. Thus the substitution for the sum of the continuity equations generates terms involving $k+1,l$ and $k,l+1$. Terms with common values of k and l can be gathered and the resultant equation takes the form

$$0 = A^c_{Ek,l}u_{k+1,l} + A^c_{Wk,l}u_{k,l} + A^c_{Nk,l}v_{k,l+1} + A^c_{Sk,l}v_{k,l} + B^c_{k,l} \qquad (6.16)$$

where

$$A^c_{Ek,l} = A^c_{Ei,j} + A^c_{Ei,j-1} \tag{6.17}$$

$$A^c_{Wk,l} = A^c_{Wi,j} + A^c_{Ei-1,j} + A^c_{Wi-1,j} + A^c_{Wi,j-1} + A^c_{Ei-1,j-1} + A^c_{Wi-1,j-1} \tag{6.18}$$

$$A^c_{Nk,l} = A^c_{Ni,j} + A^c_{Ni-1,j} \tag{6.19}$$

$$A^c_{Sk,l} = A^c_{Si,j} + A^c_{Ni,j-1} + A^c_{Si,j-1} + A^c_{Si-1,j} + A^c_{Ni-1,j-1} + A^c_{Si-1,j-1} \tag{6.20}$$

and the source term B_u is simply the sum of all the other terms involving the fine grid fluid velocities. Note that this is precisely the total mass residual given by the best current estimates of the fluid velocities on the fine grid.

The new equation, in terms of the coarse grid corrections, is of the same form as the continuity equation on the fine grid (where the source term is zero).

The same procedure is adopted for the fluid velocity equations. For the u velocity component, an equation on the coarse grid is constructed by summing the u equations for the four fine grid cells in the block. The coarse grid equation will have exactly the same form as the fine grid ones, i.e.

$$A^u_{Pk,l}u_{kl} = A^u_{Nk,l}u_{k,l+1} + A^u_{Sk,l}u_{k,l-1} + A^u_{Ek,l}u_{k+1,l} + A^u_{Wk,l}u_{k-1,l} + B^u_{k,l} - C^u_{k,l}(p_{k,l} - p_{k-1,l}) \tag{6.21}$$

where

$$A^u_{Pk,l} = A^u_{Pi,j} - A^u_{Ei-1,j} - A^u_{Ni,j-1} + A^u_{Pi-1,j} - A^u_{Wi,j} - A^u_{Ni-1,j-1} + A^u_{Pi,j-1} - A^u_{Ei-1,j-1} - A^u_{Si,j}$$
$$+ A^u_{Pi-1,j-1} - A^u_{Wi,j-1} - A^u_{Si-1,j} \quad , \tag{6.22}$$

$$A^u_{Ek,l} = A^u_{Ei,j} + A^u_{Ei,j-1} \quad , \tag{6.23}$$

$$A^u_{Wk,l} = A^u_{Wi-1,j} + A^u_{Wi-1,j-1} \quad , \tag{6.24}$$

$$A^u_{Nk,l} = A^u_{Ni,j} + A^u_{Ni-1,j} \quad , \tag{6.25}$$

$$A^u_{Sk,l} = A^u_{Si,j-1} + A^u_{Si-1,j-1} \quad , \tag{6.26}$$

$$C^u_{kl} = C^u_{i-1,j} + C^u_{i-1,j-1} \tag{6.27}$$

and the source term B_{kl} is the sum of the coefficients of the fine grid values for u and p and the fine grid source terms. It would be pedantic to show the full form of the coarse grid source term but note that A_P terms from the fine grid change sign when included in it.

The equations for v produce a coarse grid equation of the same form. The coarse grid set is identical in form to the fine grid set and they can be placed on a staggered coarse grid. Everything is therefore in place to apply the CELS solution technique on the coarse grid. The results of this calculation are the corrections that should be applied to the fine grid blocks in order to improve convergence at that level.

As in the general multigrid case, the procedure is then to solve the equations on the fine grid until some measure of progress shows that the efficiency of the solution is decreasing. At that point the equations are constructed for the coarse grid and activity is transferred to this. Corrections are calculated until some measure of the degree of convergence is met and then the fine grid values are corrected and the solution proceeds on the fine grid. The cycle is repeated whenever the fine grid efficiency falls.

It is easy to see, again since the coarse grid equations have the same form as the fine ones, that the procedure of transferring to a coarse grid can be nested. Blocks of coarse grid cells can be assembled and the equations for a still coarser grid can be constructed. The next block of 2×2 coarse grid cells would consist of 16 fine grid cells. The nesting process can be continued in this way until the entire domain is spanned by a small number of cells in a very coarse grid.

6.9   Coding of the ACM technique

Coding was provided as subroutines to be called by the CELS routines already implemented.

The coefficients for the coarse grid have to be calculated first. These are divided into the fixed coefficients and the source terms. The fixed coefficients are calculated once per coefficient update cycle and after that, whenever information is transferred to the coarse grid, only the source terms for the coarse grid equations have to be computed.

The solution of the coarse grid is carried out in a routine which is similar in structure to the fine grid one. On the fine grid level the fixed coefficients for the CELS equations are calculated once before beginning the iterations which update variables and source terms. The coarse grid routines are called repeatedly during the calculations for a fine grid coefficient update. Because the fixed coarse grid CELS coefficients do not vary within a coefficient update cycle these were removed to a subroutine which is called only once per coefficient update.

Thus, the minimum effort is expended on fixed coefficients and source terms only are updated on calling the routines for the coarse grid and while iterating on it.

Allowance has to be made for the inclusion of the continuity source terms that are not present in the solution routines for the fine grid. In practice, with the explicit use of the continuity equation on the fine grid, these source terms are small.

Chapter 7

Extension of the CELS and ACM methods to Three Dimensions

7.1 A Plane-Based Solver for CELS3D

The SIMPLE-based three-dimensional solver, developed in extending TEACH, formed the basis for the coding of CELS in three dimensions (CELS3D). The calculation of the coefficients was essentially the same as for the TEACH code. As in two dimensions, the pressure differences were removed from the source terms and handled explicitly because in the solution procedure the pressures are to be calculated and updated as soon as new values for the fluid velocity components are available.

It was found in the research for this thesis that the solver could be extended to three-dimensional problems by retaining the main features of the two-dimensional case. In two dimensions the solver acted on one plane. In three dimensions, the calculation domain is viewed as sets of planes upon which the solver can operate. To accomplish this the variables have to be paired whereas in two dimensions the only possibility was the plane defined by the u and v velocity components, there are now uv, vw and uw planes that can be defined. The solution operated in two dimensions in by calculating the variables in the orders uv and vu. In three dimensions there are 6 possible orders uv/vu, vw/wv and uw/wu.

The planes spanning a three-dimensional space are shown in figure 7.1. The positive velocity directions are indicated for each component and the component pair defining each plane is also shown.

7.2 The formulation of equations for CELS3D

The derivation of the penta-diagonal equations and the associated pressure and continuity relations could be carried out in a similar manner to the two-dimensional case presented in Chapter 5. The starting point was to take the variables in pairs and consider the solution across a plane with off-plane values regarded as constants. The equations for each plane could be moved into the same form as for the two-dimensional case by absorbing the terms involving off-plane values and their coefficients into modified source terms.
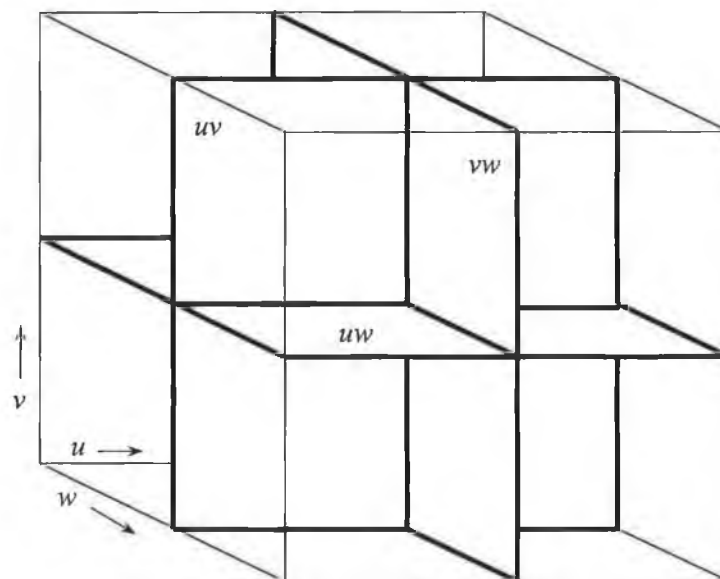
**Figure 7.1. The sets of planes spanning the calculation domain in
three dimensions.**

To demonstrate this, consider the three-dimensional forms of the continuity equation and the u, v and w velocity component equations.

For the control volume at position (i,j,k) in the solution domain

$$0 = A_N^c v_{ij+1k} + A_S^c v_{ijk} + A_E^c u_{i+1jk} + A_W^c u_{ijk} + A_F^c w_{ijk+1} + A_B^c w_{ijk} \tag{7.1}$$

$$A_P^u u_{ijk} = A_N^u u_{ij+1k} + A_S^u u_{ij-1k} + A_E^u u_{i+1jk} + A_W^u u_{i-1jk} + A_F^u u_{ijk+1} + A_B^u u_{ijk-1}$$
$$+ b^u - C^u(p_{ijk} - p_{i-1jk}) \tag{7.2}$$

$$A_P^v v_{ijk} = A_N^v v_{ij+1k} + A_S^v v_{ij-1k} + A_E^v v_{i+1jk} + A_W^v v_{i-1jk} + A_F^v v_{ijk+1} + A_B^v v_{ijk-1}$$
$$+ b^v - C^v(p_{ijk} - p_{ij-1k}) \tag{7.3}$$

$$A_P^w w_{ijk} = A_N^w w_{ij+1k} + A_S^w w_{ij-1k} + A_E^w w_{i+1jk} + A_W^w w_{i-1jk} + A_F^w w_{ijk+1} + A_B^w w_{ijk-1}$$
$$+ b^w - C^w(p_{ijk} - p_{ijk-1}) \tag{7.4}$$

where the superscript c denotes the continuity coefficients. In order to start the solution of these equations on a plane, one set of velocity components is held constant.

To show the relationship between the two-dimensional and the three-dimensional situation clearly, take the plane defined by the u and w velocity components and hold v constant. The resultant symmetry between the forms of the coefficients derived and the uv set from two-dimensions can then be clearly seen.

Thus, j is first held constant to establish a plane defined by u and w component directions. We can then derive the solver relations by considering a line of constant k along which w

142

will be calculated. The equations for u, w and continuity will then take the following form, with off-plane and off-line terms making up the modified source terms,

$$A_P^u u_{ij} = A_F^u u_{ijk+1} + A_B^u u_{ijk-1} - C^u p_{ijk} + B^u \qquad (7.5)$$

$$A_P^w w_{ijk} = A_F^w w_{ijk+1} + A_B^w w_{ijk-1} + C^w (p_{ijk-1} - p_{ijk}) + B^w \qquad (7.6)$$

$$0 = A_W^c u_{ijk} + A_F^c w_{ijk+1} + A_B^c w_{ijk} + B^c \qquad (7.7)$$

The solver can proceed, as before, by calculating a solution along a line of constant k, i.e. solving for w and producing values for u by the continuity equation.

The pressure relation is given by

$$A_P^p p_{ijk} = A_{FF}^p w_{ijk+2} + A_F^p w_{ijk+1} + A_B^p w_{ijk} + A_{BB}^p w_{ijk-1} + B^p \qquad (7.8)$$

where the coefficients without subscripts refer to the ijk cell and are given by

$$A_P^p = -C^u, \qquad (7.9)$$

$$A_{FF}^p = \frac{A_F^u A_{Fijk+1}^c}{A_{Wijk+1}^c}, \qquad (7.10)$$

$$A_F^p = \frac{A_F^u A_{Bijk+1}^c}{A_{Wijk+1}^c} - \frac{A_P^u A_F^c}{A_W^c}, \qquad (7.11)$$

143

$$A_B^p = \frac{A_B^u A_{Fijk-1}^c}{A_{Wijk-1}^c} - \frac{A_P^u A_B^c}{A_W^c} \quad, \tag{7.12}$$

$$A_{BB}^p = \frac{A_B^u A_{Bijk-1}^c}{A_{Wijk-1}^c} \tag{7.13}$$

and

$$B^p = \frac{A_F^u}{A_{Wijk+1}^c} B_{ijk+1}^c - \frac{A_P^u}{A_W^c} B^c + \frac{A_B^u}{A_{Wijk-1}^c} B_{ijk-1}^c - B^u \quad. \tag{7.14}$$

The penta-diagonal equation for w is given by

$$a_P^w w_{ijk} = a_{FF}^w w_{ijk+2} + a_F^w w_{ijk+1} + a_B^w w_{ijk-1} + a_{BB}^w w_{ijk-2} + b^{*w} \tag{7.15}$$

where

$$a_P^w = A_P^w + C^w \frac{A_B^p}{A_P^p} - C^w \frac{A_{Fi1jk-1}^p}{A_{Pijk-1}^p} \tag{7.16}$$

$$a_{FF}^w = -C^w \frac{A_{FF}^p}{A_P^p} \tag{7.17}$$

$$a_E^w = A_F^w + C^w \frac{A_{FFiik-1}^p}{A_{Pijk-1}^p} - C^w \frac{A_F^p}{A_P^p} \tag{7.18}$$

$$a_B^w = A_B^w + C^w \frac{A_{Bijk-1}^p}{A_{Pijk-1}^p} - C^w \frac{A_{BB}^p}{A_P^p} \tag{7.19}$$

144

$$a_{BB}^{w} = -C^{w} \frac{A_{BBijk-1}^{p}}{A_{Pijk-1}^{P}} \qquad (7.20)$$

$$b^{*w} = B^{w} - C^{w} \frac{B^{p}}{A_{P}^{p}} + C^{w} \frac{B_{ijk-1}^{p}}{A_{Pijk-1}^{p}} \qquad (7.21)$$

The same type of derivation is applied to each of the planes to derive a complete set of coefficients that can be used to calculate all the variables in the three-dimensional space by sweeping the appropriate line solvers through all the plane sets.

The last line of the plane sweeps requires a mass conservation approach, as used in two-dimensions (see Section 5.3) and the same type of last line pressure calculation.

7.3   Extension of ACM to Three Dimensions (ACM3D)

The extension of the Additive Correction Multigrid for the fluid flow equations was tackled in as straightforward a manner as possible. As in two dimensions, a coarse grid cell was constructed from a combination of fine grid cells. Two cells in each grid direction from a given point were included in the production of a single coarse grid cell and its equations. Thus the total number of fine grid cells per coarse grid cell is 8 $(2 \times 2 \times 2)$.

The same labelling system for the fine grid and corresponding coarse grid cell as was applied in the two-dimensional case (see Section 6.8). The arrangement of cells and their coarse grid counterparts is shown in figure 7.2

The groups of 8 cells are in a square arrangement with grid point indices ranging from i, j,k to i-1, j-1, k-1. The coarse grid cells are designated by l,m and n as indices. The group

145

of fine grid cells represented by the root cell i,j,k is mapped into the coarser grid cell l,m,n. The labelling system for coarse grid neighbours is shown in figure 7.2.
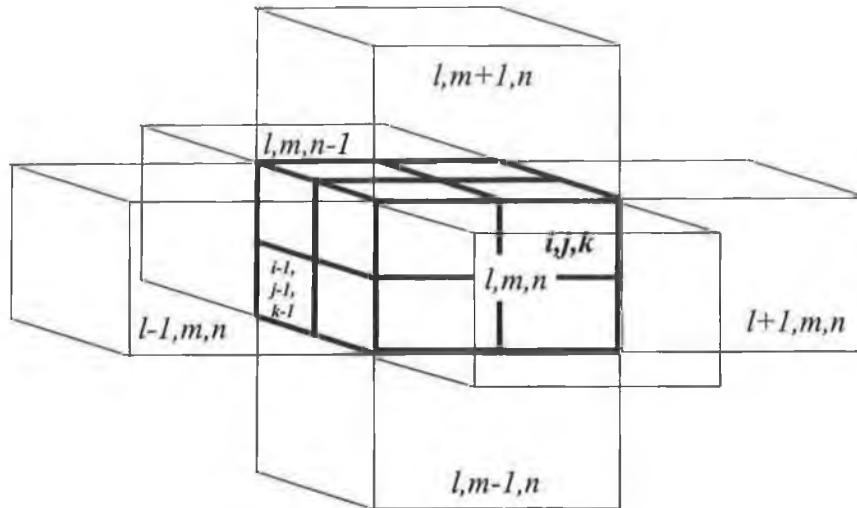
Figure 7.2. Arrangement of fine grid cells used to construct coarse grid equations.

The construction of the three-dimensional coarse grid equations starts, as in two dimensions, from the combination of the equations for the fine grid cells.

The three-dimensional continuity equation runs through the three indices from the

i,j,k cell:

$$0 = A^c_{Ei,j,k}u_{i+1,j,k} + A^c_{Wi,j,k}u_{i,j,k} + A^c_{Ni,j,k}v_{i,j+1,k} + A^c_{Si,j,k}v_{i,j,k}$$
$$+ A^c_{Fi,j,k}w_{i,j,k+1} + A^c_{Bi,j,k}w_{i,j,k} \tag{7.22}$$

to the i-1,j-1,k-1 cell:

$$0 = A^c_{Ei-,j-1,k-1}u_{i,j-1,k-1} + A^c_{Wi-1,j-1,k-1}u_{i-1,j-1,k-1} + A^c_{Ni-1,j-1,k-1}v_{i-1,j,k-1}$$

$$+ A^c_{Si-1,j-1,k-1}v_{i-1,j-1,k-1} \tag{7.23}$$

These eight equations are added together and substitutions are made for each of the variables (in this case velocity components only). As before, the substitution consists of the latest estimate for that variable plus a correction. For, say, the u velocity component, this would be of the form

$$u_{ijk} = \overline{u}_{ijk} + u_{lmn} \tag{7.24}$$

where u_{lmn} is the block correction derived by solving on the coarse grid and $\overline{u}$ is the best available u estimate from the fine grid. The u_{lmn} correction is applied to all the u values within the group of eight fine cells, i.e. those that share the same indices as the grid points from (i,j,k) to (i-1,j-1,k-1). Velocities outside these cells are corrected by the appropriate block correction from an adjacent coarse grid cell. Similar corrections are applied to the other variables.

When terms with common values of coarse grid variables are collected, a coarse grid continuity equation results. This is

$$0 = A^c_{El,m,n}u_{l+1,m,n} + A^c_{Wl,m,n}u_{l,m,n} + A^c_{Nl,m,n}v_{l,m+1,n} + A^c_{Sl.m.n}v_{l,m,n}$$

$$+ A^c_{Fl,m,n}v_{l,m,n+1} + A^c_{Bl,m,n}v_{l,m,n} + B^c_{l,m,n} \tag{7.25}$$

where

$$A^c_{El,m,n} = A^c_{Ei,j,k} + A^c_{Ei.j-1,k} + A^c_{Ei,j,k-1} + A^c_{Ei.j-1,k-1} , \tag{7.26}$$

147

$$A^c_{WI,m,n} = A^c_{Wi,j,k} + A^c_{Ei-1,j,k} + A^c_{Wi-1,j,k} + A^c_{Wi,j-1,k} + A^c_{Ei-1,j-1,k} + A^c_{Wi-1,j-1,k} + A^c_{Wi,j,k-1}$$
$$+ A^c_{Ei-1,j,k-1} + A^c_{Wi-1,j,k-1} + A^c_{Wi,j-1,k-1} + A^c_{Ei-1,j-1,k-1} + A^c_{Wi-1,j-1,k-1} \; , \tag{7.27}$$

$$A^c_{NI,m,n} = A^c_{Ni,j,k} + A^c_{Ni-1,j,k} + A^c_{Ni,j,k-1} + A^c_{Ni-1,j,k-1} \tag{7.28}$$

$$A^c_{SI,m,n} = A^c_{Si,i,k} + A^c_{Si-1,j,k} + A^c_{Ni,j-1,k} + A^c_{Si,j-1,k} + A^c_{Ni-1,j-1,k} + A^c_{Si-1,j-1,k} + A^c_{Si,j,k-1}$$
$$+ A^c_{Si-1,j,k-1} + A^c_{Ni,j-1,k-1} + A^c_{Si,j-1,k-1} + A^c_{Ni-1,j-1,k-1} + A^c_{Si-1,j-1,k-1} \tag{7.29}$$

$$A^c_{FI,m,n} = A^c_{Fi,j,k} + A^c_{Fi-1,j,k} + A^c_{Fi,j-1,k} + A^c_{Fi-1,j-1,k} \tag{7.30}$$

$$A^c_{BI,m,n} = A^c_{Bi,j,k} + A^c_{Bi-1,j,k} + A^c_{Bi,j-1,k} + A^c_{Bi-1,j-1,k} + A^c_{Fi,j,k-1} + A^c_{Bi,j,k-1} + A^c_{Fi-1,j,k-1}$$
$$+ A^c_{Bi-1,j,k-1} + A^c_{Fi,j-1,k-1} + A^c_{Bi,j-1,k-1} + A^c_{Fi-1,j-1,k-1} + A^c_{Bi-1,j-1,k-1} \tag{7.31}$$

It should be noted that the coefficients of the block corrections from adjacent coarse grid points consist of the sum of the coefficients over the fine grid interface. The source term S_{lmn} is the sum of all the terms that occurred in the fine grid equations and, as before, this is seen to be the mass residual for the block defined by the eight fine grid cells.

The coarse grid equations for the u velocity component are also constructed following the same procedure as in two dimensions. The coarse grid u equation is given by

$$A^u_{PI,m,n} u_{l,m,n} = A^u_{NI,m,n} u_{l,m+1,n} + A^u_{SI,m,n} u_{l,m-1,n} + A^u_{EI,m,n} u_{l+1,m,n} + A^u_{WI,m,n} u_{l-1,m,n}$$
$$+ A^u_{FI,m,n} u_{l,m,n+1} + A^u_{BI,m,n} u_{l,m,n-1} + B^u_{l,m,n} - C^u_{l,m,n}(p_{l,m,n} - p_{l-1,m,n}) \tag{7.32}$$

where

$$A^u_{Pl,m,n} = A^u_{Pi,j,k} - A^u_{Si,j,k} - A^u_{Wi,j,k} - A^u_{Bi,j,k} + A^u_{Pi-1,j,k} - A^u_{Si-1,j,k} - A^u_{Ei-1,j,k} - A^u_{Bi-1,j,k}$$

$$+ A^u_{Pi,j-1,k} - A^u_{Ni,j-1,k} - A^u_{Wi,j-1,k} - A^u_{Bi,j-1,k} + A^u_{Pi-1,j-1,k} - A^u_{Ni-1,j-1,k} - A^u_{Ei-1,j-1,k} - A^u_{Bi-1,j-1,k}$$

$$+ A^u_{Pi,j,k-1} - A^u_{Si,j,k-1} - A^u_{Wi,j,k-1} - A^u_{Fi,j,k-1} + A^u_{Pi-1,j,k-1} - A^u_{Si-1,j,k-1} - A^u_{Ei-1,j,k-1} - A^u_{Fi-1,j,k-1}$$

$$+ A^u_{Pi,j-1,k-1} - A^u_{Ni,j-1,k-1} - A^u_{Wi,j-1,k-1} - A^u_{Fi,j-1,k-1} + A^u_{Pi-1,j-1,k-1} - A^u_{Ni-1,j-1,k-1}$$

$$+ A^u_{Ei-1,j-1,k-1} - A^u_{Fi-1,j-1,k-1} \quad , \tag{7.33}$$

$$A^u_{El,m,n} = A^u_{Ei,j,k} - A^u_{Ei,j-1,k} - A^u_{Ei,j,k-1} - A^u_{Ei,j-1,k-1} \tag{7.34}$$

$$A^u_{Wl,m,n} = A^u_{Wi-1,j,k} - A^u_{Wi-1,j-1,k} - A^u_{Wi-1,j,k-1} - A^u_{Wi-1,j-1,k-1} \tag{7.35}$$

$$A^u_{Nl,m,n} = A^u_{Ni,j,k} - A^u_{Ni-1,j,k} - A^u_{Ni,j,k-1} - A^u_{Ni-1,j,k-1} \tag{7.36}$$

$$A^u_{Sl,m,n} = A^u_{Si,j-1,k} - A^u_{Si-1,j-1,k} - A^u_{Si,j-1,k-1} - A^u_{Si-1,j-1,k-1} \tag{7.37}$$

$$A^u_{Fl,m,n} = A^u_{Fi,j,k} - A^u_{Fi-1,j,k} - A^u_{Fi,j-1,k} - A^u_{Fi-1,j-1,k} \tag{7.38}$$

$$A^u_{Bl,m,n} = A^u_{Bi,j,k-1} - A^u_{Bi-1,j,k-1} - A^u_{[Bi,j-1,k-1} - A^u_{Bi-1,j-1,k-1} \tag{7.39}$$

$$C^u_{l,m,n} = C^u_{i-1,j,k} - C^u_{i-1,j-1,k} - C^u_{i-1,j,k-1} - C^u_{i-1,j-1,k-1} \tag{7.40}$$

The source term for the coarse grid u equation is the sum of all the terms in the eight fine grid equations but it should be noted that the fine grid central coefficient A_p terms change sign when they are included in the coarse grid source terms.

Coarse grid equations for the v and w velocities are constructed in exactly the same manner. As in two dimensions, since the coarse grid equations have the same form as the

149

fine grid ones, the CELS3D solver can be applied at this level. The only modification required is the inclusion of the source term for the continuity equation in the coding of the CELS3D sources.

The same remarks about the coding of the solver apply as in the two-dimensional case (see Section 6.9).

Chapter 8

Use of TEACH Code

8.1 The application of the TEACH code in two dimensions

A program for the calculation of the steady, two-dimensional, turbulent, re-circulating fluid flows, TEACH-2E (Gosman and Ideriah, 1976), was obtained. The program uses Cartesian co-ordinates and a rectangular grid. TEACH was first modified to set up a square geometry with inlet and outlet in much the same manner as its application for the calculation of flows in animal housing (Choi, 1988 and 1990). This could be used to provide a two-dimensional approximation to the square section growing rooms in the Netherlands and in North America.

The first step towards its application for the modelling of the geometry of an Irish mushroom tunnel was to approximate the arc of the curved wall as a boundary within the solution domain. This was done by using the technique of modifying the linearised source term as described in Section 4.5. It was therefore possible to set air velocities outside the shape of the house to zero. Zero fluxes across the wall could also be achieved by setting the convection/diffusion coefficient for the appropriate interface to zero. The set of grid points defining the wall was set up in the program and the effective wall used in the first calculations is shown in figure 8.1. This is on a 22×22 grid (the maximum grid dimension that was set in the array definitions in the version of TEACH supplied) and is a coarse approximation but Pantankar (1980) notes that, depending on what is required from the model, surprisingly good results can be obtained with such approximations.
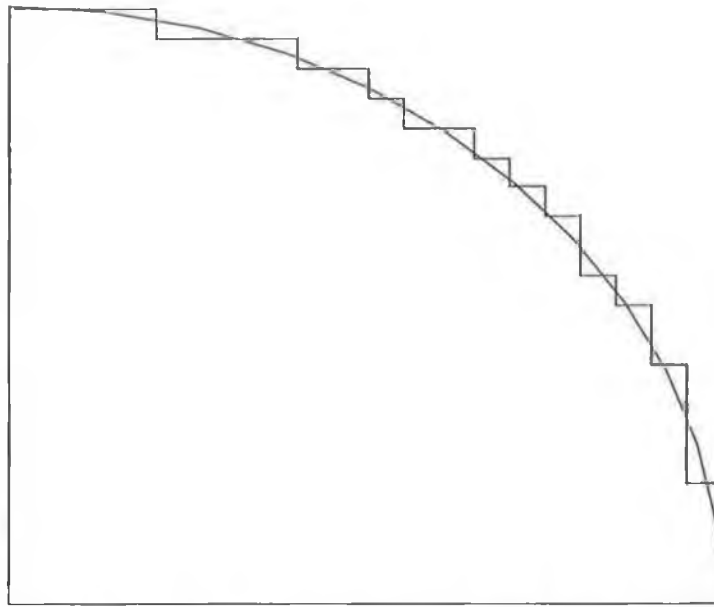
Figure 8.1. 22×22 grid representation of the tunnel wall.

As mentioned above, the growing rooms generally show bilateral symmetry but within a cross-section there is no exit for the air flow which is provided at the inlet. The centre of the house is actually an axis of symmetry so that an idealised model should allow no flow across it. In two dimensions it was necessary to set up an outlet to satisfy continuity. This approximation served as a useful testing ground for the ability of the code to meet the main requirements for the modelling. Boundary conditions were as set out in Section 4.10 and at this stage the hybrid differencing scheme (supplied in TEACH) was employed for the calculations. The dimensions of the calculation domain were set at 3 m square. For the calculations, the normalised residual target (Gosman and Ideriah, 1976) for the discretised equations was set at 0.005 (as supplied). Experiments with different targets showed that gains in accuracy beyond this target were minimal for the air flows calculated.

8.2   Preliminary flow calculations

Despite the coarseness of the approximation, the calculations on the 22×22 grid showed that varying the selection of points describing the wall (where inclusion or exclusion was marginal) produced no more than a 3% variation in the air speeds across the floor. A typical flow pattern is shown in figure 8.2 and the outlet is located where the flow markers cross the boundary. In this case, the Reynolds number was 5000 and the inlet velocity was 0.5 ms^{-1}. The inlet was set on the line $j = 18$ and the outlet from $j = 5$ to $j = 8$. Choi et al. (1988) reported that they could not achieve convergence at this Reynolds number but no difficulty was encountered in this instance.
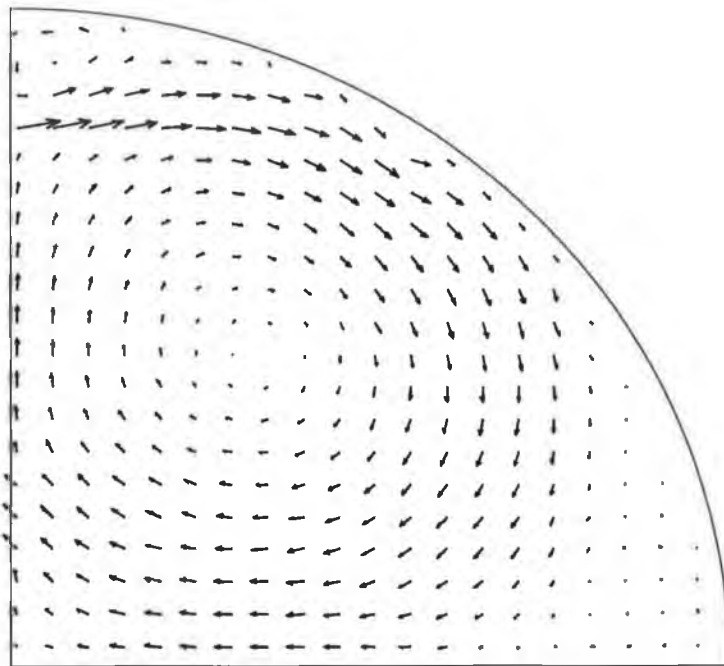
Figure 8.2. Recirculating flow on a 22×22 grid (inlet velocity = 0.5ms^{-1}).

The effect of the coarse wall representation was clearly seen on the 22×22 grid. There were obvious deflections due to the corners in the wall approximation. A point of particular concern was the first step down at the top of figure 8.1. While a small circulation zone is expected above the inlet, the corner situated on the wall shape just a short distance away is having a very clear effect on the flow pattern.

After first modifying the TEACH code (the array definitions in the COMMON blocks) to allow greater grid densities, further calculation was carried out for a 34×34 grid. This corresponds to the maximum grid density that could be achieved later in three-dimensional calculations and is also convenient for the telescoping of grids that is to be done in the application of ACM. For this flow, the Reynolds number was again set to be 5000 by setting the inlet air velocity was 0.9 ms^{-1}. The dimensions were set as 3 m square and the inlet was on the line $j = 30$ and outlet was from the line $j = 8$ to $j = 11$. The flow in figure 8.3 shows, while deflections are still present, that the effect on the flow at a given distance from the wall is reduced. It is believed that operation on this grid or higher densities could be useful in terms of studying the general flow around obstacles in the body of the flow. TEACH was not supplied with the code for body-fitted co-ordinates but within this limitation the description of the topmost corner could be refined by adding extra horizontal lines of grid points to refine the grid in that area. This sharp corner and any other particularly troublesome corners could thus be broken down into a set of smaller ones.
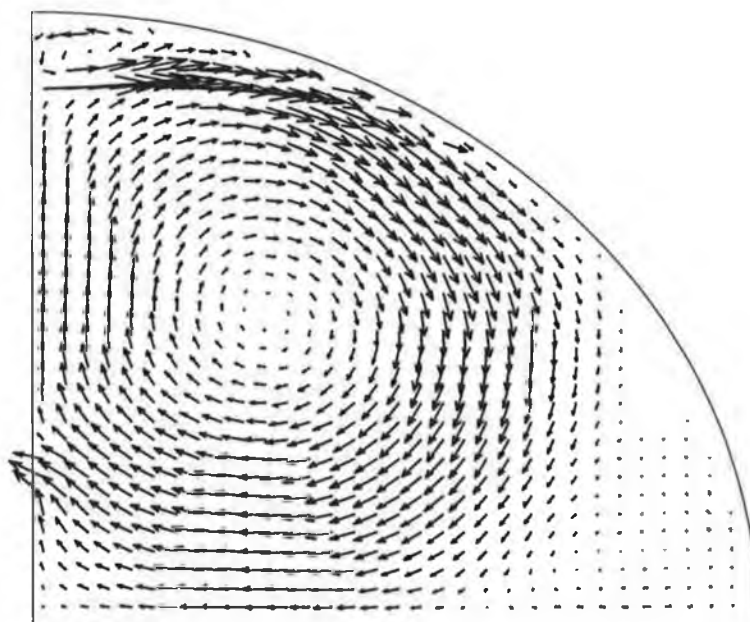
Figure 8.3. Re-circulating flow on a 34×34 grid (inlet velocity = 0.9ms^{-1}).

At this stage, a final test was performed. The most complex flows in mushroom growing are those around the multi-tier shelf or tray systems. Using SIMPLE, the inclusion of shelving in the flow was accomplished straightforwardly by the method of modified source terms. Using the same flow configuration as that for figure 8.3, grid points internal to shelves had air velocity components set to zero. The shelves were set 2 grid units high and 9 units long, with bottom left hand corners located at $i = 7$ and $j = 6$, $j = 12$ and $j = 18$. The modified TEACH code converged and a result for three shelves is shown in figure 8.4. This result is qualitatively similar to the flows observed in practice. The highest speeds are recorded on the lowest shelf while those on the upper ones are reduced and of similar magnitude.
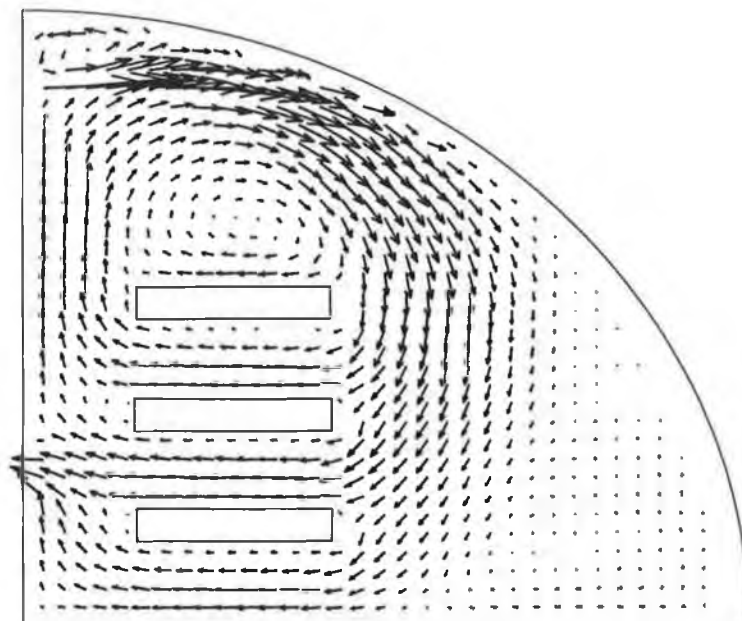
Figure 8.4. Re-circulating flow around shelves on a 34×34 grid (inlet velocity = 0.9ms⁻¹).

8.3　Extension of the TEACH code to three dimensions

Having verified that there was no difficulty in principle in calculating the main features required for modelling purposes, the code had to be extended to three dimensions in order to model the air flow along as well as across the growing rooms.

In the two-dimensional approximation a 'false' outlet was required to allow satisfaction of the continuity relations built into the pressure correction equation. Continuity in reality is maintained by a flow along the tunnel or growing room, i.e. into and out of the plane already modelled in two dimensions. Air is delivered at the top of a tunnel and moves back along its length to be re-circulated or exhausted at one of the end walls.

There are numerous studies in three dimensions using the same general procedures as TEACH. Early work by Patankar and Spalding (1972) laid down principles for the solution of the three-dimensional equations that developed into the general techniques used in TEACH. It described a generally applicable, accurate and economic method for heat, mass and momentum transfer in three dimensions and it used the staggered grid. Patankar et al. (1977) used this technique to study the deflection of a turbulent jet in three dimensions.

Extension of the code involved adding two extra faces to the existing square-geometry, two-dimensional finite volume cell making it a cuboid. The new sides added to the north, south east and west set were labelled front (F) and back (B) following a left handed convention with distance or displacement in the z direction (the third dimension) increasing positively from back to front.

The air velocity component in the third dimension was labelled z by convention. The equations for the three-dimensional case are a direct extension of those used in TEACH for the two-dimensional model. Following Patankar (1980), with F and B as variable indices for the new terms, the general discretised equations can be stated as

$$a_P\phi_P = a_E\phi_E + a_W\phi_W + a_N\phi_N + a_S\phi_S + a_F\phi_F + a_B\phi_B + b \qquad (8.1)$$

where ϕ is a general variable denoting those to be solved and now including the velocity component w. The expressions for the all coefficients as used in this work are stated below for completeness.

$$a_E = D_e A(|P_e|) + [[-F_e, 0]] \qquad (8.2)$$

$$a_W = D_w A(|P_w|) + [[F_w, 0]] \qquad (8.3)$$

$$a_N = D_n A(|P_n|) + [[-F_n, 0]] \qquad (8.4)$$

$$a_S = D_s A(|P_s|) + [[F_s, 0]] \qquad (8.5)$$

$$a_F = D_f A(|P_f|) + [[F_f, 0]] \qquad (8.6)$$

$$a_B = D_b A(|P_b|) + [[-F_b, 0]] \qquad (8.7)$$

$$b = S_C \Delta x \Delta y \Delta z \qquad (8.8)$$

$$a_P = a_E + a_W + a_N + a_S + a_F + a_B - S_p \Delta x \Delta y \Delta z \qquad (8.9)$$

Where the Péclet number P is taken as the ratio of F and D (flow and conductance) expressions, e.g. $P_e = F_e / D_e$. The symbol [[ ]] stands for the largest of the quantities within it and function $A(|P|)$ represents the hybrid scheme [[0, 1-0.5|P|]].

For each cell boundary the appropriate expressions for flow rates and conductances are as follows:

where $\Delta x, \Delta y, \Delta z$ are the dimensions of the finite volume and $\delta x, \delta y, \delta z$ are the distances between the grid points involved.

In the coding of the two-dimensional case there was an implied third dimension in the finite volume cells and these were of unit depth. The introduction of a non-unit depth means that account of it must be taken when calculating areas in existing terms for the east, west, north and south walls of the calculation volumes. The expressions for the new F and B faces were added.

No mathematical difficulty was encountered in the extension of the TEACH code to three dimensions. The generation term for the turbulence (see Section 3.4) was expanded to its three-dimensional form but otherwise all the expressions varied little in form. After a period of debugging of the code and the testing, convergence to physically realistic flows was achieved.

There was one change made to the operating procedure to that used in two dimensions. When all the equations were under-relaxed using the relaxation factors as supplied by TEACH, it was observed that the mass/continuity residuals decreased very much more rapidly than did those for the air velocity components. This had the effect of slowing the rate of convergence because there was little change in the air velocity due the adjustments generated from the pressure correction equation. Under-relaxation of the pressure correction equation was set to be 0.8 in order to correct this by slowing down the changes in these corrections. This value was found to significantly improve the rate of convergence.

A cross-section of a three-dimensional flow is shown in figure 8.5. It should be noted that the west wall that was used in two dimensions has been replaced by an axis of symmetry where the air velocity components just outside the solution domain are set to the same values as those just inside the boundary. This effectively sets up a model of the bilateral symmetry that is observed in actual flows. The dimensions of the solution domain were set to 3 m square (approximating the small tunnels used for research in growing techniques) and the Reynolds number was 5000. The inlet in this diagram is located at $i = 28$.
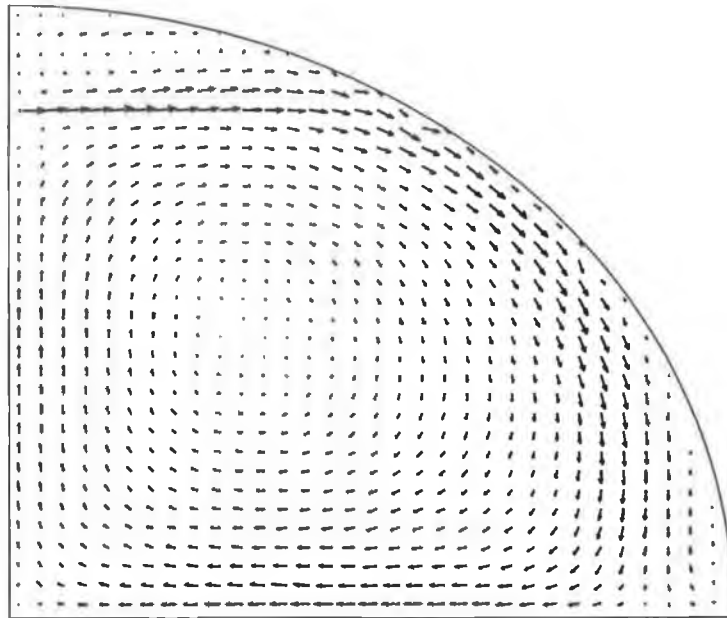
**Figure 8.5. Cross-section of a three-dimensional
velocity field (34×34×34 grid) with inlet velocity 0.2ms^{-1}.**

At this stage excessively long convergence times were making it difficult to test and further develop the model and therefore it was decided to seek an alternative solution procedure. Experiments with the Correction Storage multigrid scheme were carried out in two dimensions in order to determine its potential. The technique was applied to the inner iterations on the current set of linear equations, i.e. between the coefficient updates. The maximum overall saving achieved was 15%. Only one coarse grid was used because no extra savings were made by moving the problem to coarser grids. The limiting factor was the number of coefficient updates required by the SIMPLE technique. The benefit of improving the interim linear solutions was very limited because the effort required per update is small relative to the total effort required for convergence.

It was judged that there was little point in extending the CS scheme to three dimensions. It is reported that other researchers have found no significant savings when attempting to improve the performance of the solver by similar approaches (Phillips and Schmidt, 1985 and Miller and Schmidt, 1988).

The structure of the SIMPLE procedure has come to be called a segregated or sequential solver as opposed to a simultaneous type of solver because interim solutions are produced for each variable in turn before moving on to solve the next. The result of holding some variables at fixed values while another is solved is to place a limit on the degree of convergence required for each set of linearised equations generated during coefficient updates. With imperfect solutions for the other variables there is no point is solving any given equation to a high degree of convergence because the solution is limited by the accuracy of the other variables. In practice, no more than two or three sweeps of the solution domain with the tri-diagonal solver were required.

The other limiting factor that applies to the available segregated solvers is in the updating of the variables with the correction equations. The corrections, arising from an interim solution of the pressure correction equation, are applied to the solution of the air velocity equations and produce an interplay between the equations that slows down the overall rate of convergence. This is the pressure-velocity coupling that has been dealt with in a number of publications (e.g. Van Doormal and Raithby, 1974 and Galpin et al., 1985). It was observed during this work that it can also happen that the pressure correction equation converges more quickly than the air velocity equations. This then moves through a sequence of near convergent solutions that are prescribed by the velocity solutions. The corrections become small and no longer strongly influence the velocity solutions and this also acts to slow down the rate of convergence.

During experimentation, it was found that the convergence could be made more efficient by re-tuning the relaxation factors to produce a slowing down in the rate of convergence of the pressure-correction equation. However, this improvement was

very limited due to the segregated nature of the solution procedure and the inter-play between the pressure and the air velocity.

Chapter 9

CELS in Two Dimensions

9.1 Application of CELS in Two Dimensions

It was not clear to what extent the coupling between the velocity/pressure equations and those for the turbulence would influence the efficiency of the CELS method. In order to justify the extension of CELS to three dimensions, which presents its own difficulties, it was decided to experiment in two dimensions. Also, the use of boundaries internal to the solution domain could present a difficulty in the standard CELS approach and this had to be investigated.

Savings in computational effort in the ACM study (Hutchinson et al., 1988) were presented for relatively dense grids (64×64 and 96×96) and, as the performance of SIMPLE becomes markedly poorer as the grid density increases, it was not clear how beneficial the ACM method would be on the coarser grids which would probably suffice for the requirements of the modelling of mushroom growing rooms.

A simple square flow domain (figure 9.1) was set up as the basis for the first comparison of the relative efficiencies of the CELS and the SIMPLE methods with the turbulence model included in the equation set.
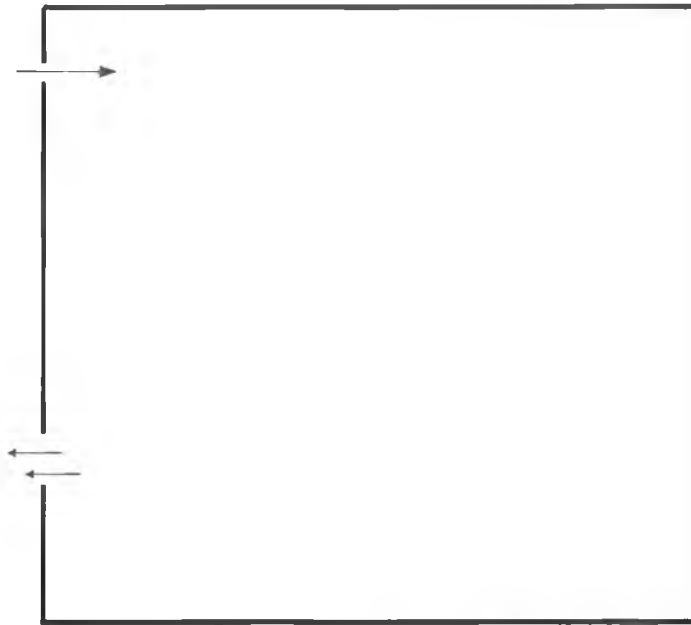
Figure 9.1. Flow geometry for the initial comparison of CELS and SIMPLE for turbulent flow.

The flow enters from the left side near the top and an exit is provided on the same side nearer the base of the room. This provides the main circulation features for the flow in a mushroom building in a two-dimensional approximation.

The first calculations were made on a 34×34 grid with flow entry one grid cell wide at $j = 30$ and the bottom of the outlet at $j = 8$ and extending upward for 3 cells. The issue of resolution of the flow at the entrance and exit would need careful examination when validating a model of a real flow but, in view of the good agreement presented by Choi (1988) in a square geometry with similar entry and exit descriptions, this aspect was not examined for the purposes of testing the solver. The boundary conditions were set in the same manner as outlined in Section 4.10. The entry air velocity was set at 1 ms^{-1} and the exit air velocity component was adjusted to balance the inflow. The dimensions of the flow domain were 3 m square and the resultant Reynolds number was 5500.

Before making the comparison between SIMPLE and CELS for turbulent flows it was necessary to examine the general behaviour of CELS and to optimise its operation. There are a number of factors that influence the convergence time of CELS.

(i)   As with SIMPLE there is the possibility of under-relaxation of the flow equations themselves.

(ii)   A further equation under-relaxation can be introduced during the solution of CELS (referred to later as CELS relaxation). This was found to be necessary by Galpin et al. (1985) to enhance the performance of the solver.

(iii)   The degree of convergence per coefficient update, α, has to be set. This is the factor by which equation residuals are reduced before the next coefficient update.

The effects of each of these on the time to convergence and the number of coefficient updates required were examined.

The code for CELS was constructed and debugged and a convergent solver was achieved. The coding was based on TEACH with essentially the same coefficient calculations. As before, the pressure difference terms in the velocity equations were not included in the source terms but were handled explicitly in the new solver.

New CELS coefficient calculations and solution routines replaced the tri-diagonal solver from TEACH. Hybrid differencing was used at this point in order to minimise complications and possible bugs in the coding.

There was one initial difference from the original implementation to be noted. A pressure correction, proposed by Galpin et al. (1985), was said to have a beneficial effect on the rate of convergence. They recommended that a block correction be applied based on the constraint that, for the j-line case (line solution of u), that the v equation be satisfied on average along the line. The residuals for the equation were used to construct a pressure correction as

$$\delta p_j = \left[\sum_i (A_p^v v_{ij} - \sum_{nb} A_{nb}^v v_{nb} + C^v(p_{ij+1} - p_{ij}) - b^v) \right] / \sum_i C^v \qquad (9.1)$$

and δp was added to each pressure on the line. This was applied beginning at the second last line of the sweep and then applied to each line in turn in the opposite direction to the sweep just completed.

This was not found to produce any benefit in this application. It was difficult to account for this but, as noted above, there are differences between this application and the original in the use of hybrid differencing rather than the upstream-weighted differencing scheme of Raithby and Torrance (1974).

It was also noted that, in the early stages of a given iteration, the values of the velocity components were not bounded in the locality of the flow entry, i.e. they could temporarily take on values greater than the input velocity. This occurred as the line solver passed the flow entry point. Continuity can be satisfied locally without boundedness. These values disappeared after a few iterations. The geometries used by Galpin et al. (1985) may have avoided this type of effect by having the flow entry on, or partially on, the first line of the solution and by having the exit conditions finally resolve the continuity on the last line. This may also have played a part in obtaining savings with the pressure correction procedure.

An arrangement was set up where the flow entry was shifted from the side to the top of the flow domain and thus the flow in began on the first line to be solved. It was found that boundedness improved but the pressure correction was still found to be ineffective.

The lack of boundedness may create a requirement for a minimum number of inner iterations per coefficient update. If coefficients are created using the unbounded values then divergence could result. In practice, this minimum was never defined because the degree of convergence per coefficient update that was required for an efficient operation of CELS avoided this problem. Also, the effect seemed to be

problem-dependent and, in the cases examined during the course of the experimentation, it simply acted as a factor making for poor convergence with a low number of inner iterations.

The first coefficient update was found to require a fixed number of conditioning iterations in order to bring the variables to values that would allow reasonable coefficients to be calculated, the initial guess being zero for the velocities. A value of 5 was taken as a standard for the following tests. Higher values could reduce the total number of coefficient updates but extra effort initially was not balanced by later savings. A smaller number left the fields insufficiently solved and resulted in more subsequent coefficient updates and increases in solution time.

9.2   General Characteristics of CELS

The general behaviour of CELS was first defined for the laminar case, i.e. the turbulence equations were not solved. Although CELS was originally constructed and tested for laminar flows, this was done because of the differences noted above between this application and the work done by Galpin et al. (1985). The factors to be tested were first brought by experimentation to near-optimum values and then the behaviour for the variation around these was investigated.

With α set to 0.25 and a CELS relaxation factor of 0.7, the effect of the relaxation factor on the solution of the flow equations was tested. Figures 9.2 and 9.3 show the results of solving the set of equations for a number of values of the relaxation factor applied to the velocity component equations.
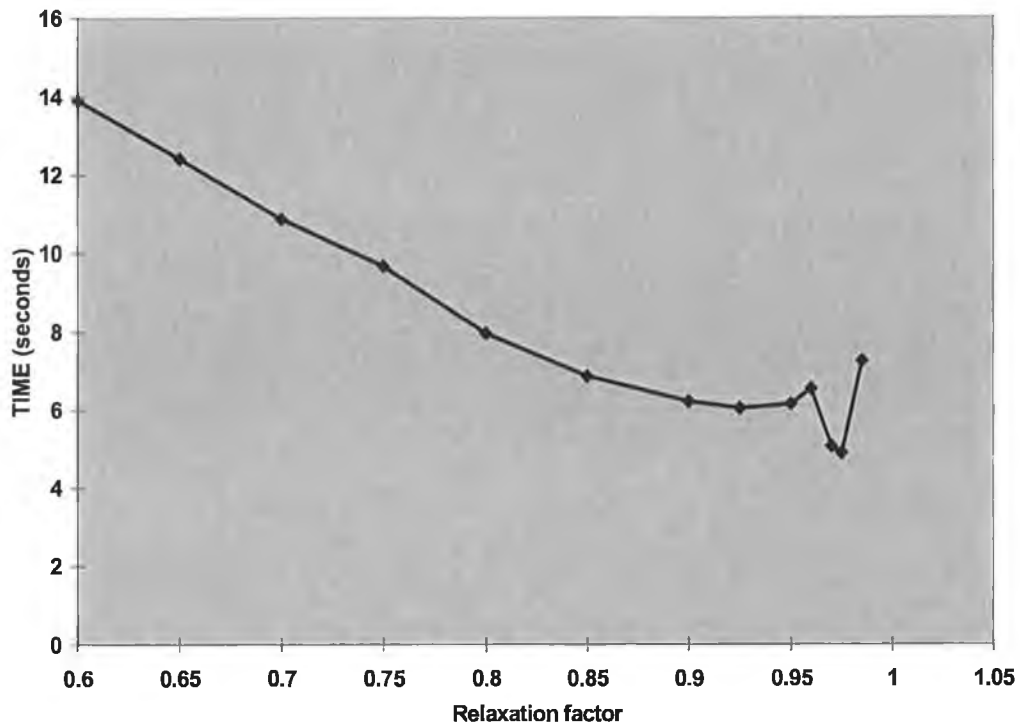
Figure 9.2. Variation of convergence time with relaxation factor.
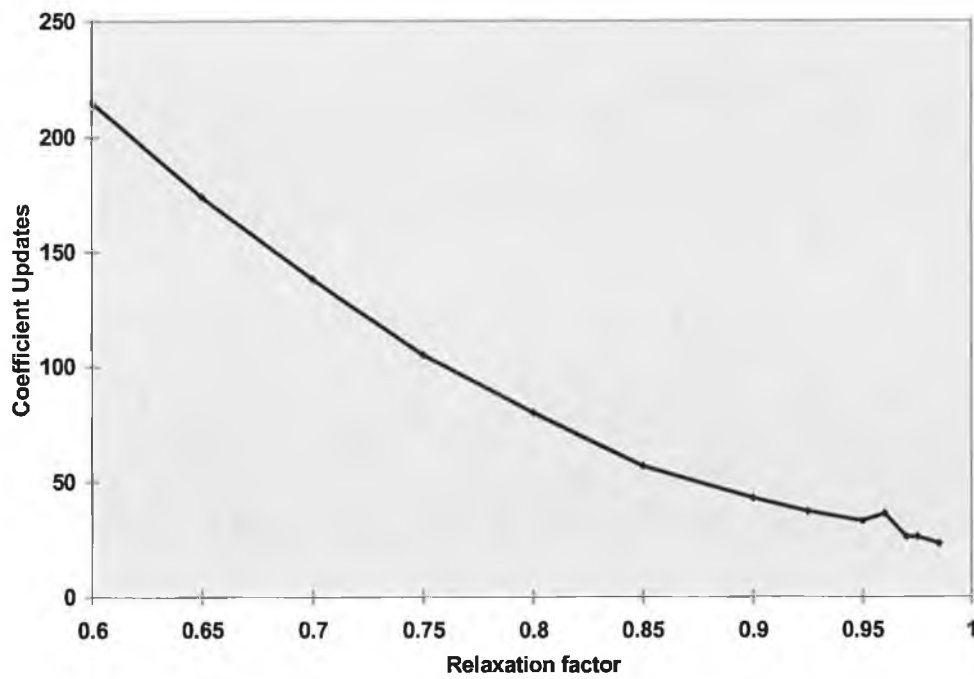
Figure 9.3. Variation of the total number of coefficient updates with relaxation factor.

The original CELS did not show the instability at high values of relaxation factor that is clear here. Values greater than those presented, i.e. greater than 0.985, provoked divergence in the solver and some of the values greater than 0.95 and intermediate to those shown also resulted in divergence. It appears that the optimum value for this application is approximately 0.925 and operating above this moves the solver into a region of instability.

The number of coefficient updates required for a solution is strongly dependent on the relaxation factor. While the number of updates was seen to continue decreasing as relaxation factor increased, there was also evidence of the zone of instability in these results at values above 0.95. For some of the situations producing divergence, the number of initial conditioning iterations was increased but this did not bring any improvement.

The general behaviour of the residuals (equation (4.2)) during the solution was as shown in figure 9.4. Convergence was initially rapid but then the rate of progress slows down and it shows the classic tail where the solver becomes less efficient. This is taken to be due to the smoothing property of the solver as discussed in Chapter 5.

The behaviour of the CELS solver at high relaxation factors was examined. The results for three values are shown in figure 9.5. A log scale is used for the vertical axis to allow comparison of the later stages of convergence. All the three cases showed the same general progress through high residual values and on down to the tailing off of the efficiency as the solution proceeded. What caused the increase in the solution time as the relaxation factor increases to the point of causing instability is that the effort required for the first few iterations increases. This occurred very dramatically in this example between 0.975 and 0.985.

The other feature to be noted here is the improvement in the rate of the residual reduction after these first iterations. Further, as the relaxation factor increases the rate increases and this suggests that a variable factor might be of value in increasing the overall rate of convergence. An initially low value would control the initial effort and
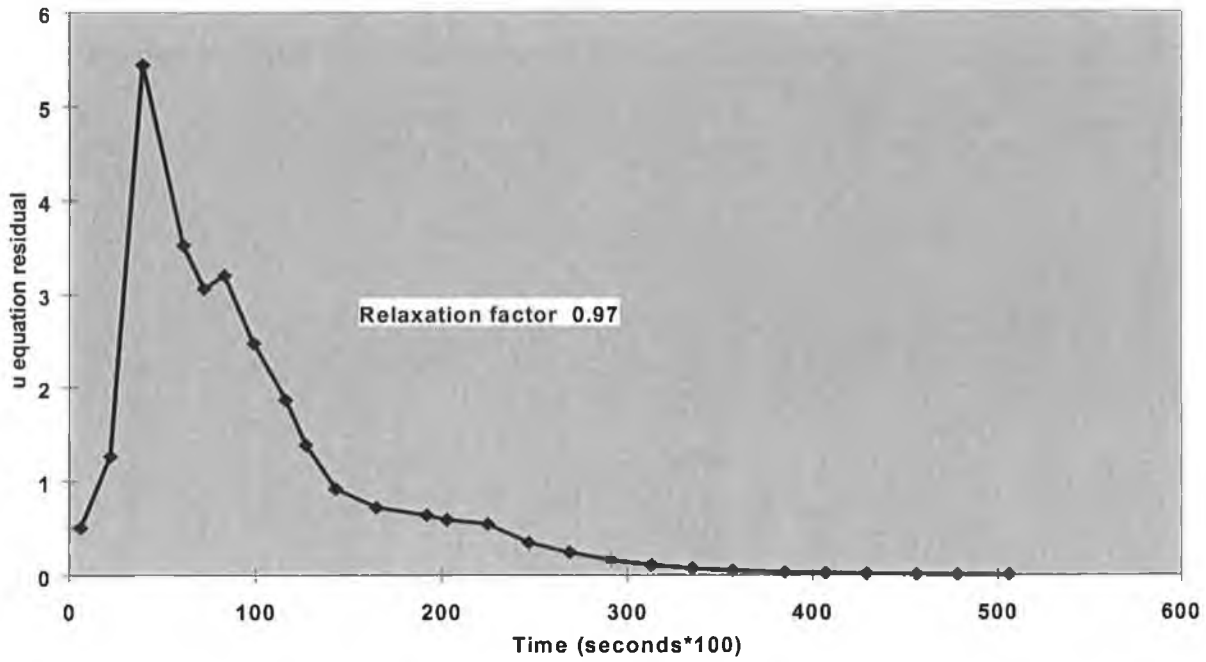
Figure 9.4. Typical convergence as monitored by the calculated residual reduction during coefficient update.
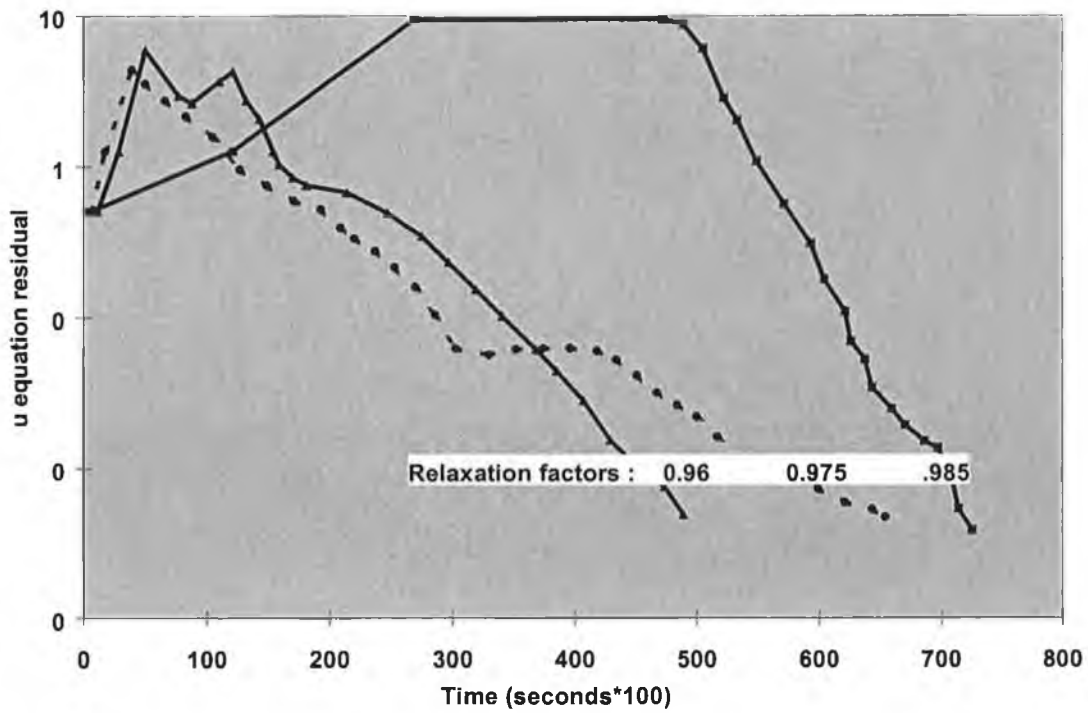
Figure 9.5. Progress of the residual reduction for a number of high relaxation factors.

170

then the process should finish with a high value to take advantage of an improvement in rate of convergence at the end.

This was implemented for the same flow. The relaxation factor started at a value of 0.9 and then increased uniformly over ten coefficient update cycles to a maximum value of 0.99. This particular pattern gave the best results. Beginning at a significantly higher value, or increasing more quickly, moved the procedure into instability. These values seem to place it far enough outside the unstable region to provide a robust solver that performs well over a range of variations in the flow geometry.

The convergence time for the current problem was 3.24 seconds, which is 66% of the best time for a fixed relaxation factor or 53% of the best time for a fixed relaxation factor set to a value low enough to avoid instability. Thus manipulation of the relaxation factor can halve the convergence time. This was adopted as best practice for
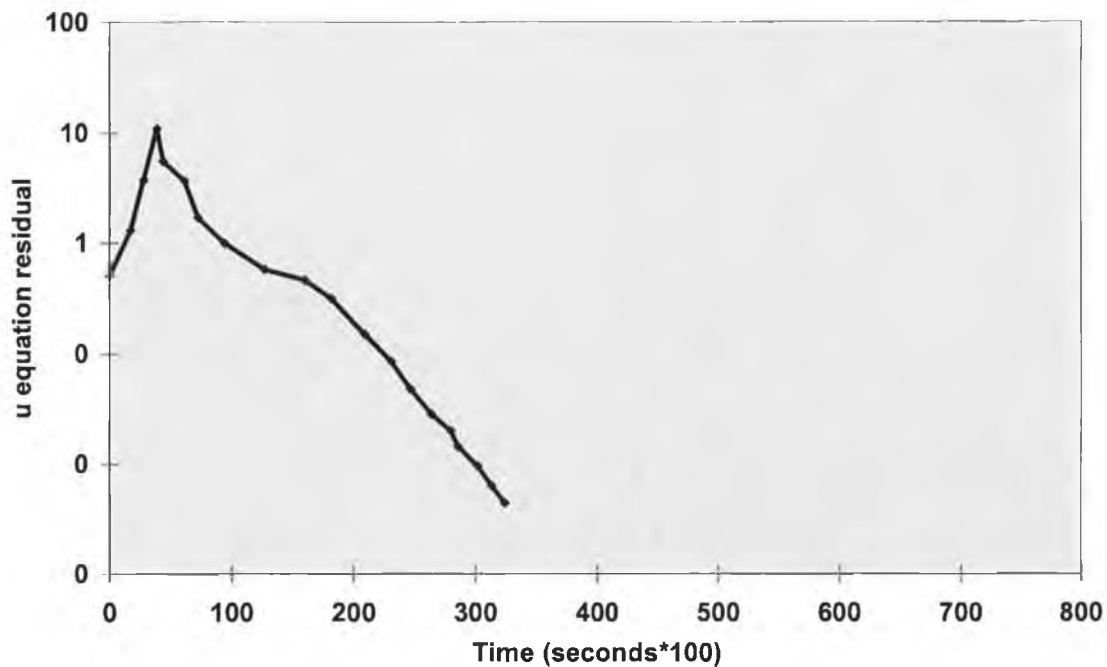
Figure 9.6. Progress of the residual reduction for a variable relaxation factor.

the current flow geometry. The progress of the residual reduction is shown in figure 9.6 where, for easy comparison, the time scale is the same as that in figure 9.5. It should be noted that while the residual reaches higher values early, the initial effort is relatively low and a good rate of convergence is rapidly achieved. The number of coefficient updates was 20, which is greatly reduced from the 37 required for the best, constant, under-relaxation factor. Minimising the number of updates is an advantage for the efficient application of multigrid procedures associated with this method.

Taking the variable relaxation factor as a standard practice, the effects of the inner, CELS relaxation factor and the degree of residual reduction, α, were investigated.

Figure 9.7 shows the number of coefficient updates versus residual reduction factor α. This general behaviour is as expected. As the residual reduction factor decreases more effort is required to meet the residual target. A better approximate solution is produced for each iteration.
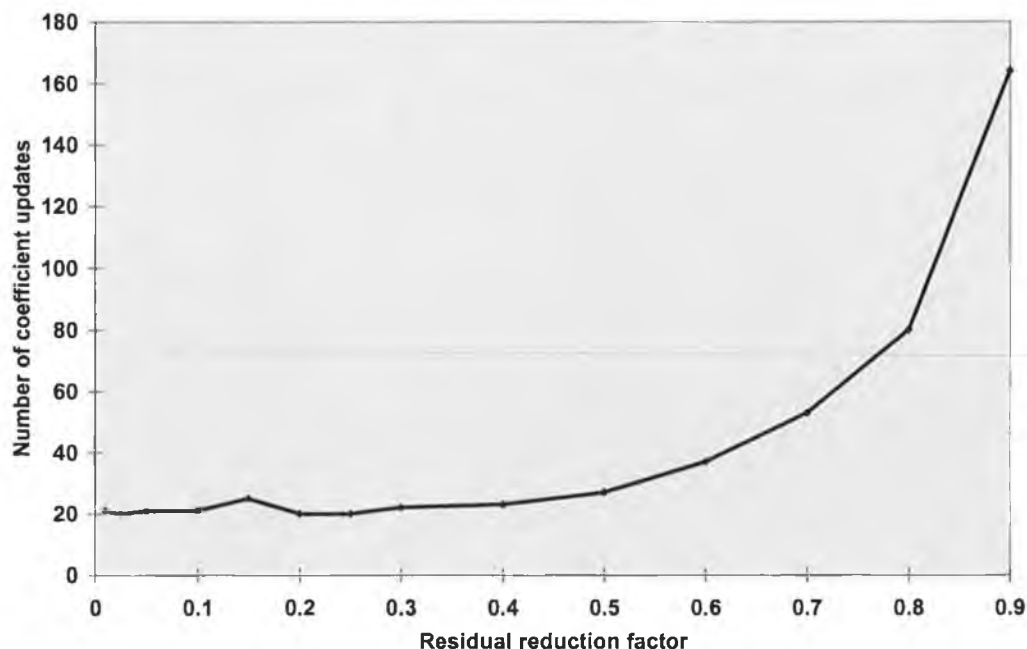
Figure 9.7. Increase in the number of coefficient updates with residual reduction factor.

As the degree of solution for each set of coefficients increases, it would be expected that the number of coefficient updates would be minimised. Below a given value (approximately 0.4 here) the return on the extra effort is minimal. There is no benefit to be had from over-solving on each coefficient update cycle and a balance has to be struck between effort per update and the overall progress.

On the other hand, at high reduction factors the equations are insufficiently solved before carrying on to the next update and the result is poorly specified fields available for the calculation of the next set of coefficients. Not using the available information correctly results in the overall inefficiency and a greater number of update cycles. As less effort is put into each cycle, more cycles are required.

In CELS, the variables used for the coefficients are calculated together within the solver and it is possible to strike a balance that produces an efficient solution technique. In the segregated solver used for SIMPLE, there is more than one equation solution used in the calculation of the velocities and the interplay between these places a limit on the effort that is productive. The result is a system that is inherently prone to a high number of coefficient updates. The low number of updates is an advantage of the simultaneous solution technique as compared to segregated techniques.

In figure 9.8, the convergence times are given for the same values of the residual reduction factor as those shown in figure 9.7. These times complete the information that is required to settle on an optimum value. Note firstly that the effort increases as the number of updates increases at the higher values. Between 0.4 and 0.9 there is an eight-fold increase in coefficient updates but the total effort is doubled between these values. This reflects the much reduced effort per iteration that follows when only approximate solutions are sought. As the target for the residual reduction falls, the effort increases quickly and this is a reflection of the smoothing property within the CELS solver linear sets. The greater the degree of solution that is required, the more the inner solver moves into its region of inefficiency.

At the lower end of the range of the residual reduction factors shown, the overall effort begins to increase again. At this stage, the number of coefficient updates is essentially fixed but more and more effort is being expended in achieving better solutions to each linear set.
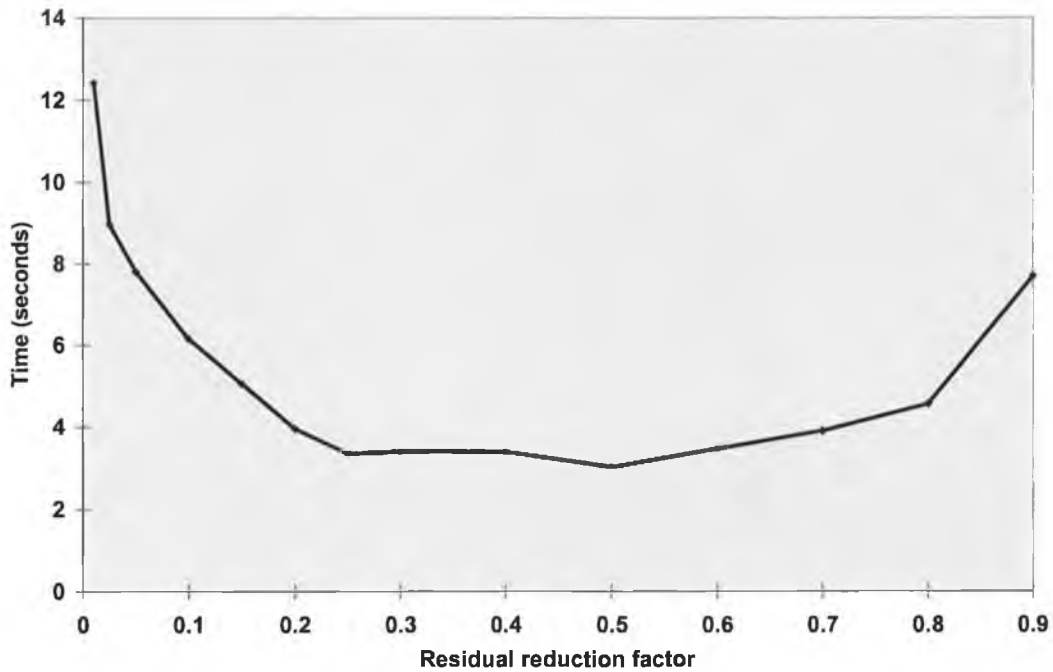

Figure 9.8. Variation of the convergence time with the residual reduction factor.

The result is wasted effort. As long as the number of updates is constant, improving the solution for each update simply results in more overall effort. The current problem can be solved in the range of reduction factors from 0.25 to 0.6 with little variation in the solution time.

The final factor to be investigated was the inner or CELS relaxation. As pointed out in section 6.4, some degree of under-relaxation was found beneficial in the original study. A determination of the best practice for the current equations and problem type was made and α was set to 0.35.

174

Figures 9.9 and 9.10 show results in terms of convergence time and number of coefficient update cycles for a range of these relaxation factors.
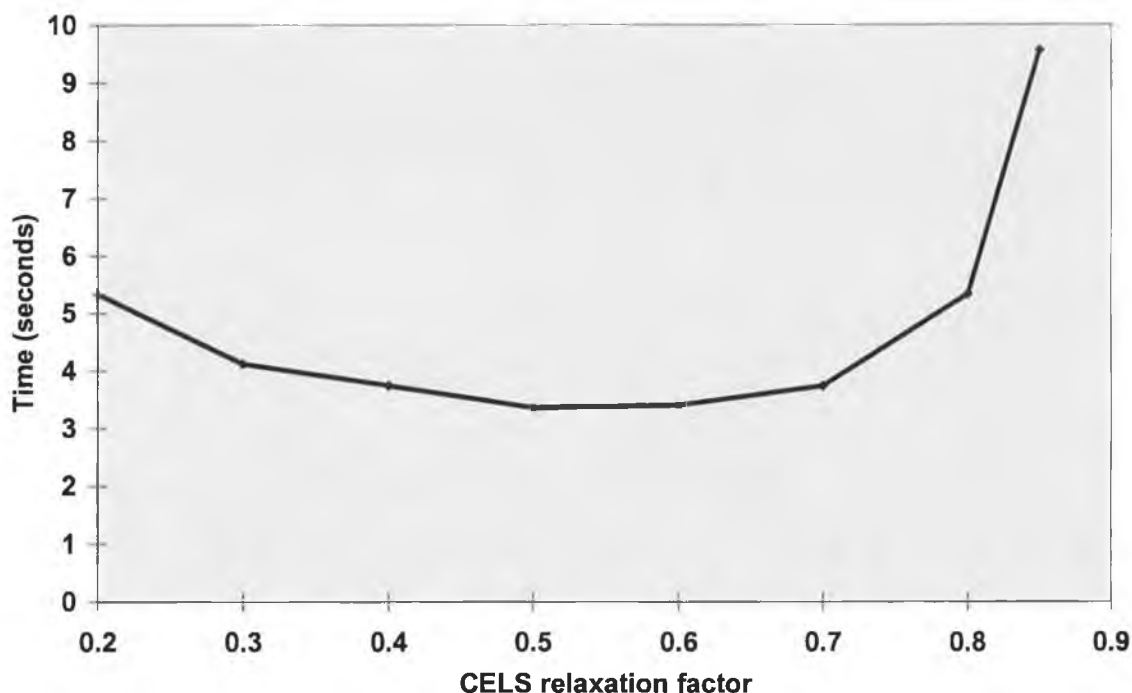
Figure 9.9. Variation of the convergence time for a range of values of CELS under-relaxation factor.

Firstly, there is a requirement to have under-relaxation in the CELS procedure in order to avoid instability. In figure 9.9, the convergence time increases sharply above 0.7 and there is no convergence above 0.85 and the solver seems to be quite robust when operated at the values that minimise the solution time. At very low values the performance is degraded significantly and the degree of under-relaxation simply serves to slow down the operation.

The number of coefficient updates is relatively unaffected compared to the variation with other controlling factors but there is a definite minimum at the same values that produce the minimum time. The fall in the number of coefficient updates at 0.825 is

of no importance as this is close to instability and corresponds to an increase in overall effort.
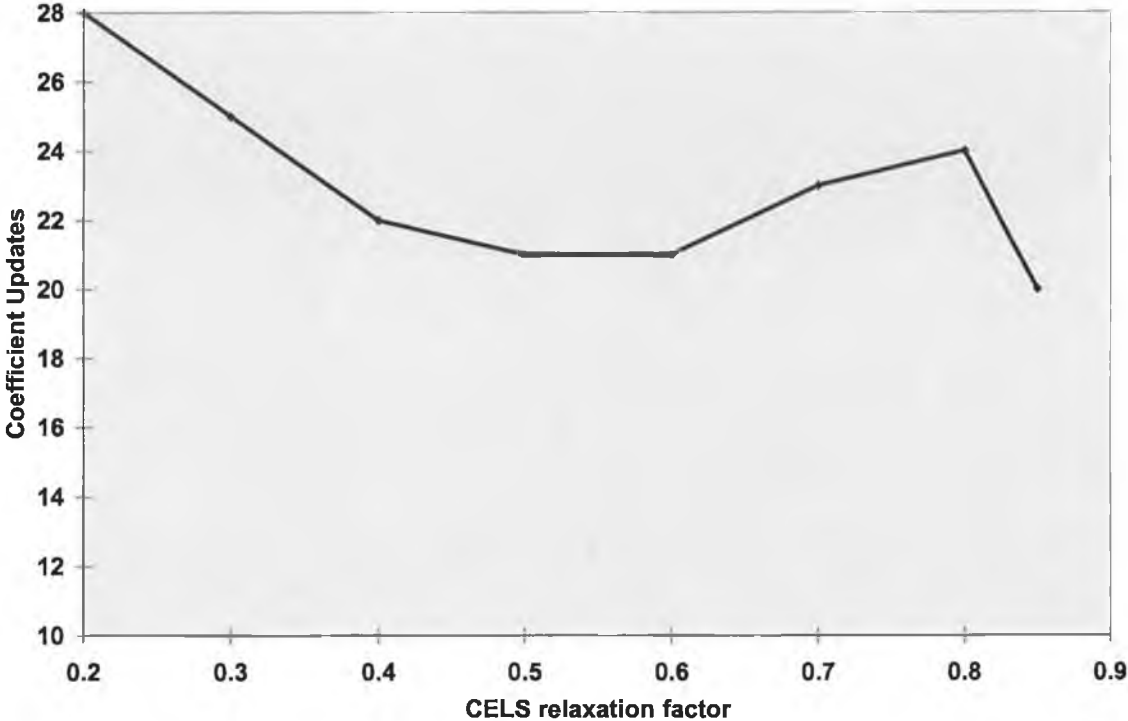
Figure 9.10. Number of coefficient updates for a range of values of CELS under-relaxation factors.

9.3 Turbulent Flow

Having established the general characteristics of CELS for the velocity component equations and the requirements for stability and efficient operation, the turbulence equations were introduced into the set to be solved.

In the first instance, the segregated nature of SIMPLE was adapted by simply replacing the velocity and pressure correction equations with the CELS procedure. After solving for u, v and p for each coefficient set, the turbulence energy and dissipation rate equations were solved, as before, in turn. The effective viscosity was then calculated and this was used in the subsequent re-calculation of the coefficients for the velocity equations.

An immediate difficulty arose because after the velocities had satisfied the residual reduction criterion there was disagreement between the results from the SIMPLE calculations and those shown by the converged CELS results. The turbulence equations did not control the convergence in TEACH in that their residuals reach the required level earlier in the calculations than do the mass and velocity residuals.

However, there were a large number of coefficient updates and this ensures that the turbulence energy and dissipation rate were sufficiently converged. The coefficients depend on the velocities but the small, fixed effort used was enough because SIMPLE reduces the residuals to their target values very slowly, allowing k and ε to reach a solution because the source terms, which depend on the currently computed estimates of k and ε, have sufficient iterations to adjust to their final values.

The errors, i.e. disagreements between the TEACH and the CELS results, were largely in the final values of k and ε. Increasing the effort devoted to their solution within the segregated solver did not significantly improve the situation. There were more fundamental aspects of the overall solution procedures that need to be revised. If the coefficient updates were continued after the velocities had converged then it was possible to produce the correct converged results. However, in order to achieve this

improvement, the computational effort required unfortunately increased greatly and eroded any savings that CELS could produce relative to SIMPLE. This was essentially moving back to a situation where a large number of coefficient updates were required.

In the two-dimensional test problem of Section 9.1, it happened that the velocities could converge quickly because the turbulence appeared to couple weakly with the velocity components and the pressure equations. In this situation, it was possible to improve the accuracy of the results for k and ε by increasing the number of sweeps of the equations in their solution. Where there was a stronger coupling the convergence of the velocity component solutions could be slowed by the poor estimates computed for the effective viscosity. This was observed in the three-dimensional flows calculated later and it limited the effort per coefficient update that was useful and limited the scope for the application of a multigrid technique that applied to the inner solver only. An improved method for calculating k, ε and effective viscosity would serve to reduce the number of updates required and maximise the effort per coefficient update.

9.4 Improvement in the turbulence calculation procedure

The turbulence parameters in the energy and dissipation rate equations are coupled to the velocity components and pressure (calculated by the CELS solver) mainly by the effective viscosity. With this in mind, the structure of the SIMPLE-based segregated solver was examined with a view to optimising the relationship between the air velocity component/pressure solver and the turbulence equations.

Firstly, there is a mutual dependence between the turbulence equations. The source terms in each of the two equations (for k and ε) incorporate information generated by the other. The k equation uses the viscosity which is calculated from the latest available values of k and ε. The ε source term incorporates k explicitly. In the segregated approach, these equations are solved in turn and then the viscosity is updated before re-entering the air velocity and pressure correction equations. The

effect of the ε calculations on the k equation is therefore postponed for a full solver cycle.

Therefore, in the first instance, there is a case to be made for breaking the turbulent flow calculations into two systems which are coupled by the effective viscosity. The first is the CELS system with its internal updating of the source terms and the second, in the same spirit, is the turbulence system with the sequence of events altered to allow updating of the source terms within the system.

Further examination of the TEACH procedures revealed another limitation that was not of great significance in the segregated system. The wall shear stresses that are used in the turbulence energy boundary conditions are calculated as part of the boundary conditions for the air velocity equations. These stresses include the turbulence energy and the current air velocity values and represent another coupling between the two equation systems.

They also represent another lag in the segregated solver approach as coded in TEACH. The wall sheer stresses are calculated at the start of the velocity-pressure system and are not updated in TEACH before the turbulence equations are solved. This is probably not a significant problem where a large number of coefficient updates are used. In the CELS system the effects of the lag can become more apparent because more effort is required per coefficient update and there are large changes in the velocity dependent coefficients of the turbulence equations.

To overcome the lag caused by these calculation sequences, the calculation of the stresses would be included in the turbulence solution cycle and within this the latest available estimates for k could be incorporated as part of the operation.

9.5 Solution sequence for the calculation of k and ε

The complete sequence of operations for the turbulence system would therefore be :

(i) Calculate the k and ε coefficients. The terms which are dependent only on the air velocity components should be calculated only once.

(ii) Calculate the wall shear stresses using the output of the CELS system and the current k values.

(iii) Include the viscosity contribution in the generation terms and calculate the full k source terms.

(iv) Solve the k equations using the tri-diagonal solver.

(v) Calculate the ε source terms, including the latest k values.

(vi) Solve the ε equations using the tri-diagonal solver.

(vii) Calculate the effective viscosity and the latest k and ε values.

(viii) Return to step (i) and repeat the cycle as required.

This set of equations could have residual monitoring in the same way as the CELS system but, in practice, a fixed number of cycles was found to perform well. The effort that would be required to calculate the residuals was saved.

9.6   Calculation of the turbulent flow

With the improved method of calculation in place it was possible to examine the efficiency of CELS in improving the solution time for a turbulent flow calculation.

The first factor examined was the residual reduction factor (α) in CELS. The turbulence equations were solved with a relaxation factor of 0.8. This seemed to give a good performance, while keeping the solution of the system of equations stable. Operating at 0.9 gave a good performance but not across the full range of values of α and the calculations diverged unpredictably as α was varied.

In addition to the initial conditioning of the air velocities and pressure mentioned above it was found that another coefficient update cycle was necessary before the introduction of the revised k and ε calculations. This was required to ensure stability and correct initial calculations of k.

Figure 9.11 shows the convergence time as a function of α. There is a more localised minimum than in the laminar case and the minimum is in the same region at a value of approximately 0.45. The general behaviour is the same as before but there is a more dramatic rise in the solution time as α increases.

The number of coefficient updates (figure 9.12) keeps falling as the value of α decreases to 0.25. The computational effort per iteration increased below a value of 0.45 but the reduction in the number of updates could be useful in the context of the potential for improved efficiency of the inner solver.

As mentioned above, the relaxation factor for the k and ε equations is important in that it must be set to avoid instability. The question arises about its optimum value from the point of view of solution time. Using $\alpha = 0.45$ in the CELS solver, a number of calculations were made using different values for the relaxation factor and the results obtained are shown in figure 9.13.
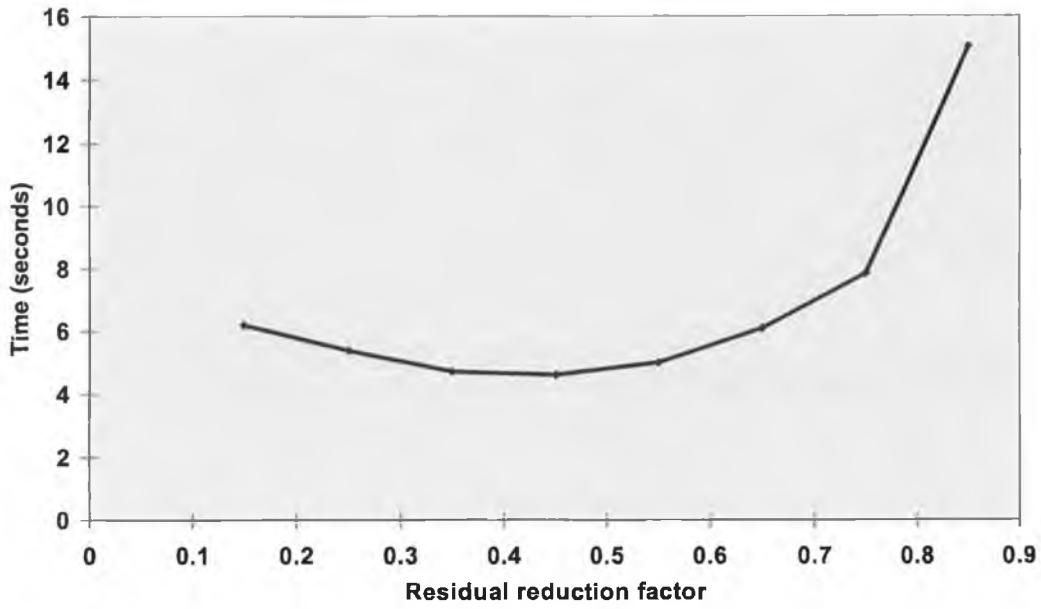
Figure 9.11. Effect of the residual reduction factor on the convergence time.
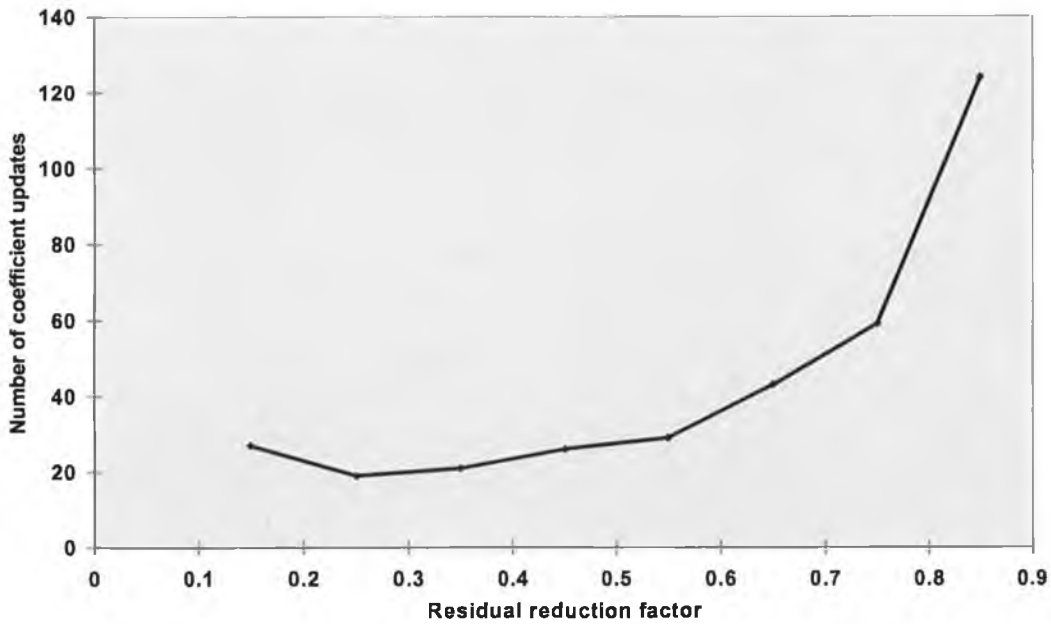
Figure 9.12. Effect of the residual reduction factor on the number of coefficient updates.
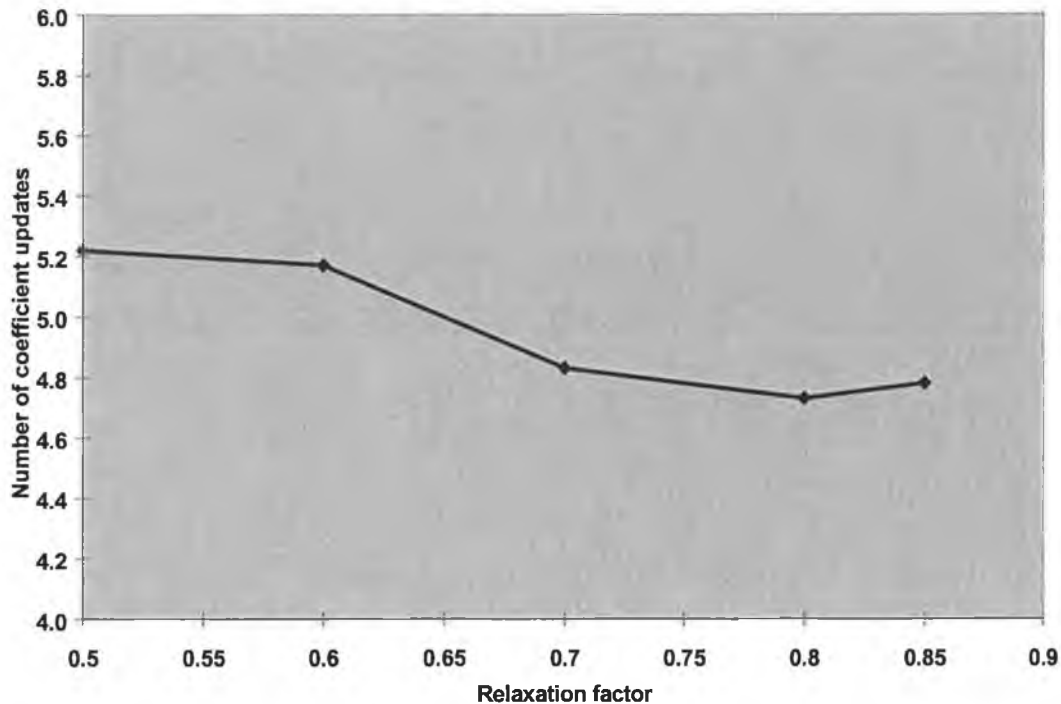
Figure 9.13. Effect on the solution time of the relaxation factor for the turbulence equations.

The effort per velocity update for the turbulence calculations was fixed at 6 iterations of the new solution scheme. There was little variation in the solution time (as shown) or coefficient updates but this did not represent an equivalent achievement. As the equations become more under-relaxed, the efficiency of the calculation procedure decreased so that while the time taken remains fairly constant there is a progressive loss of accuracy, i.e. disagreement with the SIMPLE results. Accuracy can be recovered by increasing the number of iterations in the turbulence scheme but it was more efficient to operate with the relaxation factor at approximately 0.8. This gave good agreement with the SIMPLE results and minimised the solution time.

The possibility of using varying relaxation factors for the turbulence equations was examined. The shortest times to convergence were obtained for starting values of 0.8 and increasing this uniformly to 0.9 over the course of 20 coefficient updates.

183

However, deriving savings from this required a change in the residual reduction factor. With α set at 0.45, the varying relaxation saved 4% but, with α at 0.35, the saving was 13%. This shifted the minimum shown in figure 9.11 and gave the best overall convergence time.

9.7 Comparison of CELS and SIMPLE

The TEACH code, as supplied, provided the under-relaxation factors for the operation of the SIMPLE algorithm.

The convergence behaviour of the optimised CELS and the corresponding SIMPLE solution are presented in figure 9.14 where the grid is 34×34. The CELS solver exceeded the nominal degree of residual reduction required for convergence. For the purposes of comparison, SIMPLE was set to converge to the same degree. At this level of convergence, the two methods produced air velocity fields that agreed to within 2% on average. Convergence was achieved by reducing the normalised residuals to 0.005. Using a target of 0.0005 produced no significant change in the solutions obtained.

The difference in the number of iterations was striking with SIMPLE requiring 372 iterations while CELS converged in 20. Further, the convergence time for CELS was 51% of the time required for SIMPLE.
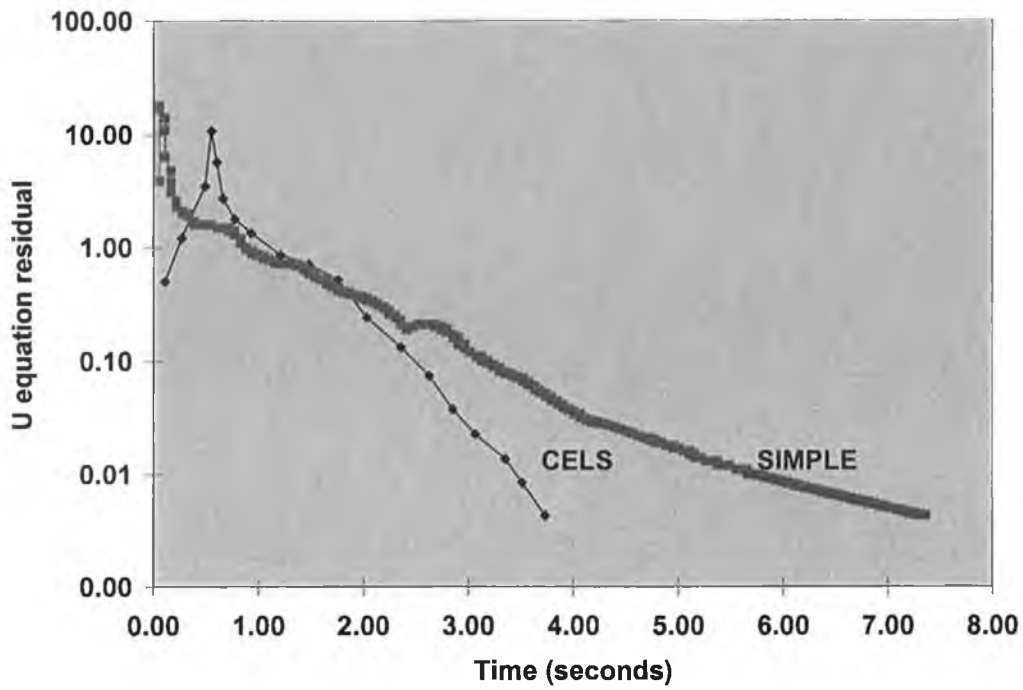
Figure 9.14. Comparison of the convergence times for CELS and SIMPLE (without turbulence model, 34×34 grid, Re = 5500).
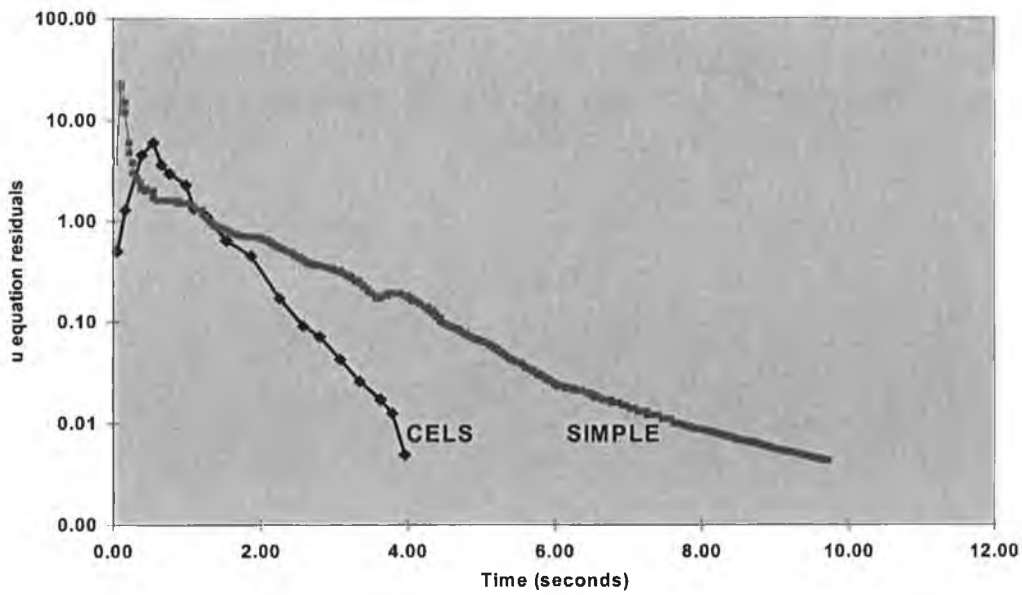
Figure 9.15. Comparison of the convergence times for CELS and SIMPLE (with turbulence model, 34×34, Re = 5500).

The comparison for the turbulent case is shown in figure 9.15. In this case CELS converged in approximately 40% of the time required for the SIMPLE solver. CELS converged in 19 iterations while SIMPLE required 337. The number of cycles in the turbulence calculations was 6 with 3 sweeps of the tri-diagonal solver per equation per cycle. This gave agreement to within approximately 2% for both the fluid velocity components and the turbulence parameter fields. Calculation of the local turbulent viscosity from the turbulence parameters, using equation (3.9), showed significant turbulent effects in the flow.

The saving in computational effort due to CELS falls within the range found by Galpin et al. (1985) but at the lower end for the laminar flow. Savings of up to 70% were reported but, as noted earlier, improvements were both problem and grid density dependent. These savings are useful and justify the application of multigrid techniques to the inner iterations in that the effort per coefficient update is relatively large.

9.8   The Application of ACM

ACM was applied to the turbulent problem discussed in Chapter 6 when comparing the performance of CELS and SIMPLE. In general terms, it was possible to achieve savings in the computational effort by the application of three coarse grid levels to the 34×34 grid. Grid density was reduced by a factor of two for each new level. The addition of each level produced a saving but the use of ACM in this situation was not straightforward in the sense that considerable tuning was required for optimum performance.

ACM was first applied using one coarse grid. The coefficients were transferred to this, the coarse grid equation solved and corrections applied to the fine grid (level 0). The residual reduction factors were tuned for best performance. The final values were 0.35 for level 0 and 0.2 for level 1. This combination reduced the convergence time to 34% of that for SIMPLE. No CELS relaxation was required on any of the coarse grids.

Coarse grids at levels 2 and 3 were added in a similar fashion. In each case the solver required re-tuning, at the level immediately above the latest one added, to find the best performance. Difficulties with the compiler prevented the implementation of level 4. It was found that data was not being passed correctly between subroutines via Fortran COMMON blocks. Testing failed to find any coding errors and it was concluded that there might be a bug in the compiler related to the level of nesting of subroutines as the routine for each multigrid level called the routine for the next coarser grid.

Results for the successful application of ACM are summarised in table 9.1, where the savings in time to convergence for the multigrid solvers invoking different numbers of coarse grids are compared with the time required for the SIMPLE solver is to converge. While these results are encouraging in general terms for the use of ACM with a turbulent flow, they are a little artificial in that a good deal of effort was put into tuning the system of grids. For example, on level 2 there is a change in the residual reduction factor from 0.05 to 0.025 as the solution proceeds. The results thus represent the best performance.

Table 9.1. Convergence times for a series of coarse grids as a percentage of the convergence time for SIMPLE.

	Convergence time (percentage of the result for SIMPLE)
SIMPLE	100
CELS	40
CELS + level 1	34
CELS + level 2	27
CELS + level 3	22

While compiler problems, as outlined above, prevented the introduction of a further level with a 4×4 grid and a direct solution at that level, it is possible that the improvements would have continued but Smith (1990) reported some instability when working at the coarsest grid levels.

The boundary conditions for the curved wall of the mushroom tunnel were inserted, as in Section 8.1, and the performance of CELS/ACM was evaluated. No difficulty was encountered at level 0 (the fine grid). At best performance, with α set to 0.8, the solution time was reduced to 62% of the SIMPLE time. This was not as effective as the square domain calculations. The high value of α on the fine grid meant that the effort per linear set was reduced and this gave rise to a larger number of coefficient updates.

The first coarse grid was introduced as above. There was a significant improvement in performance with convergence time being reduced to 53% of the SIMPLE time. The residual reduction factor at level 1 was also 0.8. The addition of a level 2 grid produced no effect. The high residual reduction factor at level 1 was probably mainly responsible for this as only one iteration was required at level 1 and level 2 was thus never required to accelerate the solution at level 1. The results are shown in table 9.2. The improvements are worthwhile but it would be desirable to find a way of reducing the number of coefficient updates.

Table 9.2. Savings relative to SIMPLE for the case of a modelled mushroom tunnel wall.

	Convergence time (percentage of result for SIMPLE)
SIMPLE	100
CELS	62
CELS+level 1	53

Given that these results are on relatively coarse grids and that the improvements due to ACM shown in table 9.1 are of the same order as those quoted by Galpin et al. (1986) for the temperature-velocity coupling, these results were regarded as sufficiently encouraging to extend the ACM method to three dimensions. Also, the practice of varying the relaxation factors for the velocity component equations was not desirable because it introduced another factor to tune for best performance and it was encouraging that initial testing of the CELS solver in three dimensions had shown that this approach was not required. No further testing was undertaken at this stage as the emphasis of this work was to be on the three-dimensional solver.

The work presented in this chapter has been published as a refereed paper in the proceedings volume of a specialist workshop (Grant and Williams, 1996).

Chapter 10

Solution of the fluid flow equations using CELS3D

10.1 Solution procedure

The problem to be solved is that of transferring a two-dimensional, plane-based solver to the three-dimensional flows that are to be solved in the mushroom growing application. Given that a plane solver is to be adapted to efficiently solve a three-dimensional flow problem, a procedure should allocate the solution effort in such a way that all equations in the velocity/pressure set receive equivalent effort. The use of planes to sweep through the flow domain throws up a number of possible variations in procedure. Those examined form a set of alternatives that can be regarded as variants or special cases of the same operation.

These variants were produced in the effort to achieve savings using ACM in three dimensions. This aspect will be treated later in the discussion of that topic.

The geometry of the flow domain used for the initial testing is shown in figure 10.1 and the boundary conditions were handled in the same manner as described in Section 4.10. The air flow enters via a long, narrow slot and exits at a low level exhaust. This is a first approximation of duct-delivered airflow as used in mushroom-growing room that would contain multiple tiers of compost. The dimensions were set to approximate the actual dimensions of a mushroom growing room, i.e. 3 m square. The computational grid was $34 \times 34 \times 34$. The inlet slot was 1 grid unit high and 30 grid units in length. It was located on the west face of the flow domain at $j = 28$. The outlet

was a square aperture on the back face of the flow domain. It was 3 grid units wide and 3 grid units high. It was located from $i = 30$ to $i = 32$ and from $j = 5$ to $j = 7$. As air is delivered as jets from a series of small apertures rather than as a continuous slot, the entry velocity was set low at 0.2 ms^{-1}. This was an approximate equivalent to flow due to the jet speeds but it was spread across the face of the slot rather the smaller total aperture area. This would give a very much more uniform initial air stream that would be the case in reality but it was used purely for convenience at this stage of the model development. Modelling of the actual flow in a mushroom growing room would be expected to require more accurate representation of the individual jets leaving a distribution duct and probably a finer grid than it was possible to use in this thesis (due to the memory limit on the Fortran compiler that was employed).

The first calculations were performed for the laminar flow situation, i.e. without the turbulence model. This was of interest because CELS improves the velocity/pressure coupling and it provided a result to contrast with the overall performance incorporating extra equations not included in the CELS procedure. The Reynolds number was set at 3,500.

The plane sets had to be solved in a particular order. Calling the planes defined by *uv*, *vw* and *wu* velocity component pairs, 1,2 and 3 respectively then the rotating orders 123, 231 and 312 produced convergence with very similar effort. 213 and 312 produced divergence and 132 gave a wide initial swing in the residuals and a very slow convergence followed. The 231 sequence was marginally faster than the other two options and was adopted as standard procedure.

It is a significant change from the two-dimensional procedure that it is no longer possible to satisfy continuity after each plane solution. An inlet, for example, was not necessarily balanced by an outlet in the same plane. Continuity 'errors' accumulate in the first sweeps of the calculation procedure and are then resolved by further sweeps in other directions. At first, the continuity relation accumulates the errors on the last line of each plane.
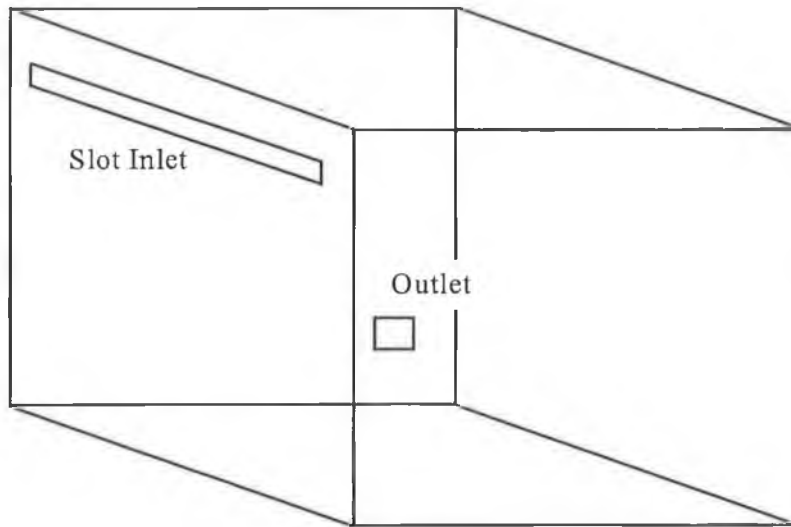
Figure 10.1. Flow geometry for the three-dimensional testing

The first iteration (coefficient update) has to be solved sufficiently to remove these errors and allow the calculation of coefficients that will allow convergence in the second and further iterations.
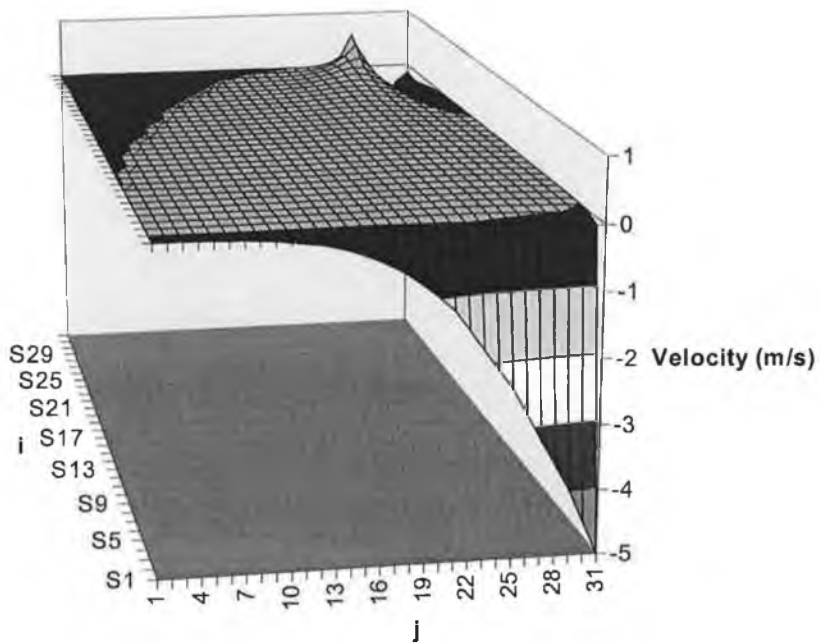
Figure 10.2. _v_ velocity component results after 1 inner iteration.
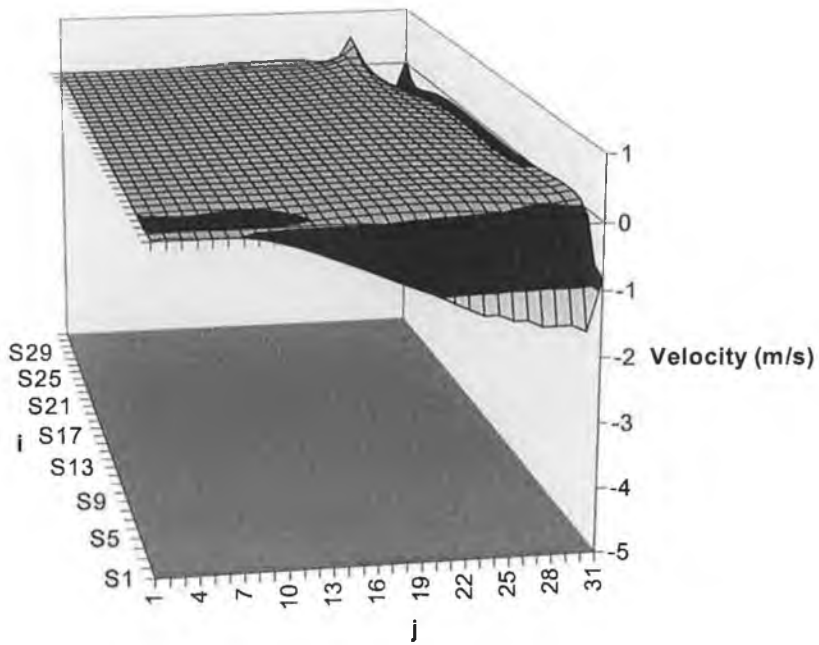
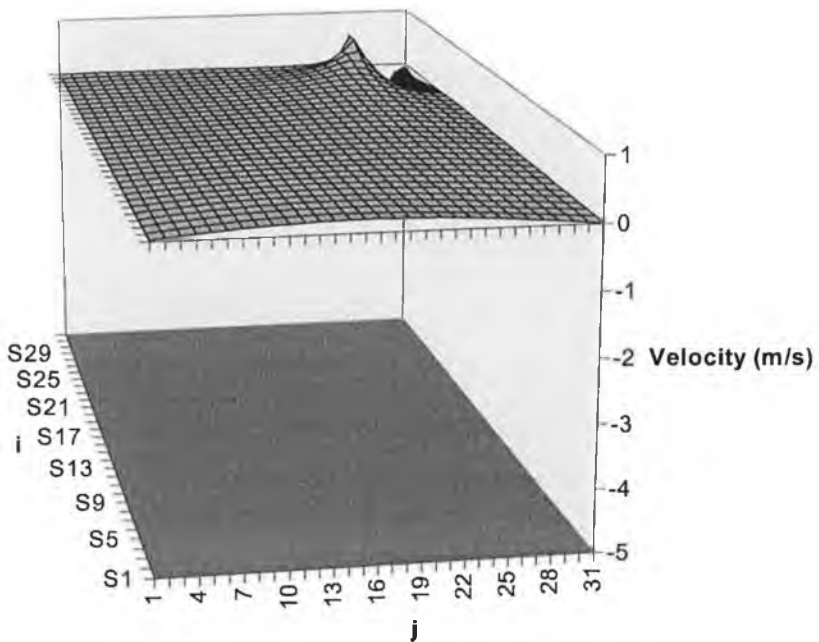Figure 10.3. *v* velocity component results after two inner iterations.

Figure 10.4. *v* velocity component after twenty inner iterations.

Figure 10.2 shows the continuity 'errors' accumulated on the last line by the mass balance calculation. This is a k plane (k = 15) , i.e. the air velocities are calculated on a vertical uv plane. The feature at the top right of the graphic is due to the slot inlet. The flow into the plane at that point is unbalanced because the outlet is on a different k plane and in this first sweep the z velocity components are zero. The continuity relation then drives the flow across the plane as the line solution sweeps across it.

Figures 10.3 shows how the flow imbalance is reduced on the second sweep after the off-plane air velocity components have been updated and figure 10.4 is the result after twenty uv plane sweeps. The errors decrease rapidly as the solution progresses.

The last line of a sweep is completed by moving from cell to cell along the line updating the velocity components by means of the continuity relation, as explained in Section 5.3. Experimentation with the direction of the solution procedure on the last lines in each plane sweep showed that, in each case, they should be solved in the same sense relative to the line solver progress through the solution plane. Experiments with different strategies, such as alternating directions between sweeps in the same plane or between planes, caused the CELS3D performance to deteriorate and in some cases to diverge.

Similarly, there is a direction associated with the solution of the penta-diagonal equation along the lines of a given plane. Coefficients are constructed by moving through the grid cells in one direction and then calculating the air velocity component values on the return. The direction of the solution procedure was examined and was found to make no significant difference to the overall performance.

As noted above, a number of variants of the procedure for sweeping the plane sets were developed.

Variant 1

The equations were solved in pairs, i.e. with one set of air velocities held constant, the other two were solved by sweeping the solution domain plane-by-plane. The solver stepped through the planes determined by a given pair of air velocities by sweeping through the flow domain in the third dimension.

The sequence of operations was that one set of planes was solved by a complete sweep in one direction, e.g. the uv planes were solved by sweeping from low to high values of k. The plane solutions were identical to the two-dimensional case in that both air velocities were solved for one plane before moving on to the next, e.g. u with v by continuity and then v with u by continuity. The next air velocity pair was then solved with a sweep in one direction and, finally, the remaining air velocity pair was done in a similar manner.

The next step then was to repeat the solution of the air velocity pairs but this time the sweeping is carried out in the opposite direction. In the uv case the planes would be swept from high to low values of k this time around. Having achieved convergence, the technique was tested to establish its general behaviour in the same way as for the two-dimensional case.

The first factor examined was the residual reduction factor, α. Residuals were calculated in the same manner as for two dimensions and all three air velocities were monitored. The equation coefficients were not re-calculated until all the velocities met the reduction criterion.

The effects of α on the convergence time and on the number of coefficient updates are shown in figures 10.6 and 10.7. There was no under-relaxation on the equations and the CELS3D relaxation factor was set to 0.7.

195
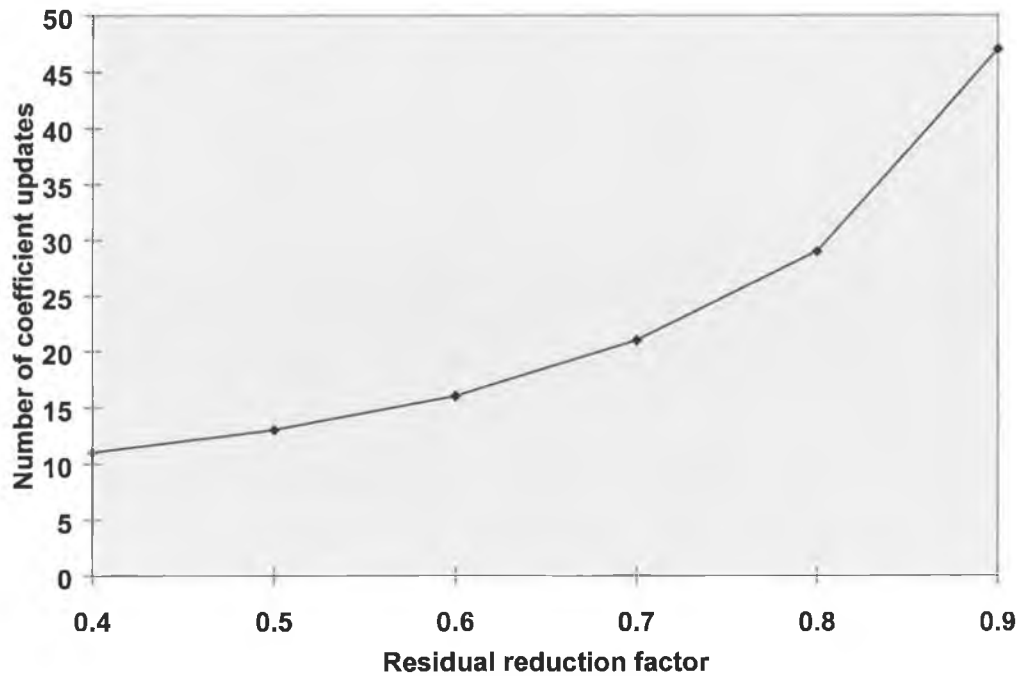
Figure 10.5. Effect of the residual reduction factor on coefficient updates required for a laminar flow solution.

It is very striking that the overall convergence time was relatively unaffected by the degree of the residual reduction but the number of coefficient updates increased quickly with increasing residual reduction factor. On the question of which value of the reduction factor should be selected for normal operation of the solver, it is the requirements of the acceleration of the inner solver that determines the choice. For this problem and Reynolds number, a small number of updates, such as the 17 required for a factor of 0.6, seems to give the right balance between the speed and the number of updates required because it allows scope for an ACM technique to achieve savings. While the time to convergence that was recorded for a reduction factor of 0.8 was not significantly higher than the time for a value of 0.6, the number of coefficient updates is almost double that required for solution at the lower value and this would greatly reduce the potential for ACM.
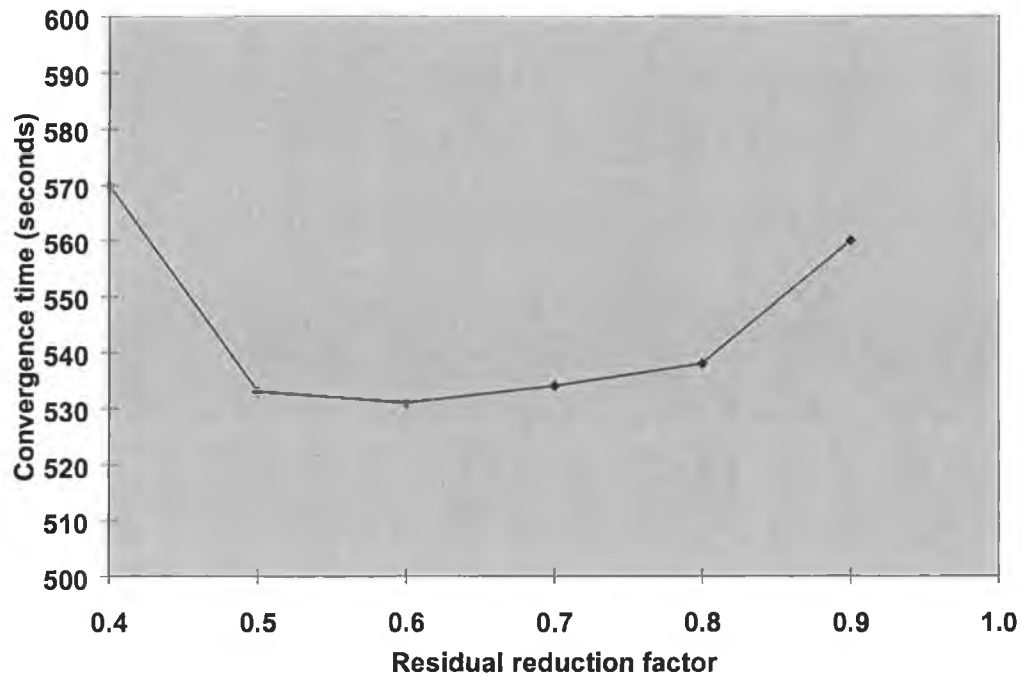
Figure 10.6. Convergence time for a number of values of the residual reduction factor.

The CELS3D under-relaxation factor (the additional relaxation applied during the operation of the solver) was varied to determine the value required to minimise convergence time (figure 10.8). The introduction of an under-relaxation of 0.95 was seen to be very beneficial. The benefits decreased slowly as the value was reduced. The number of coefficient updates was 15, increasing to 16 for values less than 0.8.
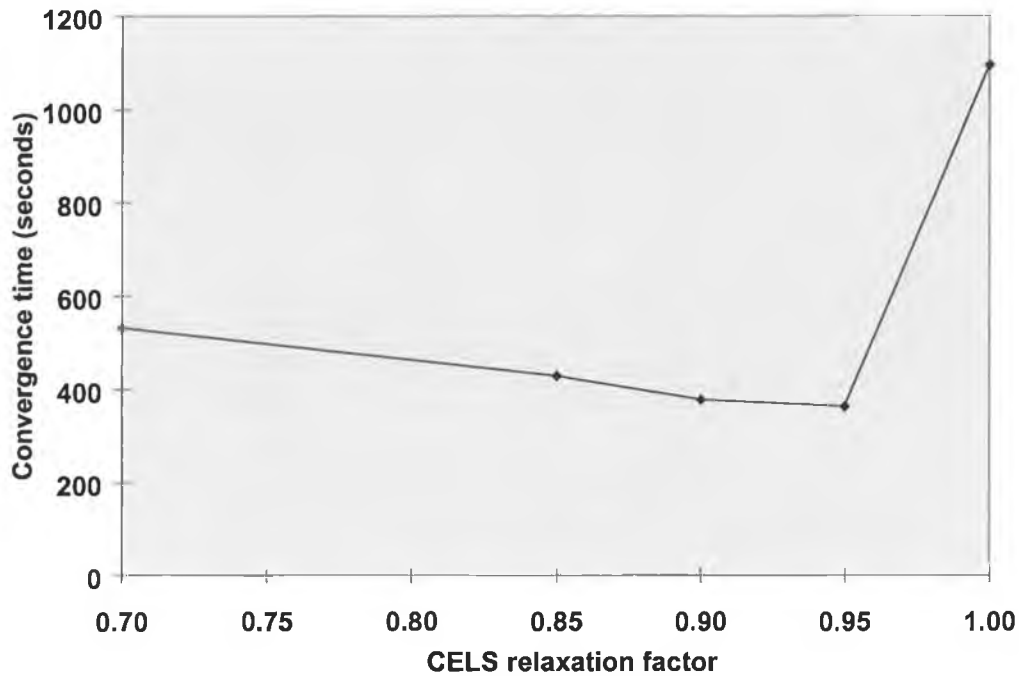
Figure 10.7. The variation of convergence time with CELS3D under-relaxation.

Under-relaxation of the velocity equations was investigated. As noted above, the results presented here were obtained with no under-relaxation (i.e. a relaxation factor of 1). The technique of varying the under-relaxation had produced good savings in the convergence time in the two-dimensional case (see Section 9.2) but no requirement for this was found in three dimensions. Figure 10.8 shows the variation in the convergence time for the three-dimensional flow calculations as the value of the under-relaxation factor is set to a number which was less than or equal to 1. The convergence time increases for all values below 1 so that there is clearly no benefit in under-relaxing the equations and the increase in the number of coefficient updates (figure 10.9) is also undesirable.

In general terms, this behaviour is more in line with that of the CELS solver as reported by Raithby et al. (1985) than was the case for the two-dimensional solver. It should be noted that the proposed pressure correction (see Section 9.1) was also found

to be ineffective in three dimensions, i.e. its application produced no reduction in the convergence time.
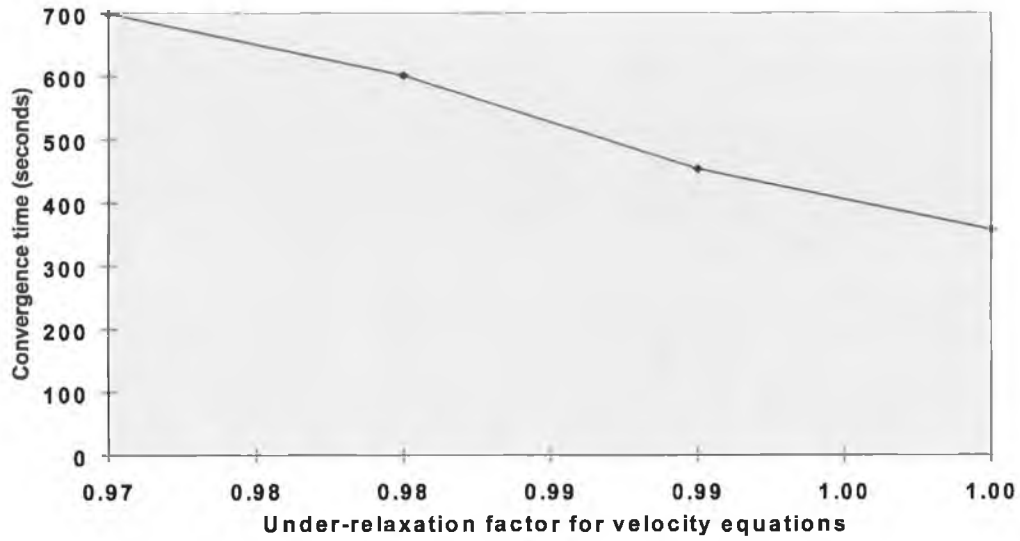
Figure 10.8. The variation of the convergence time with the under-relaxation factor on the velocity equations.
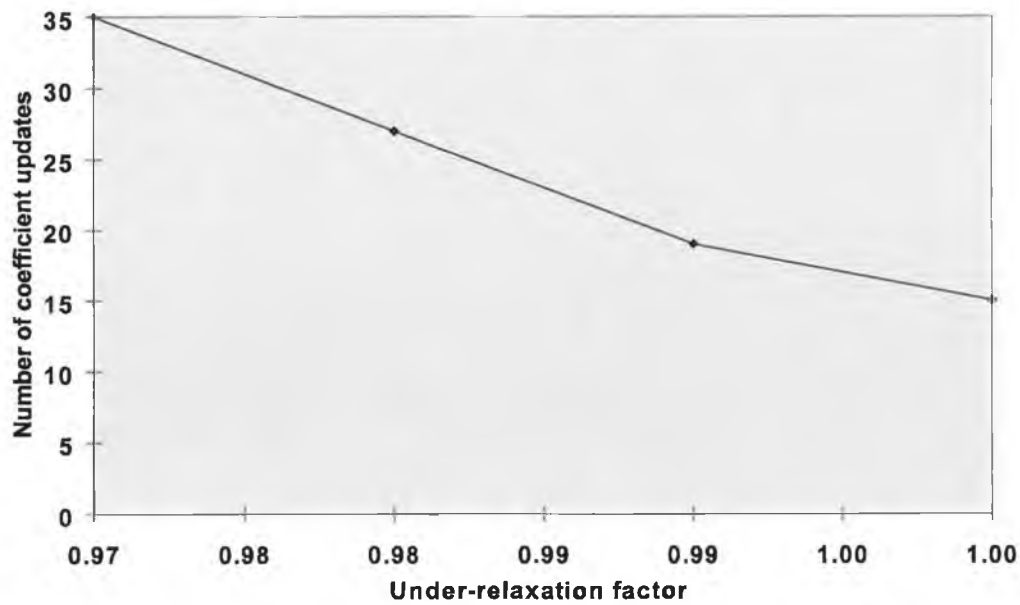
Figure 10.9. The variation in the number of coefficient updates with under-relaxation factor.

For the laminar three-dimensional case, the number of conditioning cycles required on the first coefficient update was just one. This means that the air velocity component fields for the second coefficient update had continuity errors, of the type shown in figure 10.2, on the last line of some plane sweeps but, clearly, this presented no difficulties for the solver.

Variant2

For this procedure, a departure was made from the two-dimensional solution technique where both the air velocity components defining a plane are solved using the penta-diagonal equation before moving on to the next plane. This time, on the first sweep through a plane set, one air velocity component is solved (the other being updated as usual by the use of continuity) but the line solution procedure of the second is postponed. All three plane sets are swept in this manner.

When the return sweep is made then the air velocity components skipped the first time through are solved with the appropriate penta-diagonal solution procedure. It was hoped that this would compensate for a possible excess computational effort in the solution of a given plane set while holding the third air velocity component constant, i.e. solving one plane set too tightly is wasted effort because the third air velocity component is not being updated. It was hoped that this would aid the application of the ACM.

In fact, this did slightly reduce the solution time but the solver needed 10 conditioning cycles to avoid early divergence. The saving in computational effort was 4% with one extra coefficient update.

Variant 3

In this procedure the first set of planes is selected, *uv* say. The *uv* set is then swept in both directions in the third dimension before moving on to another velocity pair.

The performance of this procedure was good for the fine grid alone. The same tuning was used as for variant 1. As with variant 2 though, the method required multiple conditioning sweeps in order to avoid divergence during the course of the first few coefficient updates. It was found that five such sweeps were required. The best convergence time for the test problem was 339 seconds. This was approximately 8% greater than for variant 1.

Variant 4

A further alternative formulation of the CELS3D solver is possible at a lower CELS under-relaxation factor. The solver using this variant diverged if this value was greater than 0.7. In variant 3, there is one solution cycle per velocity pair before moving on to the next. This variant applies more cycles per pair and the question arises about the number of cycles and how it is to be controlled.

A system with fixed a computational effort per plane set on each coefficient update is possible where each air velocity component pair is solved a given number of times before moving on to the next. Varying the number of such cycles produced differences in the convergence time that were almost insignificant. For example, five cycles per plane set per coefficient update produced a convergence time of 586 seconds for the test problem while 10 cycles produced a solution in 562 seconds. The lower number of cycles produced a greater number of coefficient updates at 39 versus 21 for 10 cycles. While significantly faster than SIMPLE, this is clearly inferior to the other procedure variants. Also, the conditioning for these timings was 5 cycles of variant 3. Applying the multiple cycles of this variant from the first coefficient set resulted in longer convergence times.

Convergence is also possible by solving one of the plane sets until the residuals are reduced by a suitable factor and then moving on to the next. Each set is solved once per coefficient update. Some care is required in residual monitoring. If, for instance, the last plane set to be solved is the *uv* set, then the pressure field will move slightly from that calculated by the *vw* and *uw* passes. When a new residual is calculated with the *w* coefficients, this can result in a higher residual than for the *u* and *v* equations and the new residual an reduction target is attained with considerably less effort than for the other cases. This results in an imbalance in the effort applied to the plane sets and severely reduces the efficiency of the solver.

For efficient convergence it was established that it is important that the effort for each plane set should be comparable. This occurs naturally in the other variants.

10.2 SIMPLE vs. CELS3D - laminar flow

The residuals for the best convergence time for the laminar case, i.e. no turbulence model, using CELS3D under variant 1 are presented in figure 10.10 with the corresponding SIMPLE results. Whereas elsewhere in this thesis SIMPLE was applied using the values for the relaxation factors and numbers of sweeps as supplied with TEACH, this was found to give rise to an excessive convergence time for the three-dimensional laminar case. The residual values as the calculation progressed were showed to oscillate strongly and the solver clearly required re-tuning. Reducing the under-relaxation on the velocity equations removed the large oscillations and gave better results. Values of the velocity component relaxation factors of greater than 0.7 produced divergence. A value of 0.7, therefore, was taken as being representative of the laminar performance of SIMPLE, although this value might be too high in some flow situations. CELS3D converged in just 34% of the time taken by SIMPLE, as shown in figure 10.10. The total number of coefficient updates for SIMPLE was 409 while CELS3D converged in just 17.
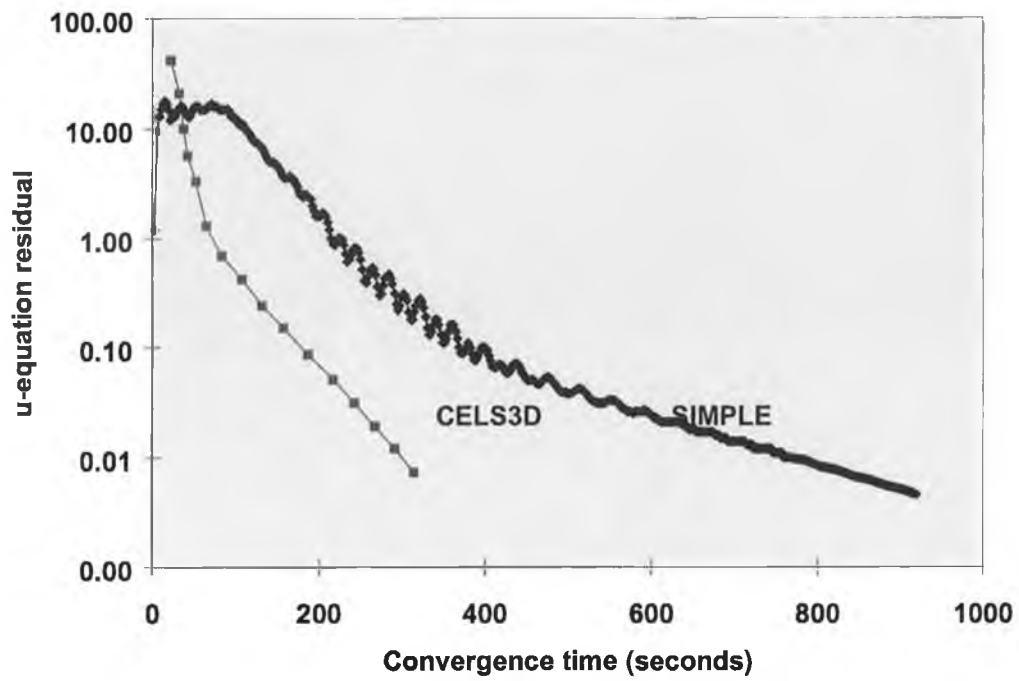
Figure 10.10. Comparison of the convergence time for CELS3D and SIMPLE for a laminar flow (no turbulence model), Re = 3500 (34×34×34).

10.3 Turbulent flow calculations

The turbulent flow calculations were carried out by including the turbulence model in the equation set to be solved and then using the same procedure developed to improve the calculation of k and ε in two dimensions, i.e. the segregated nature of the TEACH approach was revised to provide for cycling through the turbulence equations updating the effective viscosity and the source terms for the k and ε equations.

In three dimensions the use of this procedure was required not only to efficiently achieve a given degree of convergence but also to reduce the number of coefficient updates required.

The test problem was identical to that for the laminar flow examined above but including the turbulence model and, again, the Reynolds number was set to be 3500. There were two main parameters controlling the solution of the turbulence model and these were

(i)   the number of repetitions of the turbulence solution cycle, and

(ii)   the number of sweeps of the calculation domain used in the tri-diagonal solver applied to the k and ε equations.

Having first determined reasonable values for the tuning parameters by a series of tests, these parameters were systematically varied while holding reasonable values of the others and of the general CELS parameters in order to examine the convergence efficiency of the turbulent solution.

Figure 10.11 shows the number of coefficient updates required for varying values of the number of sweeps for five repetitions of the turbulence cycle. Figure 10.12 shows the convergence times for each setting. The residual reduction factor for CELS3D was set at 0.8.
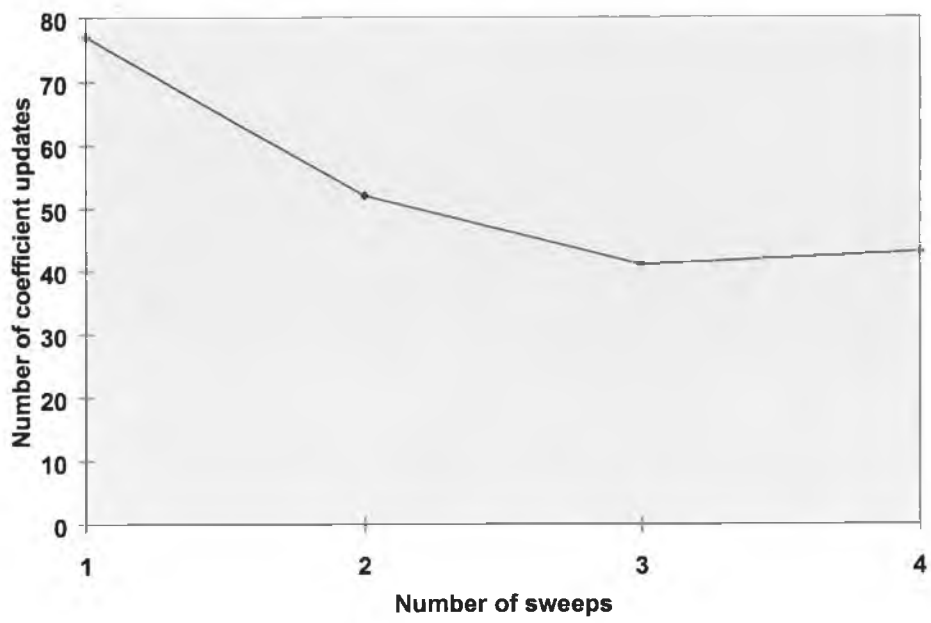
Figure 10.11. The variation of the number of coefficient updates with the number of sweeps.
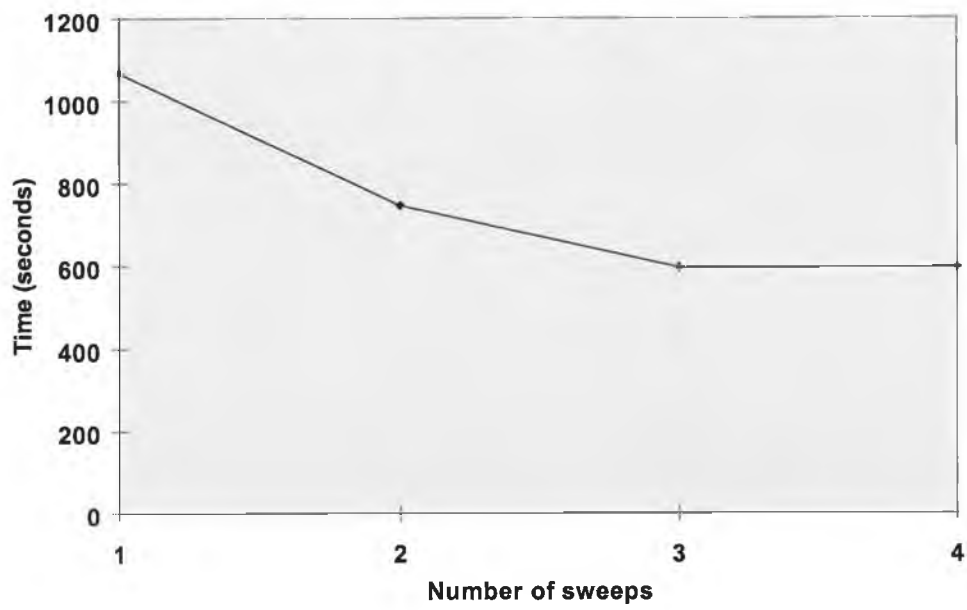
Figure 10.12. The variation of the time to convergence with the number of sweeps.

It is clear that greater efficiency is provided by more than one sweep of the domain. Three to four provides a good performance and greater than four in fact stopped the reduction of the residuals before meeting the convergence criterion and the residual values showed small oscillations and there was no further progress. This effect was removed by increasing the under-relaxation on the velocity equation coefficients but at the cost of increased computational effort for the solution procedure.

With a selection of 3 sweeps to minimise the number of coefficient updates and hence the potential for the application of ACM, the effect of the number of repetitions of the turbulence calculation cycle was investigated.

The repetitions were set at values from 1 to 6. The CELS residual reduction factor was again set at 0.8 and the under-relaxation of the velocity equations was 0.98. Figures 10.13 and 10.14 show the behaviour of the convergence time and the number of coefficient updates required for the different numbers of repetitions.
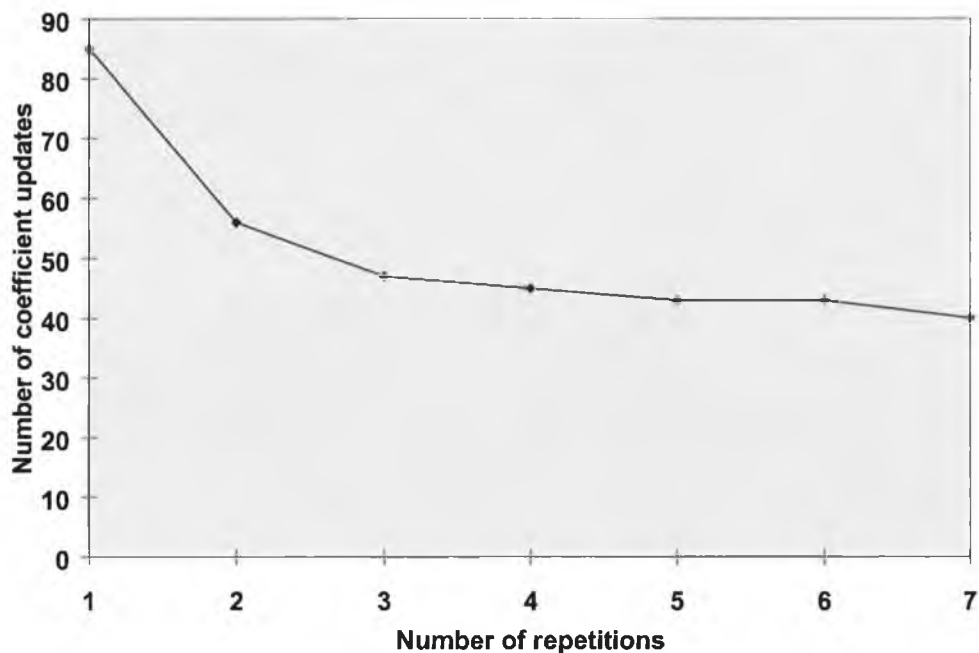
Figure 10.13. Variation of the number of coefficient updates with turbulence calculation cycles.

Minimising the number of coefficient updates requires at least three repetitions. The convergence time increased when the number of turbulence calculation cycles was set to be greater than three but the number of repetitions for the minimum value shown was not selected. In order to match the degree of convergence achieved for this problem using the three-dimensional TEACH code, the number of repetitions required was five.

Three sweeps and five repetitions were taken as the basis for further testing.
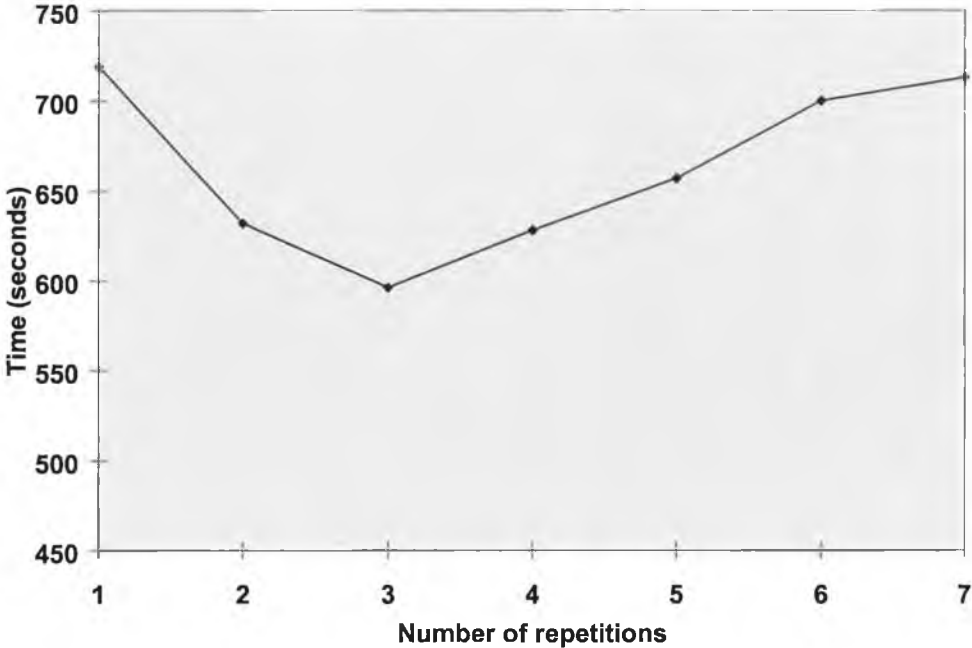
Figure 10.14. Convergence time versus the number of turbulence calculation cycles.

The general characteristics of the turbulent case were investigated in the same manner as the laminar one but full details are not reproduced as they will be re-examined after upgrading the TEACH code by implementing the QUICK differencing formulations.

It was found that the CELS3D under-relaxation had to be set to 0.7 or less to produce a convergent solution for the turbulent flow. While convergence was possible for discretised equation relaxation factors of 0.99 the results were not as predictable as for a value of 0.98. Although 0.99 gave the best convergence time, it was not taken as good practice and evaluation of the solvers performance relative to SIMPLE was made for 0.98.

Figure 10.15 shows the variation in the convergence time for a number of values of α.
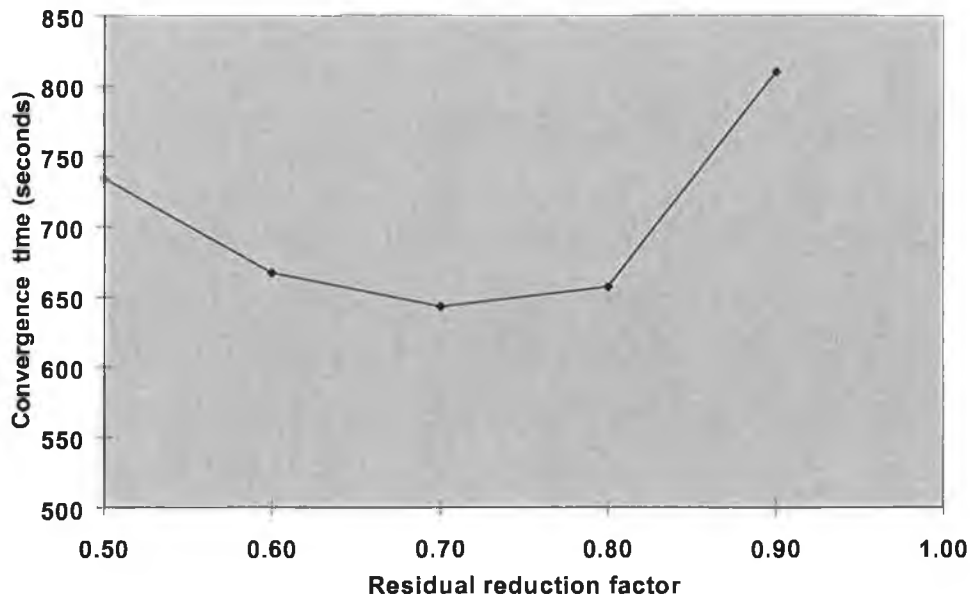
Figure 10.15. Variation in the convergence time with the residual reduction factor

Setting α at 0.7 gives the shortest convergence time and a low number of coefficient updates. For these settings the numbers of inner iterations is low. The maximum was just 5 and the average was 3. This leaves little scope for the application of an inner solver improvement such as the three dimensional form of ACM that is to be introduced. For these grid densities ($34 \times 34 \times 34$) then it is unlikely that ACM can be usefully applied to turbulent flows.
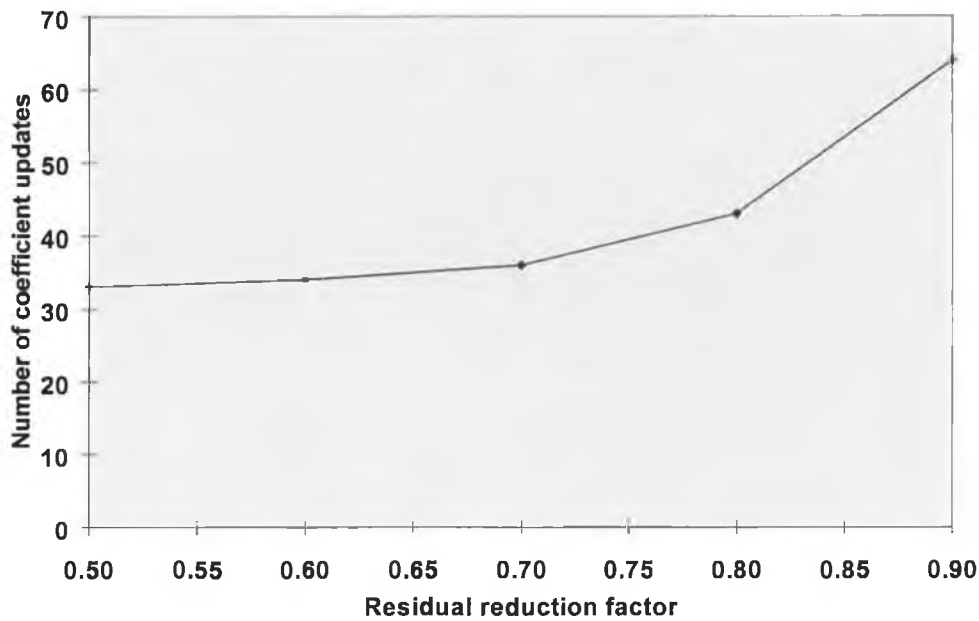
Figure 10.16. Variation in the number of coefficient updates with the residual reduction factor.

The number of initial conditioning cycles was set at twenty for the above tests and with the tuning set at stable values the requirement for this was examined. Unlike the laminar situation where convergence was possible and efficient for just one conditioning cycle, it was found that problems developed after up to five such cycles. The convergence stalled after a small number of coefficient updates.

Operation at seven conditioning cycles was judged to be sufficiently stable for present purposes and this arrangement gave the best convergence time at 602 seconds as illustrated in the following section.

10.4 SIMPLE vs. CELS3D - turbulent flow

A comparison of the convergence of SIMPLE and CELS3D is shown in figure 10.17. The gain in performance using CELS is good for this case. The convergence time for the CELS3D solver was just 28% of that for SIMPLE. The total number of coefficient updates was reduced from 627 for SIMPLE to 36.

While there is little scope for the application of ACM at these grid densities, if it were possible to make savings using ACM then the method could have some attraction. Working at increased grid density might offer some potential. For mushroom tunnel modelling at increased but still modest densities the potential is likely to be limited.

It should be noted that although the turbulence model was calculated and the inlet Reynolds number was 3500, when the flow solution was examined to calculate turbulent effects from the local values of the turbulence parameters it was found to be only weakly turbulent.   Higher values of the Reynolds number are discussed in Chapter 11 to ensure stronger coupling with the turbulence model.
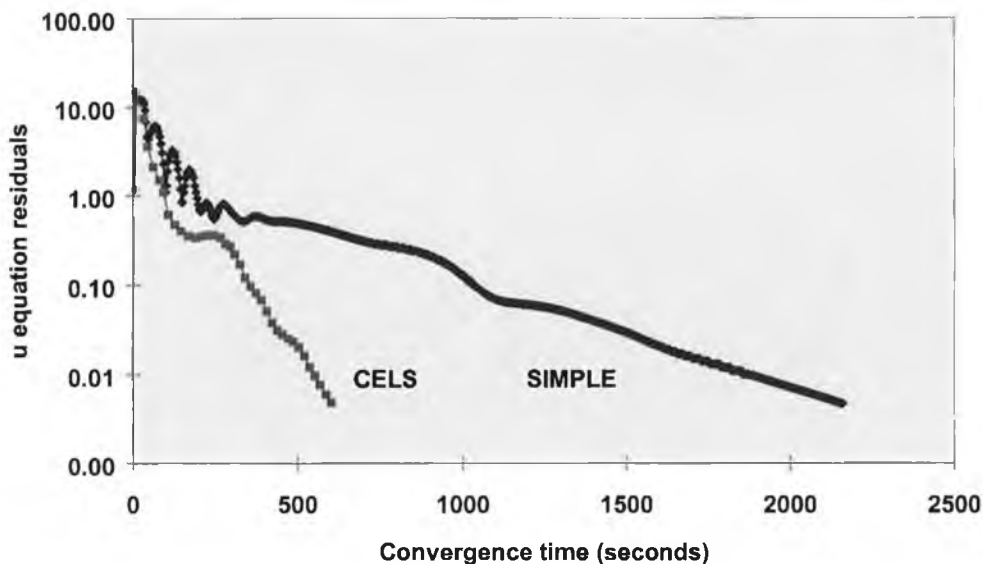
Figure 10.17. Comparison of CELS and SIMPLE for a turbulent flow (34×34×34).

210

Chapter 11

CELS3D with the QUICK differencing scheme

11.1 Accuracy of the difference scheme

Initially, in order to establish the general characteristics of the CELS3D solver, the hybrid differencing scheme was employed. This was done to limit complications in the upgrading of the TEACH code because initial experiments using QUICK formulations had not been successful but it was not clear whether this was due to a failure of the new solver code or whether it was a difficulty with the QUICK formulation. Having reverted to the hybrid scheme, it was possible to efficiently debug and test the CELS3D solver code but it then became necessary to upgrade the accuracy of the solver before further testing because of the well-known problems of numerical diffusion with first-order methods (e.g. de Vahl Davis and Mallinson, 1976 and Leonard, 1979). For turbulent flows and the study of models of turbulence, the presence of numerical diffusion in re-circulating regions means that no meaningful comparison of turbulence models can be made on coarse or moderately refined grids, especially in the case of high Reynolds number when physical diffusion is weak. De Vahl Davis and Mallinson showed that the false diffusion coefficients in steady flows were proportional to the magnitude of the fluid velocity vector, the grid size and dependent on the angle between the velocity vector and the grid lines. In principle, the grid can be refined to reduce the effects of numerical diffusion but, for three-dimensional flows in particular, the ability to do this is limited by computing resources. More accurate differencing schemes are thus highly desirable.

In view of its increasing popularity among higher-order differencing schemes the QUICK (Quadratic Upstream Interpolation for Convective Kinematics) method of Leonard (1979) was used in place of the hybrid scheme described earlier in this thesis (see Section 4.7).

As discussed in Section 4.7, there are problems associated with both central and upstream differencing techniques. According to Leonard (1978 and 1988), the problems with these techniques can be interpreted in terms of the method of estimating wall values of the dependent variables in a control volume formulation. On one hand, for central differencing the estimation is based on linear (or first-order) interpolation and while this involves only second-order truncation errors there are stability problems in the convective sensitivity (Patankar, 1980). On the other hand, for upstream differencing, the directional dependence gives very stable convective stability but the advantage of this is offset by first-order truncation error (e.g. De Vahl Davis and Mallinson, 1976). Grid refinement can alleviate the problems with these techniques but, as Leonard (1978) pointed out from a one-dimensional analysis, while the truncation error is formally proportional to the grid dimension, its actual size depends on the range of the cell-increment values of the dependent variable throughout the solution domain. Thus a small value of grid dimension does not by itself guarantee that the size of the first-order truncation order would be small.

It is necessary to construct an interpolation scheme that simultaneously possesses good accuracy and the directional properties associated with stable convective sensitivity. The three-point upstream-weighted quadratic interpolation for each wall value that is used in QUICK makes it possible to achieve a conservative formulation with stable convective sensitivity. The basic interpolation scheme is shown in figure 4.7 and an example of the formula used for interpolation is given as equation (4.34). As noted there, with the consistent quadratic interpolation used in modelling both the convective and diffusive terms, the overall truncation error in a solution obtained by the QUICK algorithm is third-order in the spatial grid size. Thus, it is possible to obtain highly accurate solutions on practical grid spacing.

In order to explore the improvements in accuracy, a series of tests were carried out that compared the results using the hybrid scheme and the QUICK scheme and grid independence of the results was investigated for a number of mesh sizes. In particular, for the flow illustrated in figure 11.7, it was demonstrated from calculations using the two schemes that grid independence of the solution is achieved at lower grid density with the QUICK scheme. For this testing, the SIMPLE solution procedure was used.

First accounts of the applications of the QUICK scheme described some difficulties in achieving convergence in all test cases (e.g. Patel and Markatos, 1986) but revised formulations were reported as giving more satisfactory performance. Of these later methods the "consistent formulation" (CFOR) of Hayase et al. (1992) and the "deferred correction" (DCOR) scheme of Li and Baldacchino (1995) were used in this thesis as well as the earlier "extended revised" (QUICKER) formulation of Pollard and Siu (1982). All three methods were coded for the modified TEACH program and third-order accurate approximations (see Section 4.10) (Hayase et al., 1992) were used throughout for those dependent variable references that fell outside the calculation domain.

The first formulation that was implemented was QUICKER and, when the method was coded as a Fortran routine and debugged, it was not possible, despite considerable effort in varying the relaxation factors and numbers of sweeps per iteration, to produce convergence for the two-dimensional square section test case of Section 9.1. In the course of investigating this difficulty it was decided to implement the scheme on a flow supplied as the example flow with the TEACH program. This approach minimised the additional coding required for a proven program and made it possible to focus on the immediate problem of the QUICKER formulation. The flow problem was a turbulent flow across a backward facing step (step duct expansion).

It was found that, to achieve convergence, the number of sweeps of the tri-diagonal solver per coefficient update had to be reduced to from 3 to 1 for the air velocity component and pressure equations. Reducing the number of sweeps in this way for the hybrid scheme was found to make the solver divergent. The flow calculated with the hybrid scheme and with QUICKER is shown in figure 11.1 The grid used was non-

uniform with an expansion factor in the x direction that was required to properly resolve the flow in the re-circulation region. The Reynolds number was 6,600 and the entering flow velocity was 0.5 ms^{-1} in the x direction. Boundary conditions were as outlined in Section 4.10. For clarity, the figures do not show the x-direction on a linear scale. The diagrams show the results as though they were on a uniform grid thus exaggerating but showing more clearly the re-circulation zone.

While it was found that differences between the results of the calculations with the two methods for this flow were small, at 190% of the computational effort for the hybrid case, the QUICK solution time was considerably longer.
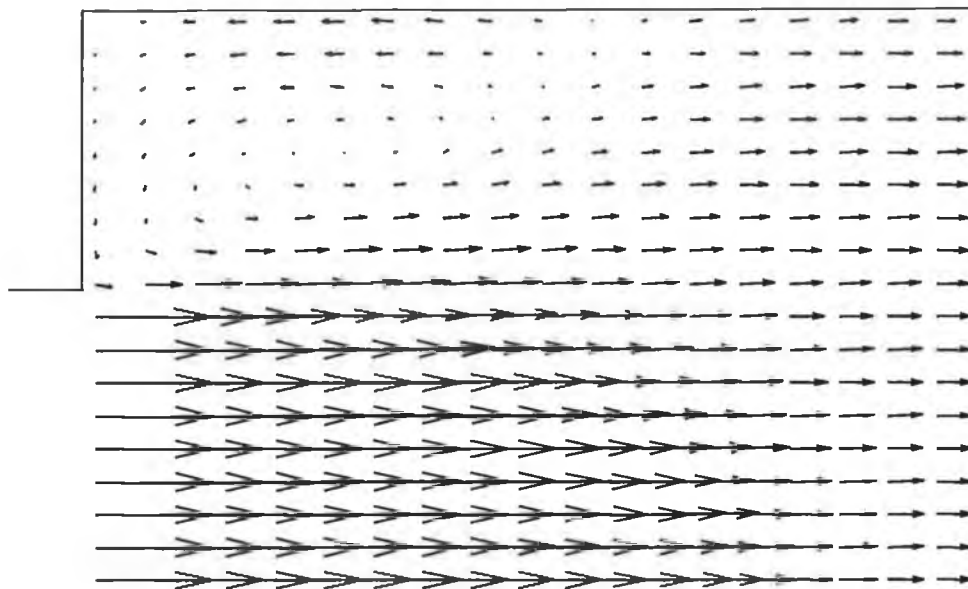
Figure 11.1. Duct flow past a step.

To examine the differences between the solutions calculated with hybrid differencing and those obtained using the QUICK scheme, the same flow problem was solved using a number of two-dimensional grid densities. This was done using the modified TEACH code so that the calculations also served as a check to ensure that no bugs that might affect QUICK had been introduced at this stage of the coding. The solver code was rewritten so that a succession of grids (10×10, 20×20, 30×30 and 40×40) could be used. The u-velocity components were calculated (using the SIMPLE method

and a residual target of 5 x 10^{-3}) for both the hybrid and the QUICKER differencing schemes.

In order to compare the results from the two schemes, the absolute difference between the two sets of results was calculated at each point in the grid and graphed as a three-dimensional surface chart for each grid density. Figures 11.2 to 11.5 show plots of the absolute differences for the series of grids. The differences are significant on the coarsest grid (figure 11.2) and, while decreasing on the next, are greatly reduced on the two densest grids. The relatively good agreement for a 30×30 grid is consistent with the results presented by, for example, Zhu and Rodi (1991).

The final accuracy of the methods can only be assessed by careful comparison with experimental data and so, while the issue of the proper resolution of the flow arises on the coarse grids and it is not clear how this would interact with the two difference schemes, it is possible to say that there are serious discrepancies between the two methods which diminish and are effectively removed by increasing the grid density. Most studies (e.g. Leonard, 1979 and Leschziner, 1980) have ascribed the differences to numerical diffusion in the hybrid scheme.
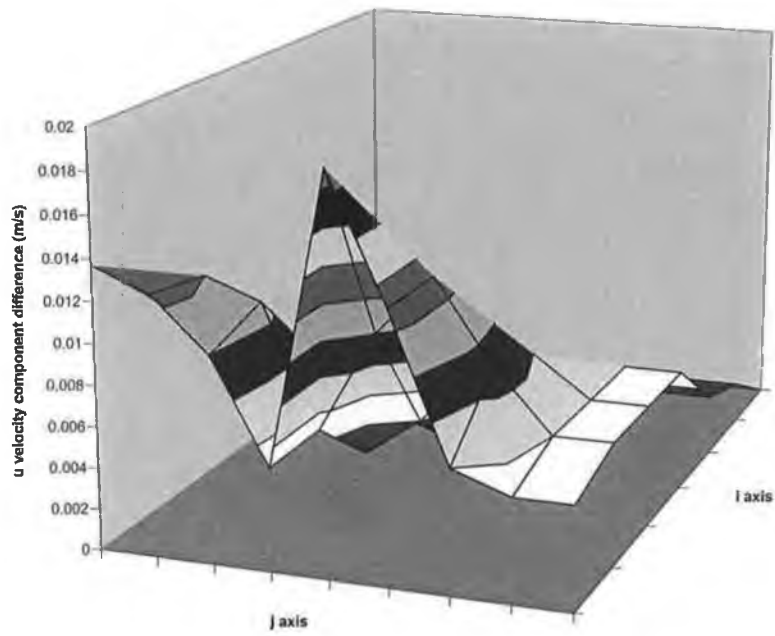
Figure 11.2. The differences beteween *u*-velocity components calculated using the hybrid and QUICK schemes for a sudden expansion duct flow (10×10 grid).
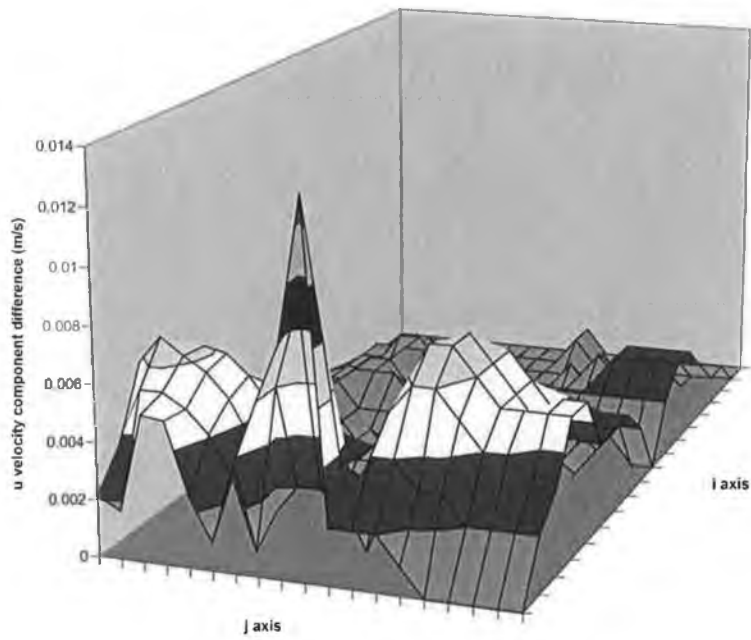
Figure 11.3. The differences between *u*-velocity components calculated using the hybrid and QUICK schemes for a sudden expansion duct flow (20×20 grid).
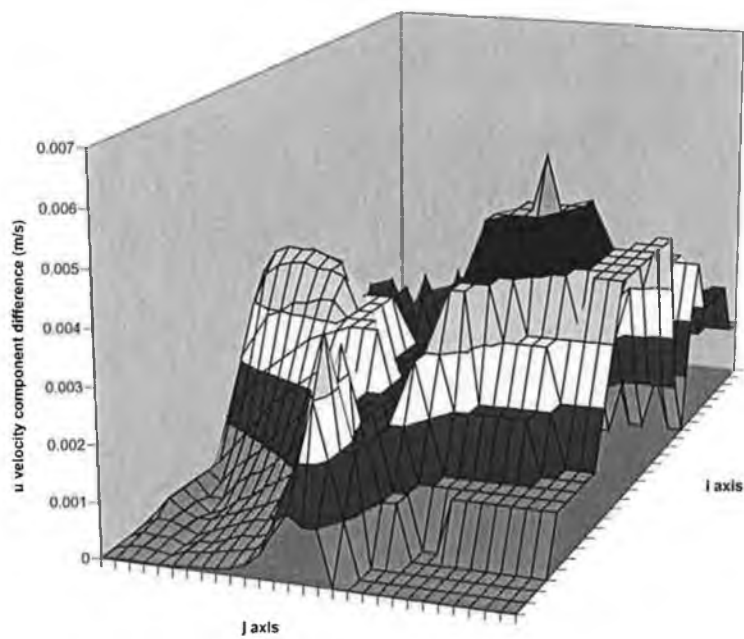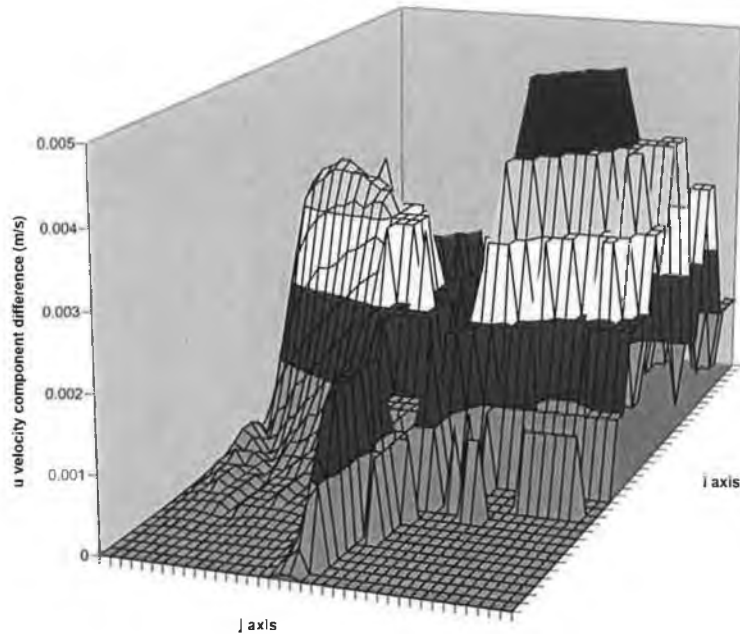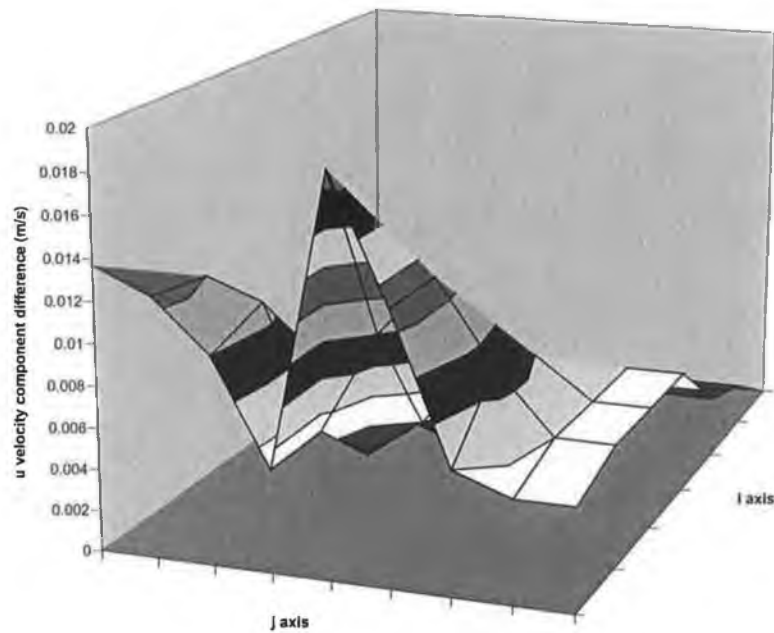

Figure 11.4. The differences between *u*-velocity components calculated using the hybrid and QUICK schemes for a sudden expansion duct flow (30×30 grid).

Figure 11.5. The differences between *u*-velocity components calculated using the hybrid and QUICK schemes for a sudden expansion duct flow (40 × 40 grid).

The convergence criterion was examined to ensure that changing to the QUICK schemes did not require an adjustment to this. Using a smaller residual target to improve the solution produced no significant change in these differences between the hybrid and the QUICK results, i.e. the discrepancy between the results is due to the change in the differencing scheme. For example, figure 11.6 contains the differences produced after solving the equations using a residual target of 5 x 10^{-4} and this was not found to be significantly different to the results in figure 11.2.

Figure 11.6. The differences between *u*-velocity components calculated using the hybrid and QUICK schemes for a sudden expansion duct flow (10×10 grid) using a residual target of 0.0005.

Having established the operating procedures that produced the satisfactory results for the duct flow, the testing for grid independence was resumed on the more relevant test case of Section 9.1 but calculations with the QUICKER formulation were proving troublesome because of long convergence times. Using the QUICKER method, convergence was only achieved for the original problem dimensions with a reduced inlet velocity and relaxation factors of 0.3 for the velocity equations and the computational effort relative to the hybrid scheme was found to increase very quickly as the Reynolds number increased. Patel and Markatos (1986) reported a failure to achieve convergence using QUICKER although it is not clear how much of this difficulty could have been due to a problem associated with the sweep direction for the line solutions. Typically, in the TEACH code for two dimensions, only one sweep direction is used. In this work it was found that using one sweep direction caused a difficulty in the application of QUICKER and this has already been noted by Leonard (1988).

219

Therefore, because of the increasing time taken to achieve convergence and as a check on the programming of the QUICKER formulation, the CFOR and DCOR formulations were also implemented and tested. Both of these methods include the QUICK terms in the source terms for the velocity equations. The intention of doing this was to achieve a stable numerical method based on the work of Khosla and Rubin (1974). The DCOR method was revised to achieve a more consistent implementation at the boundaries of the calculation domain because Li and Baldacchino (1995) reduced the DCOR formulation to first-order upwind (Patankar, 1980) for those grid cells adjacent to boundaries and it was found by the author that this approach produced results that differed markedly from those calculated using CFOR and QUICKER, particularly in the vicinity of slot inlets. Applying the third-order accurate boundary approximations of Leonard (1988, Hayase et al. 1992) with the DCOR formulation covering all the cells produced agreement between the QUICK methods. Li and Baldacchino noted that the accuracy of their results for a two-dimensional flow was lower than expected and this lack of consistency in their implementation of DCOR could explain this.

The results of the duct flow test case calculation for the 40×40 grid problem using CFOR and DCOR agreed to three significant figures and QUICKER was very similar with the same agreement over most of the flow field but it occasionally differed by less than 0.5% (one digit).

Table 11.1. Convergence times in seconds for various differencing schemes and formulations.

	Convergence time	
Formulation	**Re = 660**	**Re=6600**
Hybrid	115	46
QUICKER	183	116
CFOR	121	54
DCOR	120	53

Convergence times for the various formulations are shown in table 11.1 for two values of the Reynolds number. All the QUICK formulations required extra computational effort relative to the hybrid scheme but both CFOR and DCOR were considerably more efficient than QUICKER. The use of the low relaxation factors necessary for QUICKER made convergence very slow and therefore no further use of the formulation was made.
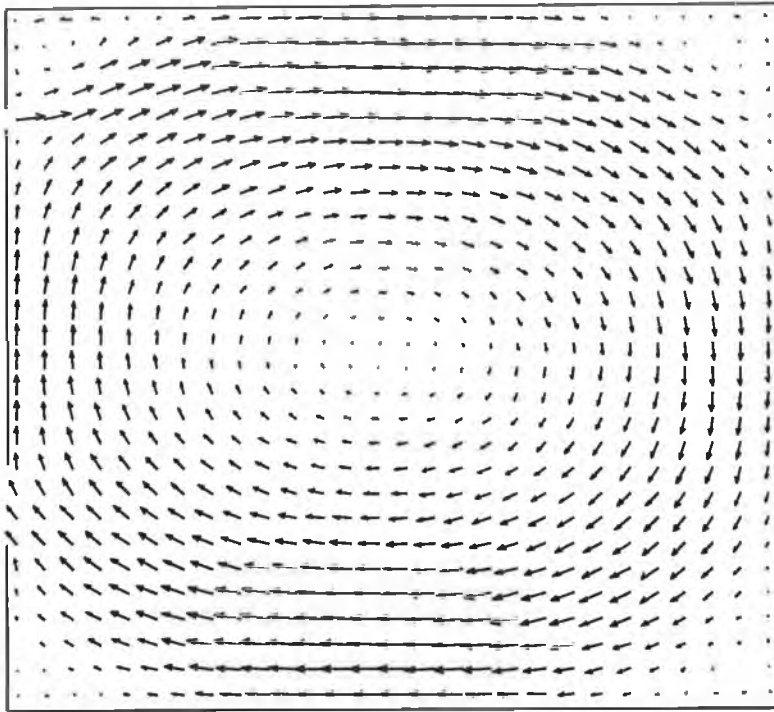

Figure 11.7. Two-dimensional flow field (Re=650).

Having established the new formulations, the flow field for the test case of Section 9.1 was calculated using the hybrid scheme and the DCOR QUICK formulation for three different values of the Reynolds number (figure 11.7).

This comparison was made between the flow fields, calculated on a range of grid densities, was made in order to investigate the requirements for the accurate solution of the flow problem. The purpose was to determine the magnitude of the errors that could be expected on the coarser grids and to determine the effect of the grid refinement on the reduction of these errors.

The grid sizes for this comparison were derived from the requirements of matching the geometry of the flow domain for each grid. The edges of an inlet, say, are located midway between two scalar grid points. In order to preserve the physical location of the inlet it is necessary to insert a minimum of two new grid points between the existing ones. Thus the inlet edge remains in the same physical location and midway between the two scalar grid points. On the coarsest grid there is one point located outside the boundary of the flow domain and after inserting the new points there would be two points located outside. To simplify the coding for the solver when using the finer grids it is convenient to operate with just one point outside the flow domain (as on the coarsest grid) and therefore the number of grid points in a given direction is two lower that what might be expected by inserting two points between each grid point.

Working in this way the flow problem was formulated on five grids. These were set and programmed as 30×30, 86×86, 142×142, 198×198 and 254×254. The flow problem was solved on each of these, for three values of the Reynolds number and using both the hybrid and the QUICK difference schemes. It was noted that the solution for Re = 650 showed very weak turbulence properties while the higher values produced significant turbulence effects. The solutions obtained were compared initially by determining the maximum difference between the solutions obtained on a given grid and the next finer one. The location of this maximum difference was also noted. This procedure is based on the proposals of Farrell et al. (2000) for a robust experimental approach to investigating the accuracy in a numerical procedure.

The results for the flow problem of Section 9.1 are presented in table 11.2 for the hybrid difference scheme and then table 11.3 for the QUICK scheme. In the table, the entries in the column for a given grid density are the differences between that grid and the next finer one. Thus, the first entry for the 30×30 grid at Re = 650 refers to the maximum difference between the solution for the 30×30 and that for the 86×86 grid.

Note that the grid references given for the finer grids are those points that correspond to the scalar grid points used on the 30×30 grid and that they are numbered as on the

coarsest grid for ease of comparison. When making such comparisons it is most important to be aware of the staggered grid arrangements. Thus, the u component of the velocity is solved for a point midway between the scalar grid point bearing the same co-ordinates and its nearest neighbour. This is quite a different physical location to the one associated with the corresponding grid point on the finest grid because the nearest neighbour is very much closer in that case. It is therefore important to interpolate the values of the u component of the velocity so that the values associated with a particular physical location are compared rather than a particular grid point co-ordinate. This, for instance, would be important in a corner where the flow direction is varying rapidly across the grid.

The first feature of note in the results is that the maximum difference is greatest when using QUICK on the coarsest grid. This is true for the three values of the Reynolds number that were used. The maximum differences decrease as the grid density increases and on the 142×142 grid the differences become very similar for both difference schemes. The maximum differences between the 198×198 and the 254×254 grids indicate that using the QUICK scheme results in an overall more accurate solution on these grids and that this is particularly true as the value of the Reynolds number increases.

For the hybrid difference scheme, on the coarsest grid there is no clear association between the maximum difference and the region of the flow field that is located near the inlet but, as the grid density increases there is an indication that the flow is most poorly resolved in the region of the inlet. The evidence is similar for the QUICK scheme but not as clear as in the hybrid case.

Table 11.2. **The maximum differences between flow fields calculated on a number of grid densities using the hybrid difference scheme.**

	Grid Density				
Re	**30×30**	**86×86**	**142×142**	**198×198**	**254×254**
650	0.0187	0.0086	0.0039	0.003	finest grid
	(5,26)	(5,27)	(3,29)	(4,26)	
6500	0.281	0.11	0.086	0.0386	finest grid
	(18,29)	(7,29)	(5,27)	(9,29)	
65000	2.98	1.16	0.89	0.6	finest grid
	(18,29)	(7,29)	(3,26)	(5,2)	

Table 11.3. **The maximum differences between flow fields calculated on a number of grid densities using the QUICK difference scheme.**

	Grid Density				
Re	**30×30**	**86×86**	**142×142**	**198×198**	**254×254**
650	0.0251	0.0111	0.0042	0.002	finest grid
	(12,27)	(22,29)	(3,29)	(20,29)	
6500	0.378	0.138	0.082	0.0127	finest grid
	(23,29)	(6,27)	(6,27)	(9,29)	
65000	4.61	1.56	0.86	0.227	finest grid
	(23,29)	(26,29)	(6,27)	(9,29)	

Note that while the maximum differences between the solutions obtained on the 30×30 and the 86×86 grids are lower for the hybrid case, the differences between the solutions for the 30×30 grid and the solutions on the finest grid at 254×254 are comparable. For example, at Re = 6500 the maximum difference between the finest and the coarsest grid is 0.407 for the hybrid scheme and 0.447 for the QUICK scheme. Thus, these tables provide information on the progress towards grid independence of the solutions as the grid is refined but may be misleading in terms of the overall accuracy of the solution on a given grid relative to the most numerically accurate one calculated on the finest grid. This point is made more explicit below in an analysis of the solution profiles on the coarser grids. Given this, it is interesting to note that the while the differences between the solutions obtained on the 30×30 and the 86×86 grids are more favourable for the hybrid scheme it is important to note that the situation tends to reverse for the finer grids and the QUICK scheme differences start to decline more rapidly than for the hybrid case. The QUICK scheme thus approaches grid independence more rapidly.

It is important to be aware that this analysis of accuracy, while it is a clear indicator of how accuracy improves as the grid is refined, can mask useful information about the difference schemes if the data is not analysed in detail. For example, the location of the maximum difference for the hybrid scheme at Re = 6500 and for the 198×198 grid is given as (5,2). This cannot be taken as an indication that the largest differences are no longer associated with the inlet because an examination of the data for the fine grids shows that there are differences of comparable but slightly lower magnitude associated with the inlet.

The maximum difference for the QUICK results in table 11.3, while it declines more rapidly than the hybrid case as the grid is refined, appears to be associated with the inlet. There is a small area of re-circulating flow just above the inlet (see figure 11.7) and the problem may be due to the adequate resolution of this area or it may lie with the calculation of the direct flow from the inlet. In order to check this, another set of tests was carried using the QUICK scheme out for a different inlet position and the results are presented in table 11.4. The inlet, a single grid unit in width on the coarsest grid, was moved to $j = 18$ and the outlet extended 3 grid units above $j = 5$.

Table 11.4. The maximum differences between flow fields calculated on a number of grid densities using the QUICK difference scheme.

			Grid Density		
Re	**30×30**	**86×86**	**142×142**	**198×198**	**254×254**
650	0.0235	0.00558	0.0028	0.00169	finest grid
	(7,18)	(3,16)	(4,16)	(4,16)	
6500	0.638	0.057	0.052	0.0365	finest grid
	(9,18)	(9,17)	(9,17)	(8,29)	
65000	7.3	0.6	0.37	0.24	finest grid
	(10,18)	(2,5)	(10,17)	(9,17)	

The overall pattern of the results in table 11.4 is very similar to those in table 11.3. It seems clear from these results that, while errors decrease on the finest grids, the maximum errors are mainly associated with the flow from the inlet and therefore in cases where the investigation centres on the inlet flow itself it may be useful to locally refine this area in order to ensure improved accuracy on the medium grid densities.

For the flow problem under discussion, then, there will be significant errors unless the solutions are carried out on a grid larger than 86×86. The usefulness of performing calculations on grids coarser than this will depend on the information that is required from a given solution. For example, the overall flow pattern is little altered for grids of 86×86 or greater.

Further testing of accuracy is carried out for the three-dimensional case in Section 11.2.

It is useful to analyse the effects of the difference scheme in more detail. The absolute values of the differences between the results from the hybrid and the QUICK schemes

for Re = 650 and for Re = 6500 on a 30×30 grid were plotted as before and the results are shown in figures 11.8 and 11.9. Comparing figure 11.8 with the duct flow results for a 30×30 grid in figure 11.4, it was clear that the discrepancies between the two differencing schemes were now considerably more pronounced. This would be expected, as the re-circulation for this flow is similar to that in the much-studied moving lid cavity flow where the angle between the flow direction and the gridlines varies throughout the flow domain and thus provides a strong source of numerical diffusion. Figure 11.9 shows results for the higher value of the Reynolds number (6500) and it can be seen that discrepancies between the results for the two schemes have increased for this case.
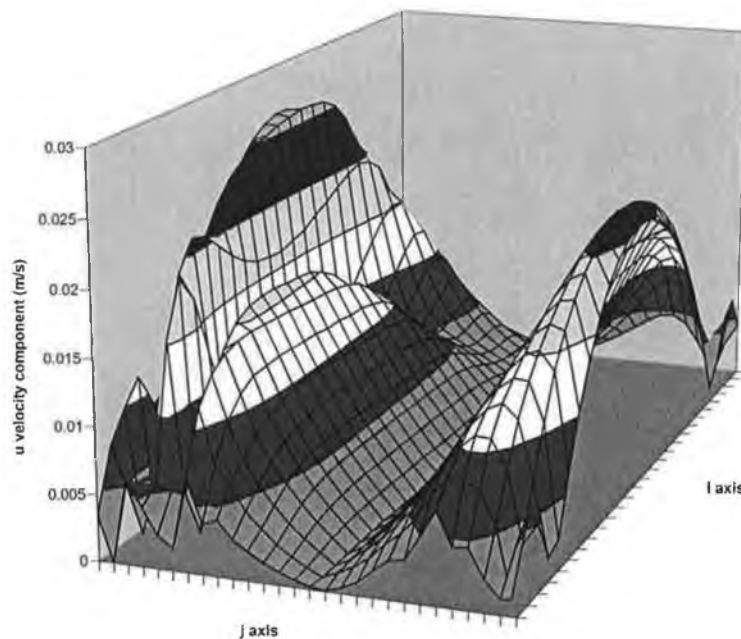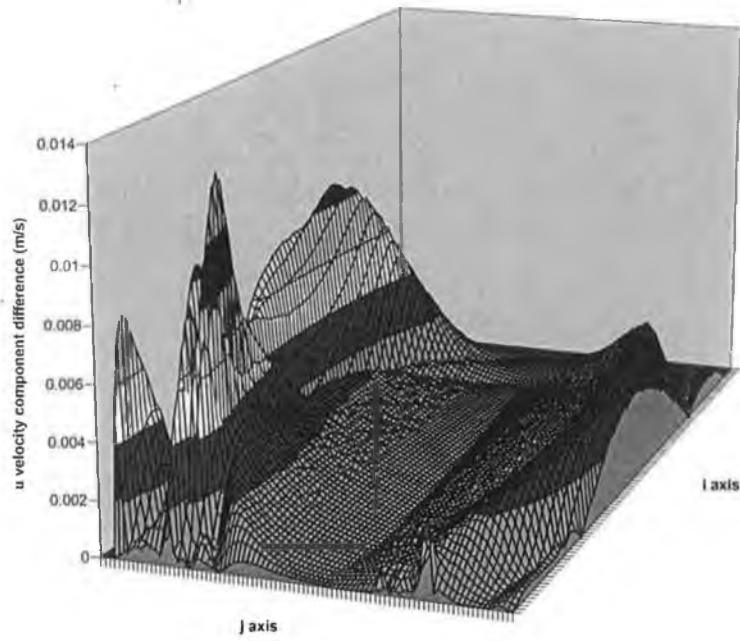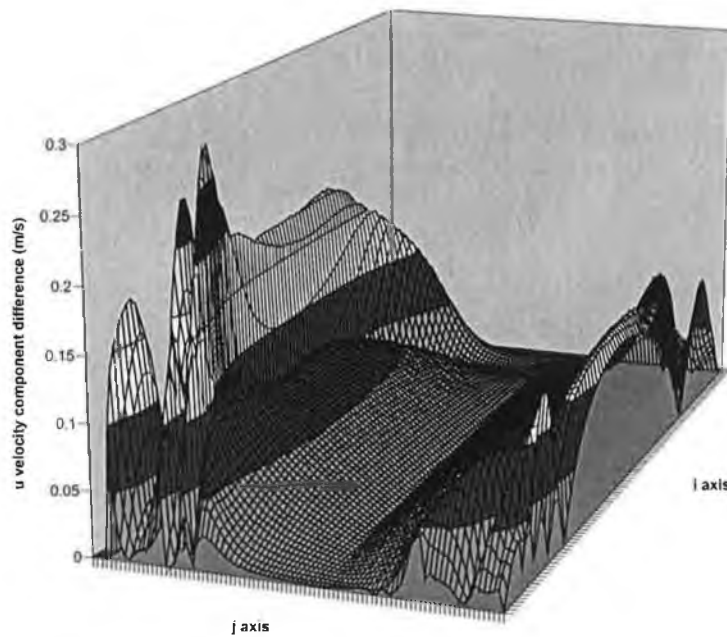

Figure 11.8. The differences between *u*-velocity components calculated using the hybrid and QUICK schemes for a slot inlet flow calculated on a 30×30 grid (Re = 650).

The solutions were calculated on higher grid densities were compared to test whether the hybrid method and QUICK solutions approached each other as the grid density increased. For the 86×86 grid, the results of the calculations for the same two values of the Reynolds number are shown in figures 11.10 and 11.11.
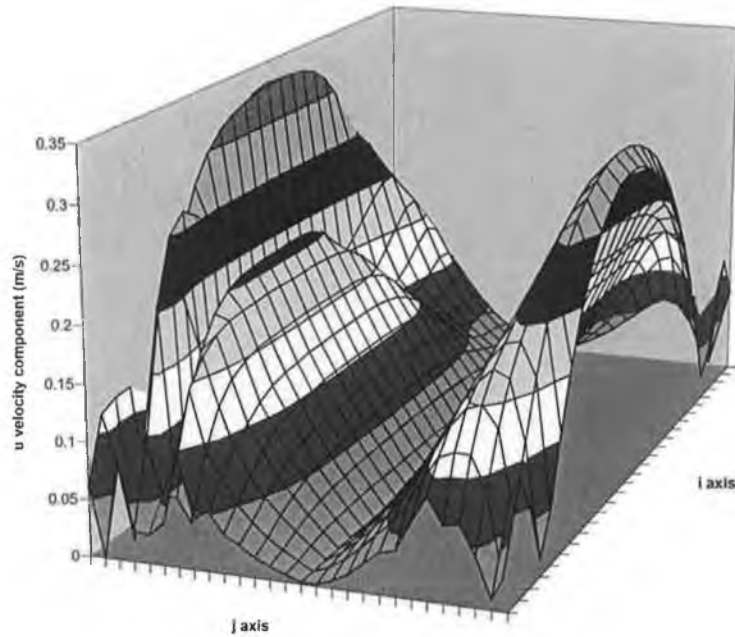
Figure 11.9. The differences between *u*-velocity components calculated using the hybrid and QUICK schemes for a slot inlet flow calculated on a 30×30 grid (Re = 6500).

Again, the discrepancies between the hybrid and QUICK methods were seen to decrease with increasing grid density but, even on the 86×86 grid, the differences are still significant. The lower Reynolds number showed reasonable agreement over most of the flow field but, for Re = 6500, the disagreement was more pronounced.

The equations for the turbulence parameters were removed from the solution sequence and the solutions for the higher Reynolds number of 6500 were recalculated. Figures 11.12 and 11.13 show the *u* velocity component differences for the 30×30 grid and the 86×86 grids respectively. The discrepancies are clearly much reduced and agreement is good for the denser grid.

Figure 11.10. The differences between *u*-velocity components calculated using the hybrid and QUICK schemes for a slot inlet flow calculated on a 86×86 grid (Re =650).

Figure 11.11. The differences between *u*-velocity components calculated using the hybrid and QUICK schemes for a slot inlet flow calculated on a 86×86 grid (Re =6500).

229
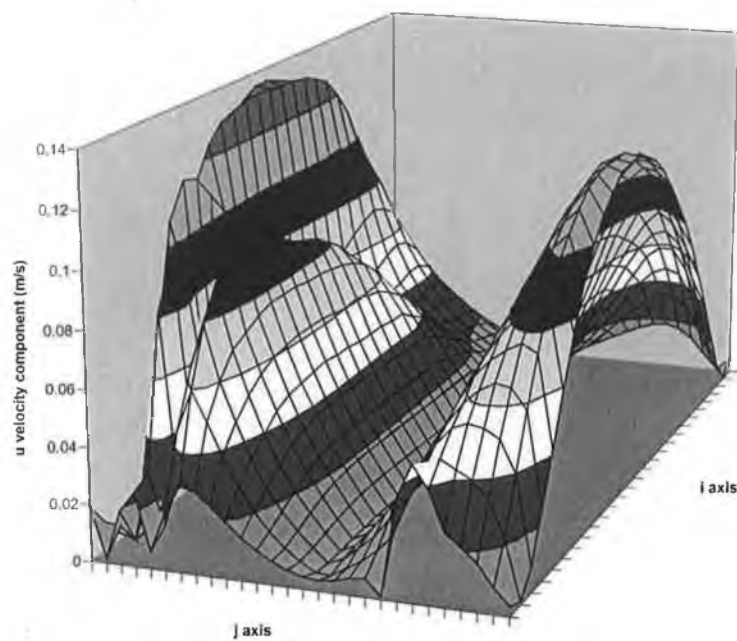
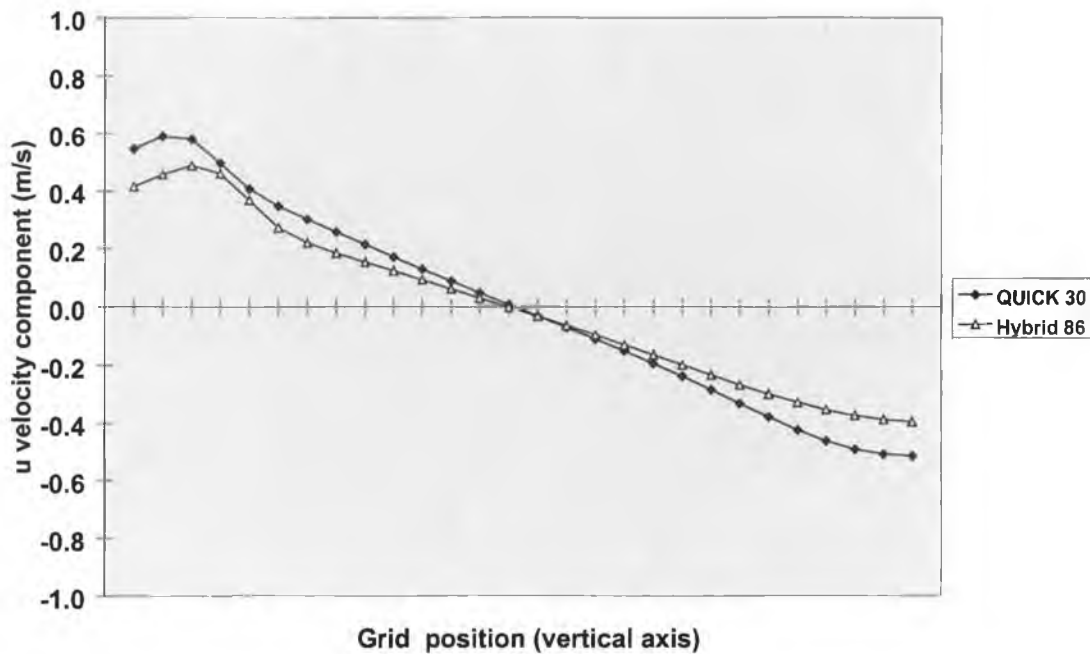**Figure 11.12. U-velocity component differences for a laminar slot inlet flow calculated on a 30×30 grid (Re =6500).**
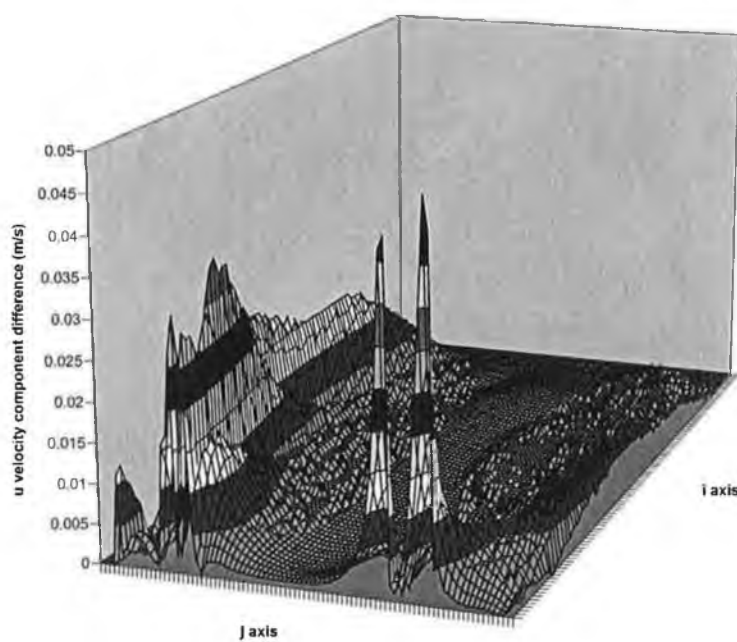
Figure 11.13. U-velocity component differences for a laminar slot inlet flow calculated on a 86×86 grid (Re =6500).

Figure 11.14. U velocity component profile for laminar flow (Re = 6500) on a 30×30 grid.

Deciding which of the differencing schemes is more likely to be the more accurate can be facilitated by comparison of the air velocity component profiles calculated by both methods. Figures 11.14 and 11.15 contain such a profile for the *u* air velocity component along a vertical line (constant *i* index) through the centre of the flow domain. The data on the left hand side of the graphs in figures 11.14 and 11.15 corresponds to the top wall of the flow as shown in figure 11.7 and the left hand side of the graphs in figures 11.12 and 11.13. Figure 11.14 is the profile calculated on the 30×30 grid and figure 11.15 is the result for the same flow geometry calculated on the higher resolution 86×86 grid. The differences between the results for the hybrid and QUICK schemes that are seen on the lower resolution grid are almost entirely removed on the finer grid. Comparison of these results shows that the velocity components calculated on the 30×30 grid for the QUICK scheme are almost identical

to those on the 86×86 grid with a slight under-estimate at the top wall on the coarser grid.
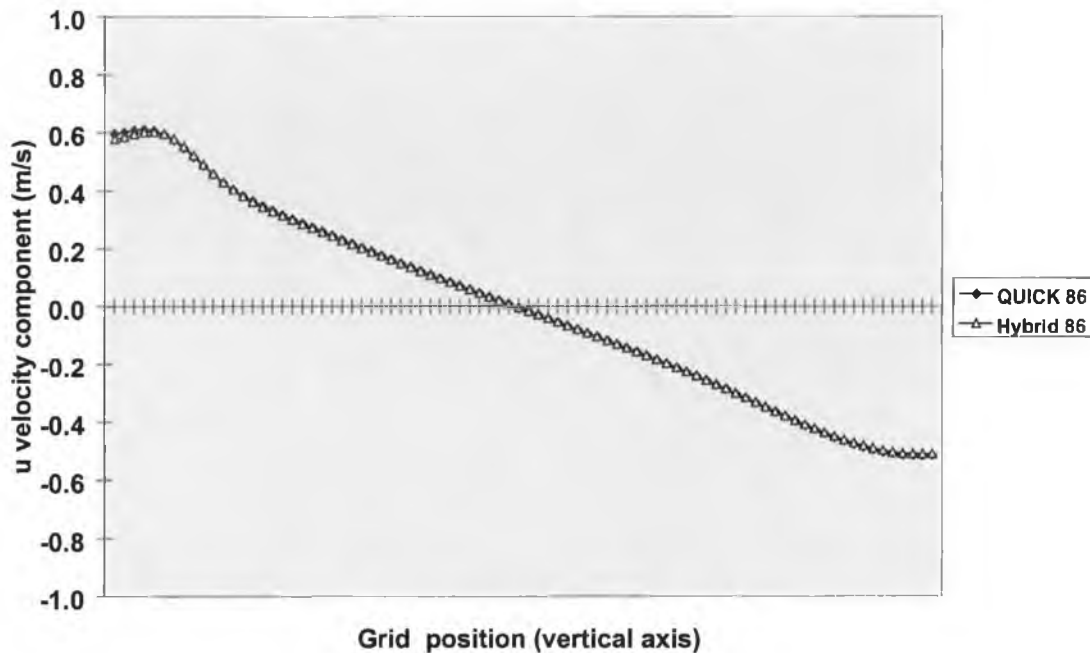
Figure 11.15. U velocity component profile for laminar flow (Re = 6500) on an 86×86 grid.

The results for the 86×86 grid are seen to be essentially method independent and, on this basis, the lower magnitude results from the hybrid scheme on the 30×30 grid can be said to under-estimate the magnitude of the velocity component. Thus, it is clear that the QUICK scheme provides accurate solutions on modest grid densities for the laminar flow.

The same type of comparisons can be made in examining the results for the same flow geometry with the incorporation of the turbulence calculations. The *u* air velocity component profiles can be plotted for the centre line of the flows and these are presented in figures 11.16 and 11.17 for the 30×30 and 86×86 grids, respectively.
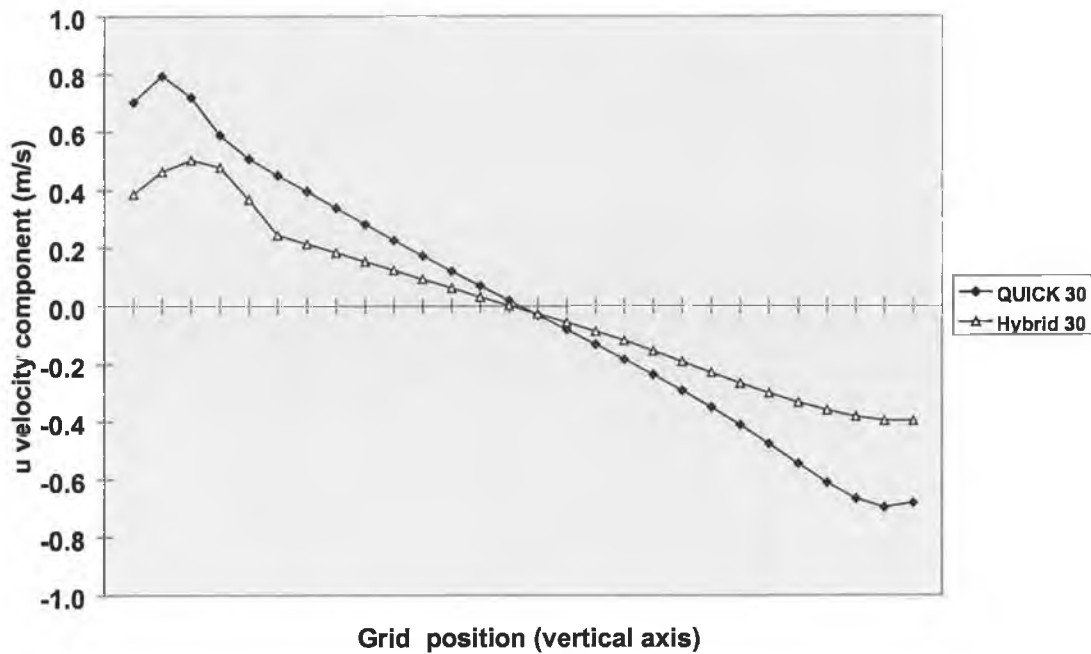
Figure 11.16. u velocity component profile for turbulent flow (Re = 6500) on a 30×30 grid.

It is clear that there is again an underestimate by using the hybrid scheme. However, in contrast to the laminar case, even on the 86×86 grid there are large discrepancies between the hybrid and QUICK-based predictions. Also, on the finer grid, there is a region of higher speed near the top wall that is predicted using the QUICK scheme and this is not as clear from the solution based on the hybrid scheme.
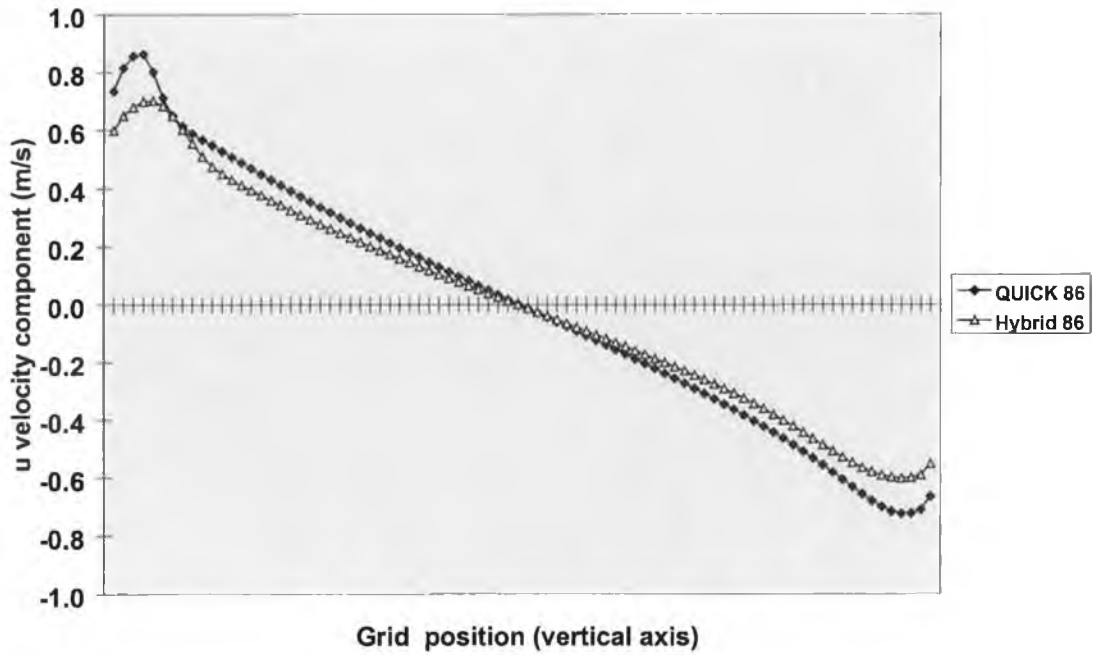
Figure 11.17. *u* velocity component profile for turbulent flow (Re = 6500) on an 86×86 grid.
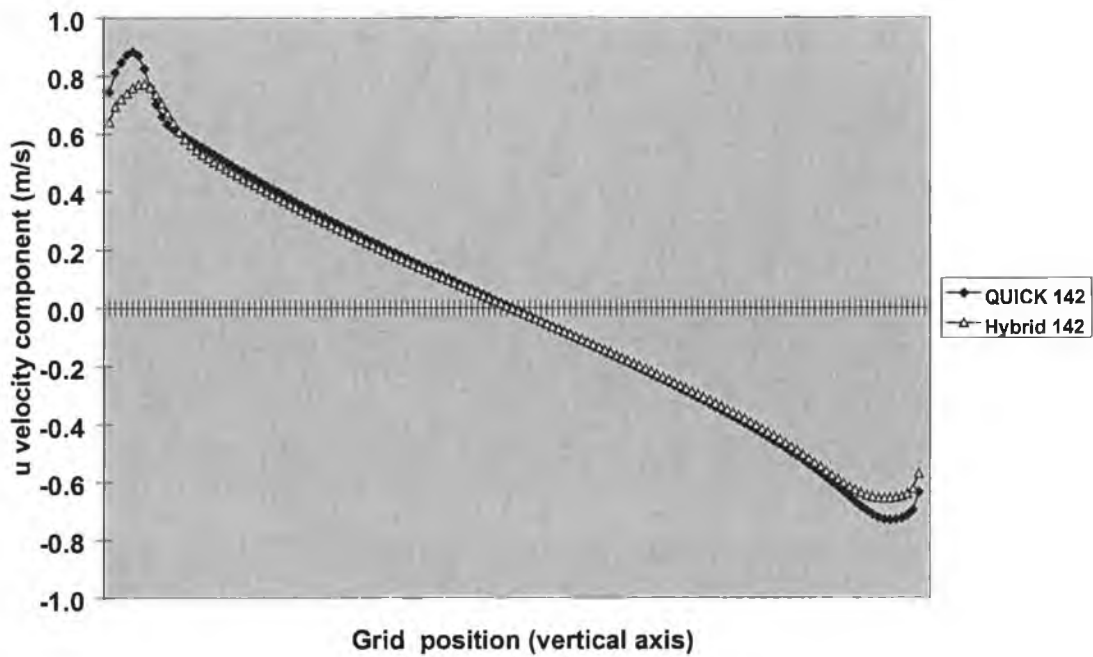
Figure 11.18. *u* velocity component profile for turbulent flow (Re = 6500) on a 142×142 grid.

In order to check that the discrepancies between solutions obtained using the hybrid and the QUICK schemes could be reduced as before, it was necessary to examine the solutions obtained on the 142×142 grid. At this resolution, the situation was much improved in that the two sets of results showed reasonable agreement as seen in figure 11.18. It should be noted that the region of higher speed flow near the top wall is now more apparent on the profile produced using the hybrid scheme.

The effect of grid density can be illustrated by plotting together the velocity component results for the three grids. The flow domains were constructed so that the 30×30 grid points map directly onto corresponding points for the 86×86 and 142×142 grids. The u air velocity components for these corresponding points were collected and are presented in figures 11.19 and 11.20 for the QUICK and hybrid schemes respectively. The results using QUICK are very similar for all three grids while those using the hybrid scheme are seen to approach the QUICK solution as the grid is refined.

The QUICK scheme provides results that are reasonably grid independent over a large part of the flow field at grid densities greater than 30×30 and this supports Leonard's assertion that accurate calculations should be possible on practical grid densities (Leonard, 1978). The errors using the hybrid scheme are such as to make it necessary to test for grid independence in the testing of, say, turbulence models because it would not otherwise be possible to attribute discrepancies in the results to the turbulence models alone. The high grid densities required would make such testing computationally intensive and expensive. The QUICK scheme, even with longer run times that the hybrid scheme at a given grid density, is more attractive for such work because the lower grid density required for accurate results would more than compensate for the extra computational effort per grid point.
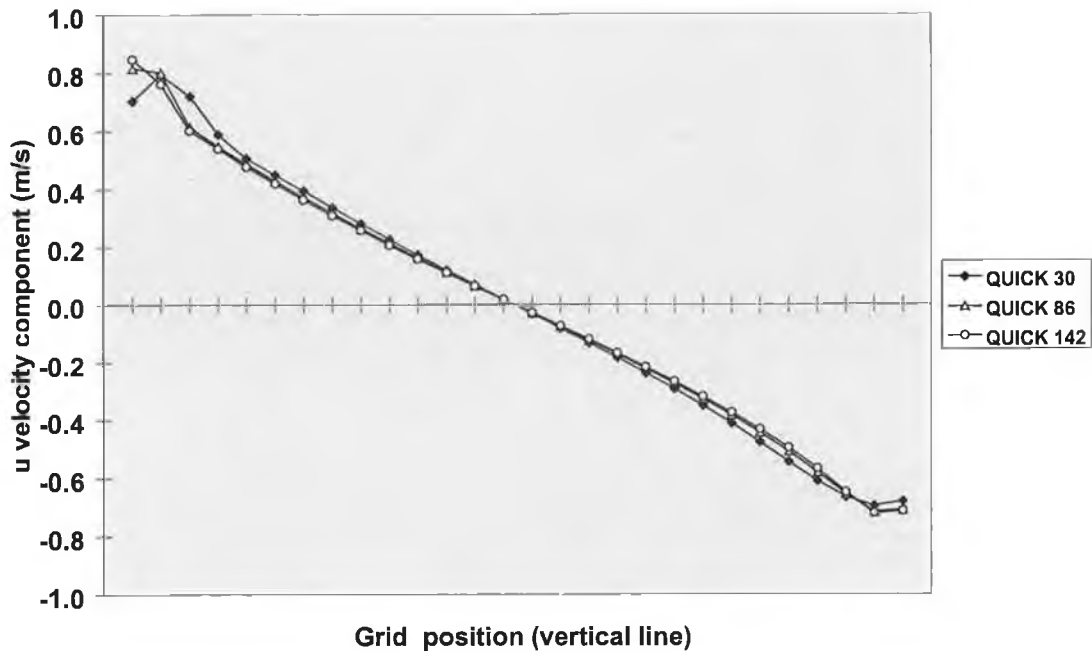
Figure 11.19. *u* velocity component profiles for turbulent flow (Re = 6500) on a number of grid densities using QUICK scheme.
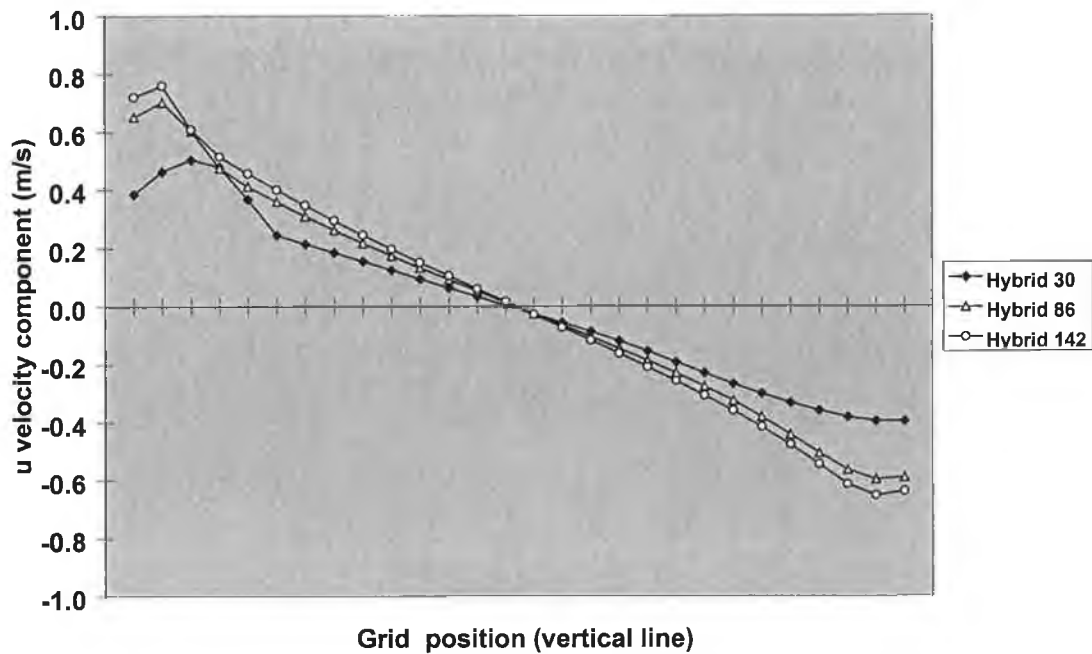

Figure 11.20. *u* velocity component profiles for turbulent flow (Re = 6500) on a number of grid densities using hybrid scheme.

To ensure that the boundary conditions did not play a significant role in producing the discrepancies in the sets of results produced when using the hybrid and the QUICK schemes, the flow due to the lower inlet position was also examined in more detail. The flow configuration is shown in figure 11.21.
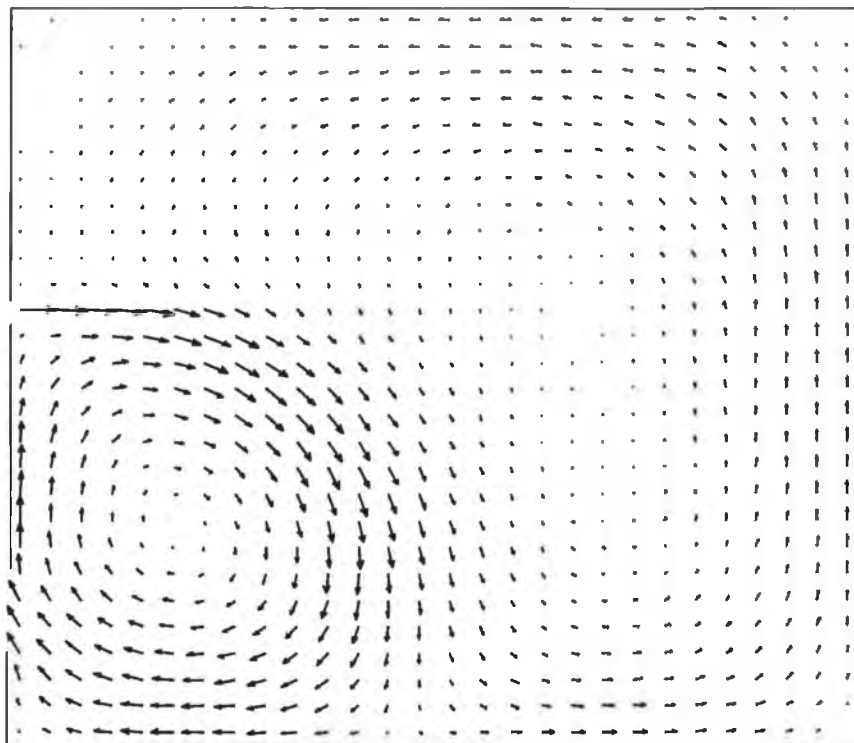

Figure 11.21. Two-dimensional flow field (Re=650).

The u air velocity component profiles for a vertical line through the centre of the main re-circulation zone (the lower left region of figure 11.21) were plotted and are presented in figures 11.22 and 11.23 for a flow with Reynolds number of 6500 for the 30×30 and 86×86 grids respectively. There are still considerable discrepancies on the coarser grid and, while agreement on the 86×86 grid is better than for the previous flow geometry, there are still significant discrepancies. This disagreement appears to be related to difficulty in resolving regions of high velocity gradient such as
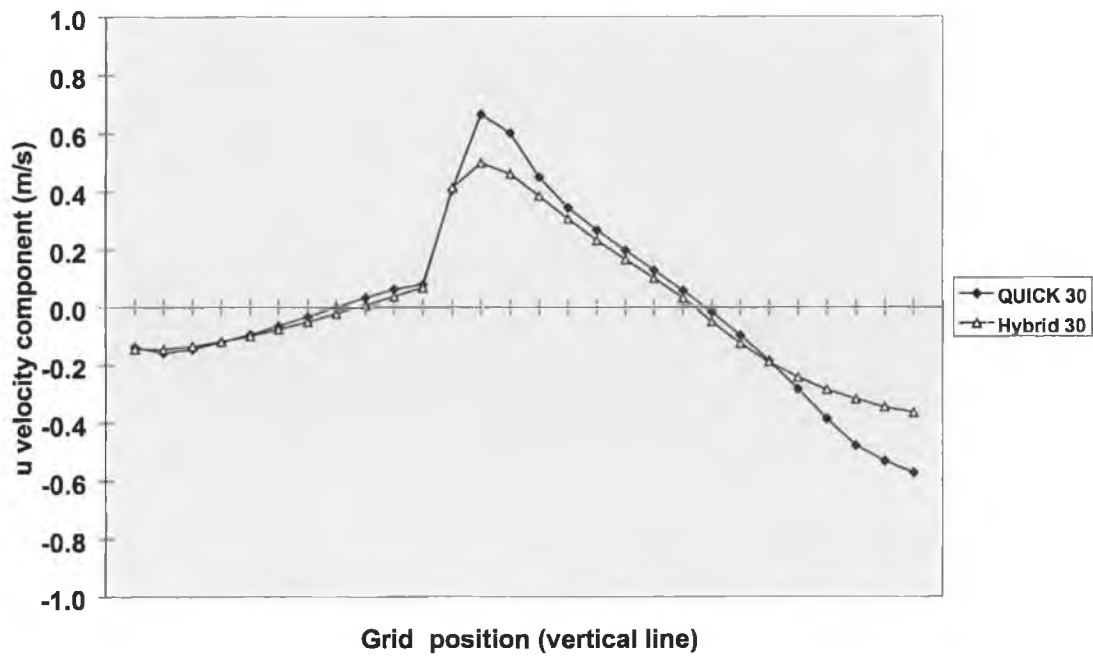
Figure 11.22. A *u* velocity component profile for turbulent flow (Re = 6000) on a 30×30 grid.

are associated with an inlet jet corresponding to one grid cell width rather than a characteristic of resolving wall flows.
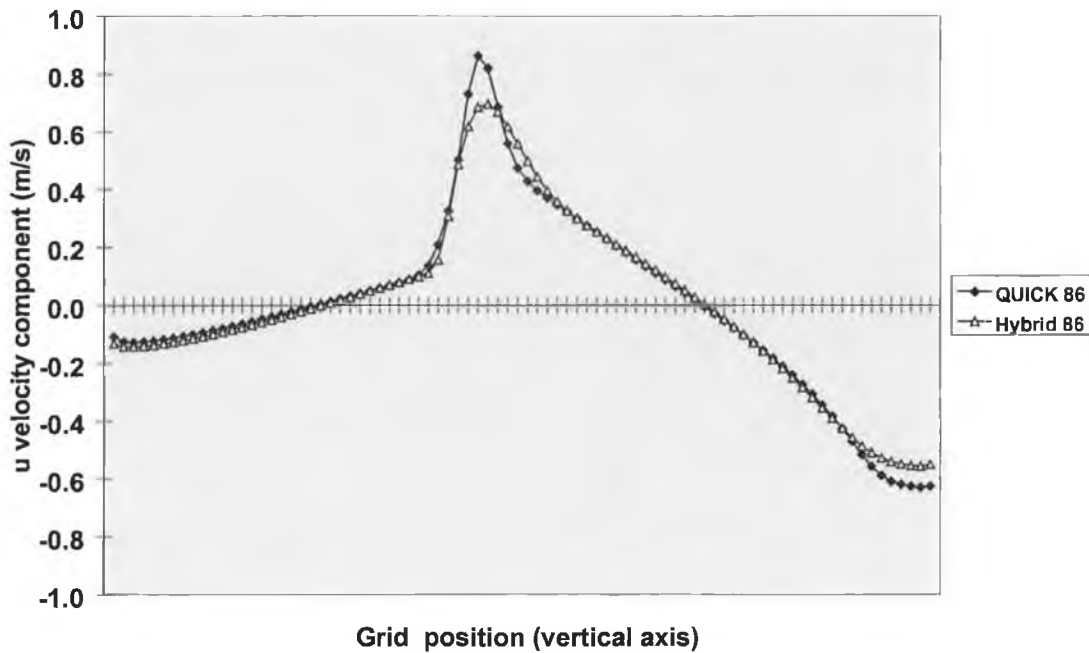
Figure 11.23. A _u_ velocity component profile for turbulent flow (Re = 6000) on an 86×86 grid

The hybrid scheme was found to be underestimating the _u_ velocity component so that even though a grid significantly denser than 30×30 would be required for grid independence of the QUICK scheme it was again clear that the use of the QUICK scheme is preferable in that it allows more accurate solutions overall on lower grid densities.

11.2 Accuracy of QUICK difference scheme for CELS3D

In order to promote rapid convergence and avoid very low relaxation factors for turbulence equations, as noted by Leonard (1988) in referring to the general application of QUICK, the difference code for QUICK was applied to the CELS3D solver code for the velocity component equations.

The memory limit of the compiler prevented a full examination of the accuracy of the three-dimensional solver (CELS3D) because it was not possible to reach grid-independent solutions with the range of grid sizes available. In order to demonstrate that the solver did in fact produce solutions that improved in accuracy at the grid was refined, the grid density in the z direction was fixed at a relatively coarse number and this enabled the grid density in the other two directions to be increased. In the three-dimensional test problem with the slotted inlet and outlet in the adjacent wall (see figure 10.1 in Section 10.1) it would be expected that, to some extent, the flow pattern would be dominated by a re-circulating flow due to the momentum from the inlet with a general flow towards the outlet in the third dimension. In this situation, the grid was refined in the plane of re-circulation due to the inlet and the grid size was fixed at 15 for the third dimension. Using 15 for the z direction facilitated the construction of three grids on which the geometry could be accurately matched.

As explained in Section 11.1, the grid sizes for this comparison were derived from the requirements of matching the geometry of the flow domain. The resulting grids were $8 \times 8 \times 15$, $21 \times 21 \times 15$ and $56 \times 56 \times 15$. The geometry of the flow problem was similar to that described in Section 10.1. The inlet was set to be 1 grid spacing in width in the y direction, 11 grid units long in the z-direction and was located at $j = 6$. The outlet was 1 grid unit high in the y direction, 1 grid unit wide in the x direction and was located at $i = 6$ and $j = 3$. The flow domain was 3 m square.

A series of calculations were made for the three grids and for a range of Reynolds Number (3500, 7000, 10000 and 20000). The maximum differences in the magnitude of the flow speeds between the calculated flow fields were determined and these and results are presented in table 11.5, using the same approach as that adopted in Section 11.1 (Farrell et al., 2000). The tables also contain the number coefficient updates required to obtain a solution for each calculation as an indicator of the computational effort required in each case. Overall, from the three cases examined in tables 11.5 to 11.7, it was determined that the relationship between the solution times t and the number of grid points N was of the form $t \propto N^{\beta}$ where β varied from a minimum of 1.36 to a maximum of 1.66 as the Reynolds number increased.

As before, when making such comparisons it is most important to be aware of the staggered grid arrangements and this is particularly true when the first grid is coarse. Thus, the u component of the velocity is solved for a point midway between the scalar grid point bearing the same co-ordinates and its nearest neighbour. This is quite a different physical location to the one associated with the corresponding grid point on the finest grid because the nearest neighbour is very much closer in that case. It is therefore important to interpolate the values of the velocity components so that the values associated with a particular physical location are compared rather than a particular grid point co-ordinate.

Also given in the table are the locations on the grids with which these maximum differences were associated. It should be noted that the grid references given for the finer grids are those points that correspond to the points used on the $8 \times 8 \times 15$ grid and that they are numbered as on the coarsest grid for ease of comparison.

QUICK differencing was used for the calculations in the three-dimensional case because the two-dimensional results showed that this method offered the best accuracy on fine grids and was less susceptible to the numerical diffusion associated with the hybrid difference scheme. The convergence criterion was set to 0.005 because the maximum differences recorded between solutions on the finest grids were found to be at least 10 times the maximum difference between these solutions and those for a convergence criterion of 0.0005 and the degree of convergence was sufficient to ensure that the grid comparisons were valid. The full set equations, including the turbulence transport equations, were solved in each case. The solutions were examined to determine whether or not the resultant flows were turbulent. Calculation of the turbulent viscosity from local values of the turbulence parameters showed marginal turbulence effects for Re = 3500 whereas it was concluded that the solutions for the higher values were turbulent.

Table 11.5. The maximum differences between flow fields calculated on a number of grid densities using the QUICK difference scheme.

		Grid Density			
	8×8×15		21×21×15		56×56×15
Re	Coeff. Updates	Max. diff.	Coeff. Updates	Max. diff.	Coeff. Updates
3500	50	0.0291 (5,7,6)	118	0.0122 (2,2,13)	175
7000	58	0.0589 (5,7,7)	180	0.022 (2,2,14)	280
10000	65	0.0888 (5,7,5)	212	0.04 (2,2,14)	341
20000	82	0.181 (5,7,5)	284	0.0773 (2,2,14)	608

It is clear from table 11.5 that, even over the relatively coarse grids used, the accuracy of the solution improved as the grid was refined.

One of the issues that arises is the modelling of the inlet and outlet. On the coarsest grid these are modelled as being one grid unit in width and one unit square respectively. On the next grid in the sequence (21×21×15) each one of these cells becomes three providing improved resolution. There is no clear evidence from the table that there are problems with the resolution of the flow at or near the inlet and outlet because no maximum error is recorded at or near those positions in the flow field.

The maximum differences between the two finer grids are not associated with either outlet or inlet. For the 56×56×15 grid there were 9 grid cells across the inlet and the outlet was 9 grid cells square, suggesting that 3 cells, in a uniform spacing

arrangement, may be sufficient to model the in and out flows. It may be the case, as in two dimensions (in Section 11.1) that an association with the inlet region would only emerge as the grids are further refined, i.e. the coarse grid differences are dominated by other resolution issues.

When the maximum differences were examined plane by plane (uv, vertical) to check for errors associated with the modelling of the inlet the maximum differences between the $21 \times 21 \times 15$ and the $56 \times 56 \times 15$ grids were all recorded midway across the ceiling of the enclosure and were not closely associated with the inlet. This is similar to the two dimensional case in Section 11.1 where large differences were noted at the top and bottom of the flow pattern. The same examination for the differences between the $8 \times 8 \times 15$ and $21 \times 21 \times 15$ grids showed little consistency in the positioning of the maximum difference on each plane and no association with the inlet.

For the differences between the $21 \times 21 \times 15$ and the $56 \times 56 \times 15$ grids all the differences recorded at the inlet were relatively small. Again this suggests that, for the uniform grids tested, that 3 grid cells is likely to be enough to accurately model the inlet flow in the context of other effects on these relatively coarse grids.

The results in table 11.5 are for a particular flow type with a particular geometry. While the type of flow is of interest for the application area of this thesis, it is important to examine what effects on accuracy of the solution may arise from variations in the geometry of the problem. Table 11.6 sets out results for a flow where all the details of the problem were the same as the one corresponding to table 11.5 but with a wider inlet. In this case the inlet was set to be 2 grid cells in width on the coarsest grid and was, therefore, 6 cells wide for the $21 \times 21 \times 15$ grid and 18 for the $56 \times 56 \times 15$. The inlet was located across $j = 5$ and $j = 6$.

Table 11.6. The maximum differences between flow fields calculated on a number of grid densities using the QUICK difference scheme (wide inlet).

		Grid Density			
	8×8×15		21×21×15		56×56×15
	Coeff.	Max.	Coeff.	Max.	Coeff.
Re	Updates	diff.	Updates	diff.	Updates
3500	50	0.023	88	0.008	181
		(5,3,14)		(5,3,14)	
7000	52	0.049	105	0.015	217
		(5,3,14)		(5,3,14)	
10000	52	0.079	109	0.031	264
		(5,3,14)		(4,2,2)	
20000	55	0.15	177	0.078	326
		(5,3,14)		(3,2,2)	

Again, the results in table 11.6 show that the accuracy of the solution improved as the grid was refined. It is of interest to note that the maximum differences recorded were very similar for all the calculations to those in table 11.5. In this case, the maximum differences between the solutions on the different grids were found to be associated with the outlet. The wider inlet appears to have made it more difficult to resolve the flow in the vicinity of the outlet.

When the maximum differences were examined plane by plane (uv, vertical) to check for any errors associated with the modelling of the inlet a number of the maximum differences between the $21 \times 21 \times 15$ and the $56 \times 56 \times 15$ grids were recorded midway across the ceiling of the enclosure and were not closely associated with the inlet. This was not as consistent as the previous case and some of the largest differences were located near the wall opposite to the inlet slot. As for the previous case, the same examination for the differences between the $8 \times 8 \times 15$ and $21 \times 21 \times 15$ grids showed

little consistency in the positioning of the maximum difference on each plane and no association with the inlet.

A further test was carried out with the original inlet and outlet dimensions (as for the flow corresponding to table 11.5) but, in this case, the inlet and the outlet were re-positioned to provide a change in the resultant flow pattern. This allowed a comparison that ensured that the choice of flow problem was not a fortuitously good or bad selection. The inlet was moved to $j = 5$ and the outlet was repositioned at $i = 6$ and $j = 5$. This would correspond to a high level return exhaust from the growing room.

Table 11.7. The maximum differences between flow fields calculated on a number of grid densities using the QUICK difference scheme (with alternative geometry).

	Grid Density				
	8×8×15		21×21×15		56×56×15
	Coeff.	Max.	Coeff.	Max.	Coeff.
Re	Updates	diff.	Updates	diff.	Updates
3500	64	0.0257	113	0.0138	211
		(5,7,4)		(3,2,2)	
7000	57	0.0563	149	0.016	300
		(4,7,13)		(2,7,3)	
10000	56	0.0887	171	0.0429	442
		(5,7,14)		(2,2,3)	
20000	56	0.185	194	0.106	358
		(5,7,14)		(2,2,3)	

The results presented in table 11.7 again show that the accuracy improves as the grid is refined and maximum differences recorded are similar to those in tables 11.6 and

11.7. There is a consistency in the accuracy of the solver when the flow geometry is adjusted.

For this particular example, the association of the maximum differences with the outlet that was seen in the for the wider inlet is no longer present and this is encouraging for the modelling of the narrow inlets (slotted or perforated ducts) that are associated with this application. It was also noted that when, as before, the maximum differences were examined plane by plane (uv, vertical) to check for the location of larger differences, the majority of the maximum differences between the $21 \times 21 \times 15$ and the $56 \times 56 \times 15$ grids were recorded midway across the ceiling of the enclosure and were not closely associated with the inlet. This was more consistent than the previous case and very similar to the flow for table 11.5.

As for both of the previous cases, the same examination for the differences between the $8 \times 8 \times 15$ and $21 \times 21 \times 15$ grids showed little consistency in the positioning of the maximum difference on each plane and no association with the inlet.

11.3 Characteristics of CELS3D with QUICK differencing

In Chapter 10 an initial examination of the convergence behaviour of CELS3D was made using one flow geometry and one value of the Reynolds number, with the hybrid differencing scheme. Further testing, required to reach more general conclusions on the behaviour of CELS3D, was undertaken after the examination of the use of the QUICK differencing scheme in the Section 11.1. Having successfully implemented the QUICK differencing, a set of tests was carried out to determine the more general characteristics of the behaviour of CELS3D, in terms of the efficiency of the convergence relative to the performance of the SIMPLE procedure. Two flow configurations without two-dimensional symmetries were used to allow an examination of the effects of the change in geometry and three values of the Reynolds number were used in each case.

11.4 First test case

The test case of Section 10.1, i.e. the approximation to the air flow in a square-section mushroom growing room, was used as the first test case for the application of CELS3D using QUICK differencing. The boundary conditions were identical to those used in that section. The results below were calculated using the formulation of Li and Baldacchino (1995). Third-order boundary treatment was used to take account of the reference to grid points outside the flow or solution domain. The formulation of Hayase et al. (1992) produced identical solutions with a slight advantage (approximately 5%) in convergence time. Testing was carried out for three values of the Reynolds number but full results in graphical form are included only for the lowest value. In all cases considered, the best convergence time was compared with the corresponding result using SIMPLE.

For a Reynolds number of 3500, the parameters controlling the convergence (relaxation, residual reduction, etc.) were tuned for a good convergence time and then they were varied in turn about these near-optimum values to obtain information about their effect on the progress of the solution.
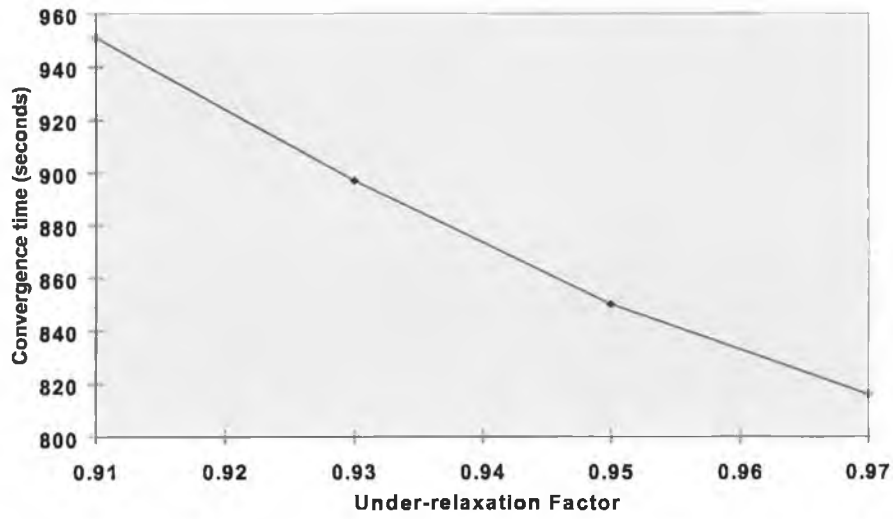
Figure 11.24. Variation of the convergence time with the under-relaxation factor for the turbulence parameter equations.
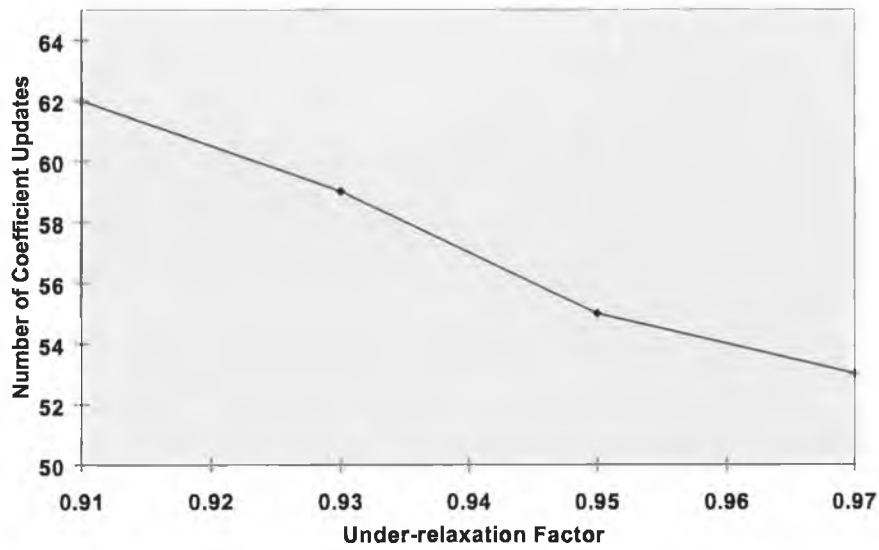

Figure 11.25. Variation of the number of coefficient updates with the under-relaxation factor for the turbulence parameter equations.
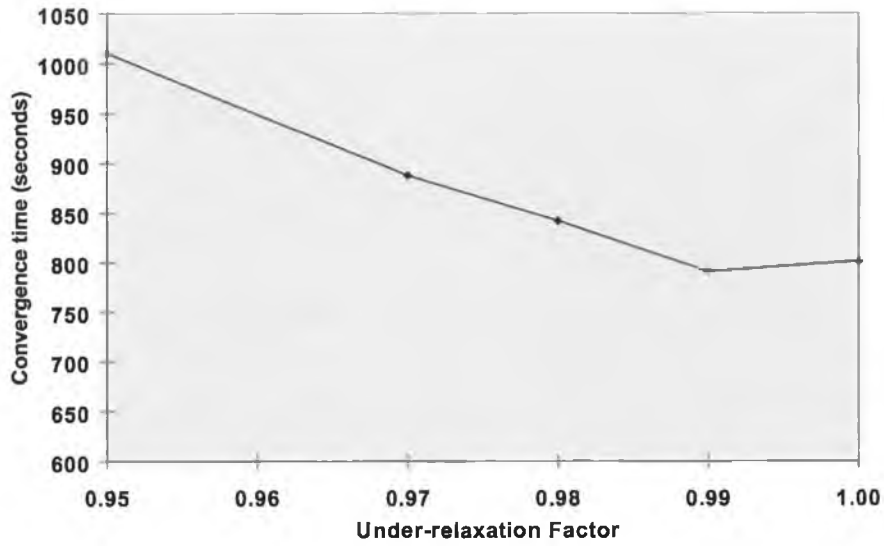
Figure 11.26. Variation of the convergence time with the under-relaxation factor for the velocity component equations.
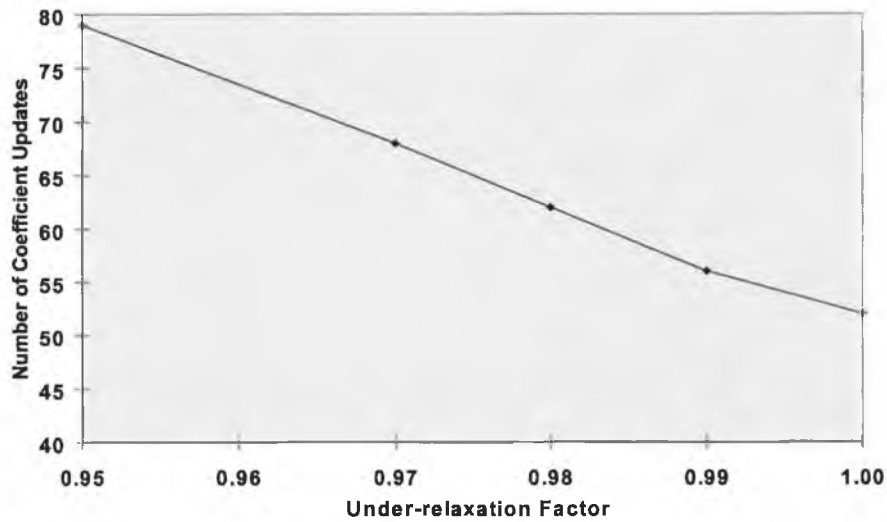

Figure 11.27. Variation of the number of coefficient updates with the under-relaxation factor for the velocity component equations.

249

The first factors examined were the under-relaxation factors for the turbulence and air velocity equations. Figure 11.24 shows the variation of convergence time with relaxation factor for the turbulence parameter equations and figure 11.25 shows the variation in the number of coefficient updates. Both of these dependent variables show a decrease as the relaxation factor is increased toward a value of 1 (no relaxation). However, the value was limited to 0.97 as the solver diverged for higher values and a value of 0.97 was selected for further testing.

Under-relaxation of the velocity component equations was addressed next. Figures 11.26 and 11.27 show the results of varying this factor. The best convergence time was given by a value of 0.99 and, in spite of a small drop in the number of coefficient updates for a value of 1, this value was chosen for further testing.
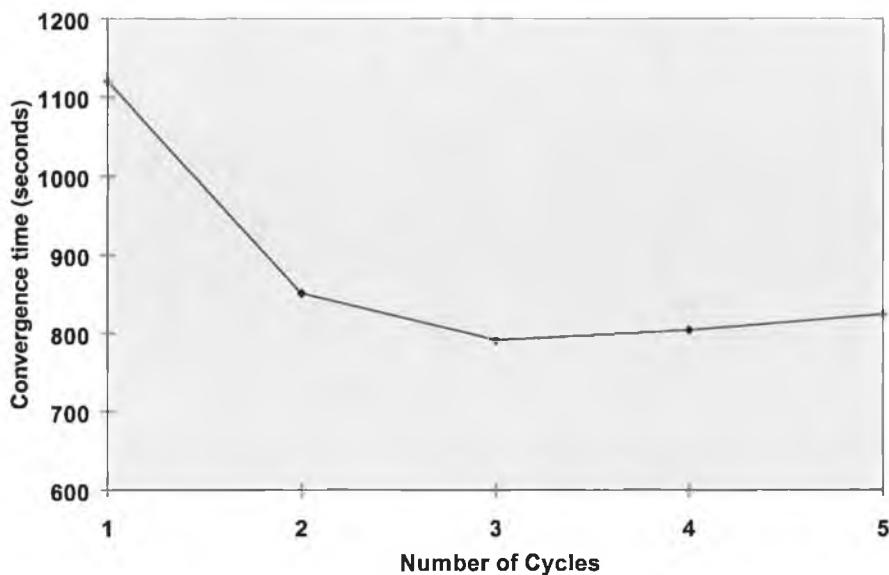

Figure 11.28. Variation of the convergence time with the number of turbulence calculation cycles.
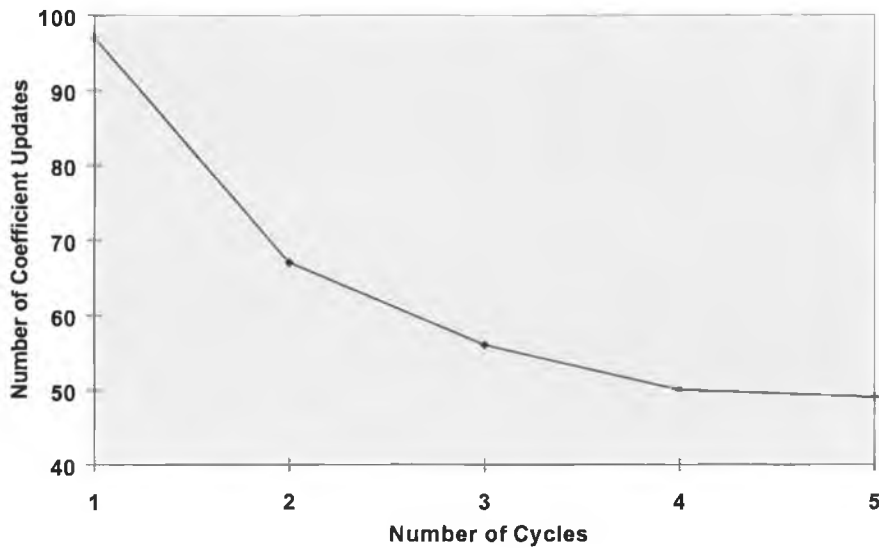
Figure 11.29. Variation of the number of coefficient updates with the number of turbulence calculation cycles.

The behaviour of the turbulence calculation procedure was tested by varying both the number of turbulence calculation cycles and the number of sweeps of each equation in the tri-diagonal line solution procedure. Figures 11.28 and 11.29 show how the convergence time and the number of coefficient updates varied as the number of calculation cycles was increased. The shortest time was provided by a value of three and this was chosen for further testing despite the continuing drop in the number of coefficient updates as the number of cycles increased.

Figure 11.30 and 11.31 show the same sort of data for different numbers of sweeps. Convergence time was minimised by 6 sweeps. As above, the number of coefficient updates continued to show a downward trend as the number of sweeps increased but convergence time increased above 6 sweeps. As the number of sweeps had initially been set at 3, the testing of turbulence cycle number was repeated but the results still set this number at 3.
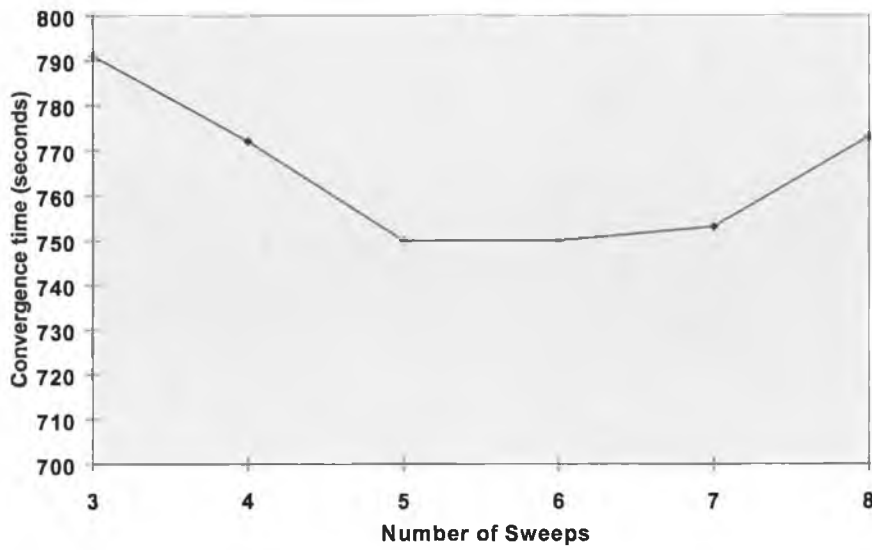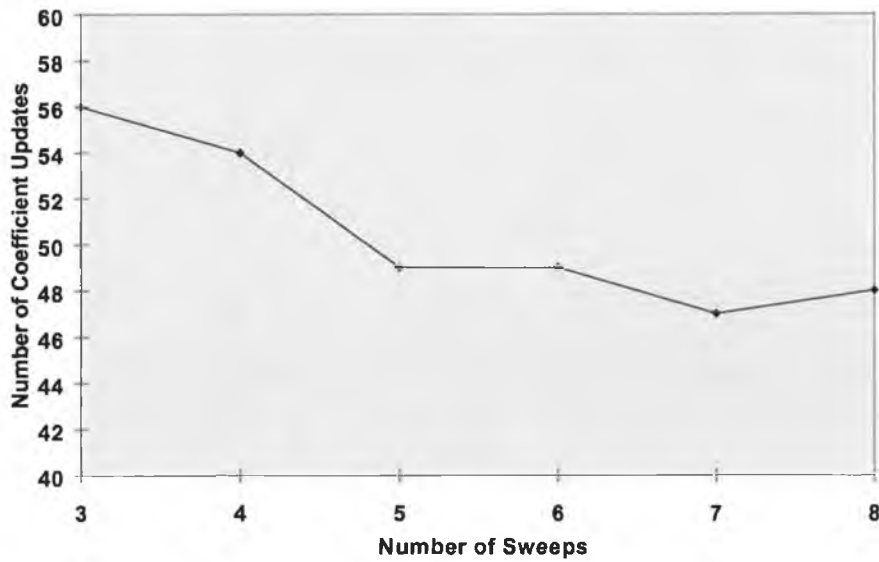
251

Figure 11.30. Variation of the number of coefficient updates with the number of sweeps of the turbulence parameter equations.

Figure 11.31. Variation of the number of coefficient updates with the number of sweeps of the turbulence parameter equations.
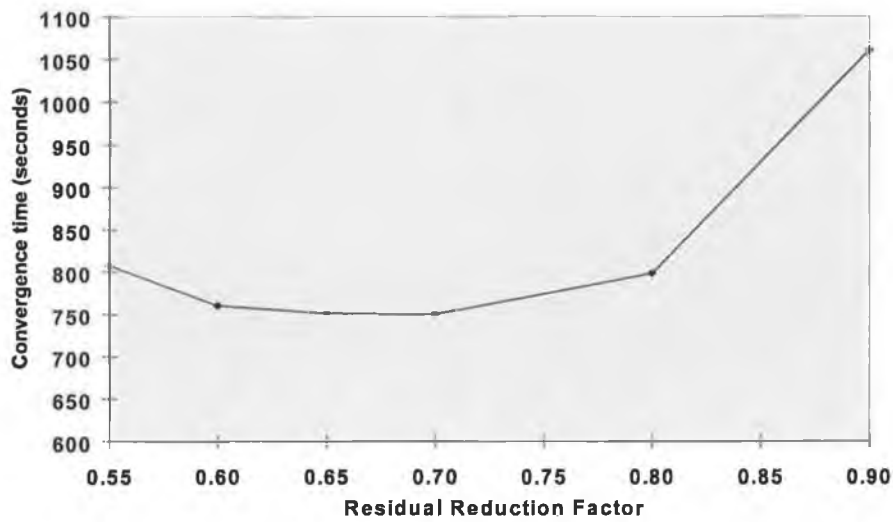
Figure 11.32. Variation of the convergence time with the residual reduction factor for CELS3D.

Testing was then focused on the CELS3D routines and the residual reduction factor was varied to determine its optimum value for this problem (figures 11.32 and 11.33). It can be seen from the results that the best combination of convergence time and number of coefficient updates was given by a value of 0.65.

Finally, the relaxation factor for the CELS3D routines was varied and produced the results presented in figures 11.34 and 11.35. A relaxation factor of 0.8 was selected as the best for this problem.
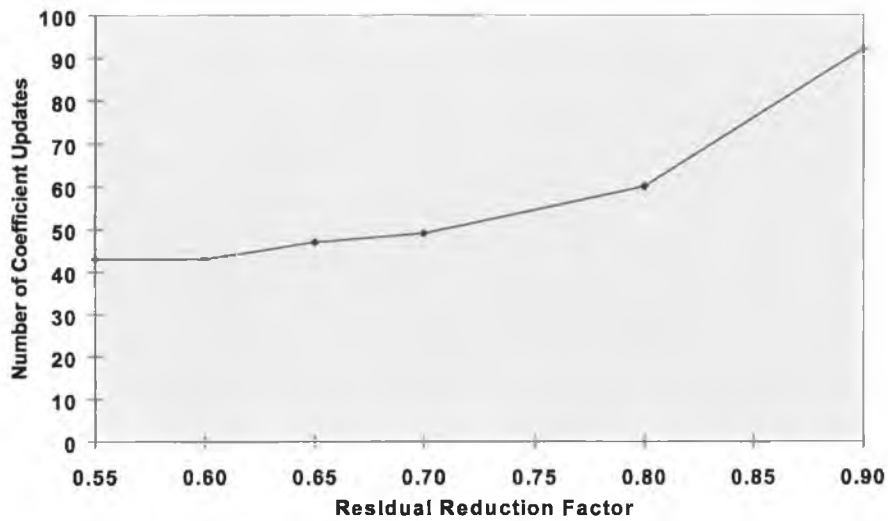
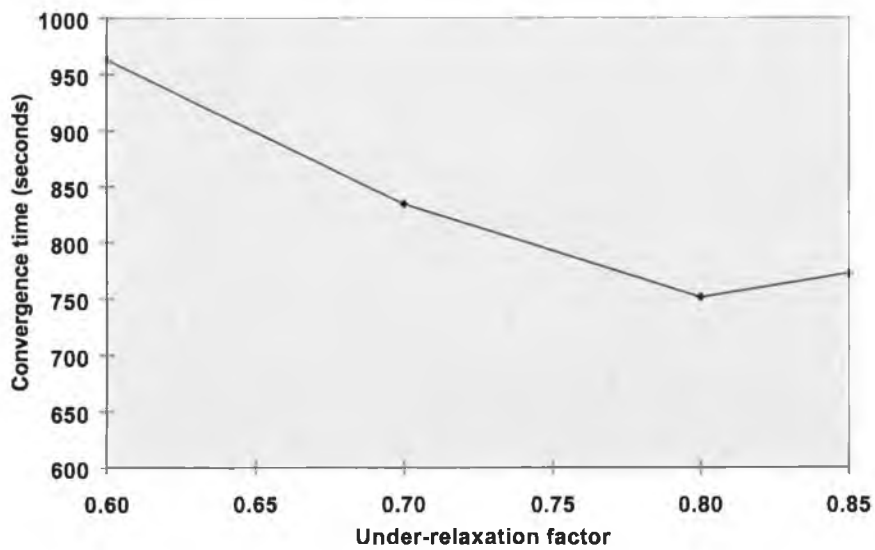**Figure 11.33. Variation of the number of coefficient updates with the residual reduction factor for CELS3D.**

Figure 11.34. Variation of the convergence time with the under-relaxation factor for CELS3D.
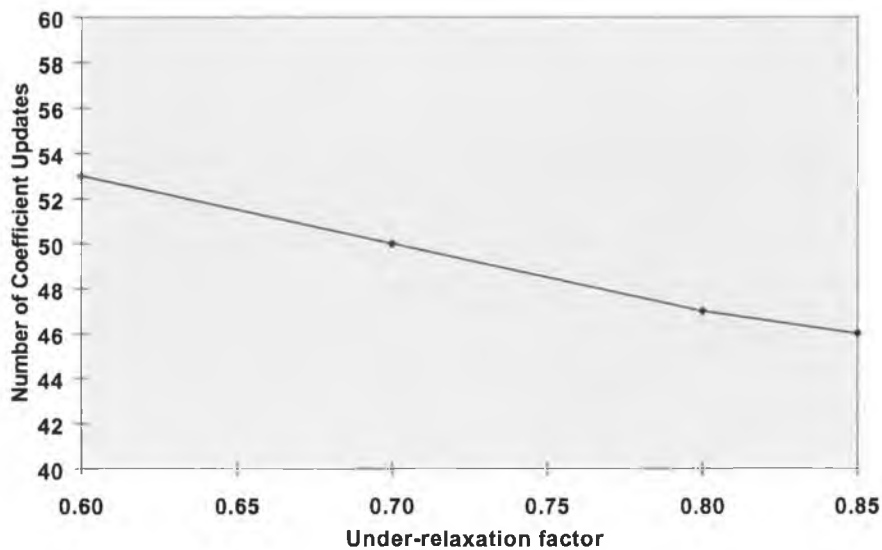
Figure 11.35. Variation of the number of coefficient updates with the under-relaxation factor for CELS3D.

Flows were calculated for two other values of the Reynolds number (7,000 and 10,000) and, in both cases, the factors controlling the convergence were tuned to give the most efficient convergence. The decrease in the residuals over time for all three values of the Reynolds number are shown in figures 11.36 and 11.37. It can be seen that the computational effort required increases with increasing Reynolds number.

The values of the factors controlling the convergence for each of the three values of the Reynolds number are presented in table 11.8. The results in table 11.8 show that more under-relaxation of the air velocity component and the turbulence equations is required as the Reynolds number increases and that greater effort per coefficient update is required. The decrease in the residual reduction factor as Re increases shows this last point. As well as the most efficient convergence being at lower values of these factors, note that convergence was not possible at Re = 7,000 and Re = 10,000 using the parameter values for Re = 3,500 and that some reduction in these values was necessary to avoid divergence.

255

Table 11.8. Convergence parameters for three values of the Reynolds number.

	Reynolds number		
	Re = 3,500	Re = 7,000	Re = 10,000
Vel. relax.	0.99	0.95	0.94
Turb. relax.	0.97	0.94	0.95
Cycles	3	3	2
Sweeps	6	3	2
Resid. reduction	0.65	0.45	0.35
CELS relax.	0.8	0.85	0.8

Table 11.9 contains the times to convergence and the saving in the computational effort that was produced by the use of CELS3D relative to the performance of the SIMPLE solver. At the lowest Reynolds number, the saving is very worthwhile with the CELS3D solver requiring only 21% of the effort using SIMPLE. However, this performance advantage is eroded as the Reynolds number increases and at the highest value the savings, while significant, are not as dramatic.

If the values of the relaxation factors and other convergence controlling parameters, as used for the highest value of Re, are applied to the lowest value of Re then the convergence time increases to 1328 seconds and the saving in computational effort for this case is reduced to 58%. Thus, using the extra relaxation required for the high Reynolds number flow still produces a useful saving across the range of values tested. Operation of the solver at the values for high Re is what might be required for general application of a solution code based on the CELS3D technique.
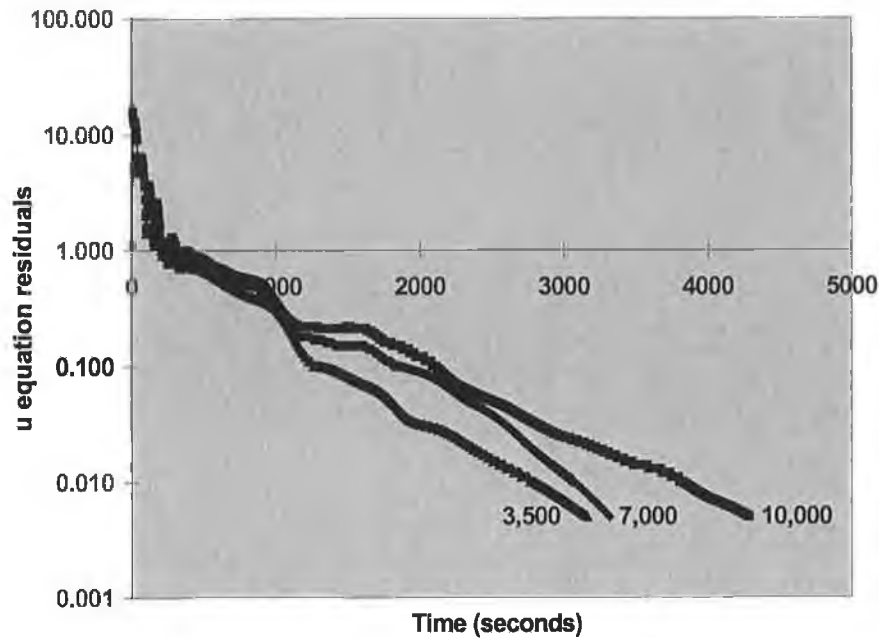
Figure 11.36. The u-equation residuals as a function of time using SIMPLE for three values of the Reynolds number.

Table 11.9. Performance of SIMPLE and CELS3D.

Reynolds number	SIMPLE (seconds)	CELS3D (seconds)	% SAVING
10,000	4,288	2741	36%
7,000	3,317	1487	55%
3,500	3,164	680	79%

The relationship between solution time t and the Reynolds number Re was of the form $t \propto \mathrm{Re}^{\beta}$ where β was approximately 0.25 for SIMPLE and 1.3 for CELS3D but this latter value would be dependent on the tuning of the solver convergence parameters. These values can be compared with those for the following test case.
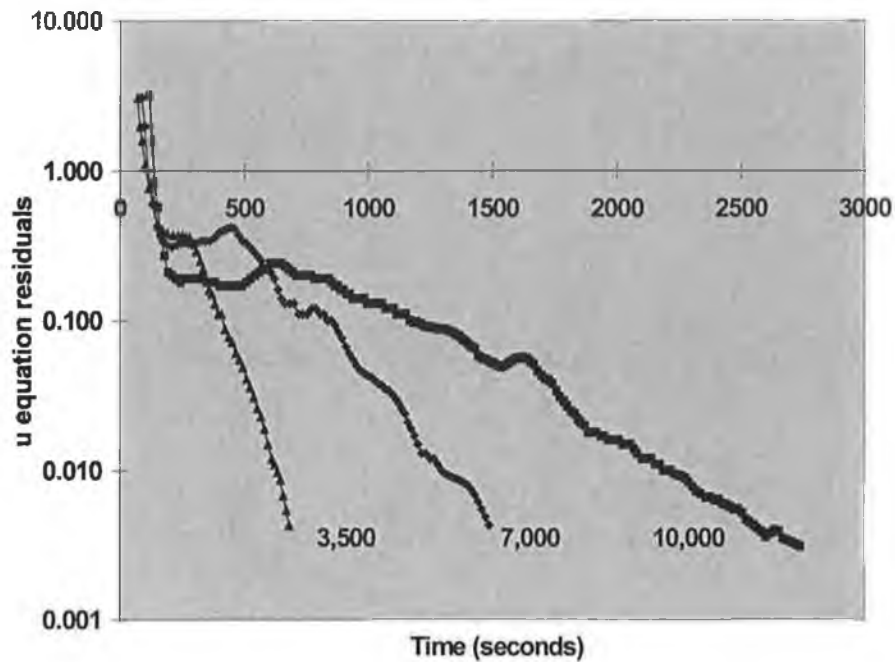
257

Figure 11.37. The reduction of u-equation residuals as a function of time using CELS3D for three values of the Reynolds number.

11.5 Second test case

A second test case is illustrated is figure 11.38 and was constructed purely as a test for the CELS3D solver with an alternative and very different geometry. This was done to determine whether the behaviour of the CELS3D solver as identified in this chapter was specific to the application-related problem or would, in fact, be consistent for the solution of other flow problems. It has no relation to any actual flow of current importance in the mushroom-growing application. Using a rectangular flow domain with the same specification as that used in Section 10.1, the entry and exit apertures for the flow were placed at the centres of the back and front walls, respectively. The inlet was square and extended from $i = 12, j = 12$ to $i = 22, j = 22$ while the outlet was square and extended from $i = 14, j = 14$ to $i = 20, j = 20$. Boundary conditions were,

again, the same in form as those described in Section 4.10. The inlet velocities were set to provide the same values of the Reynolds number.
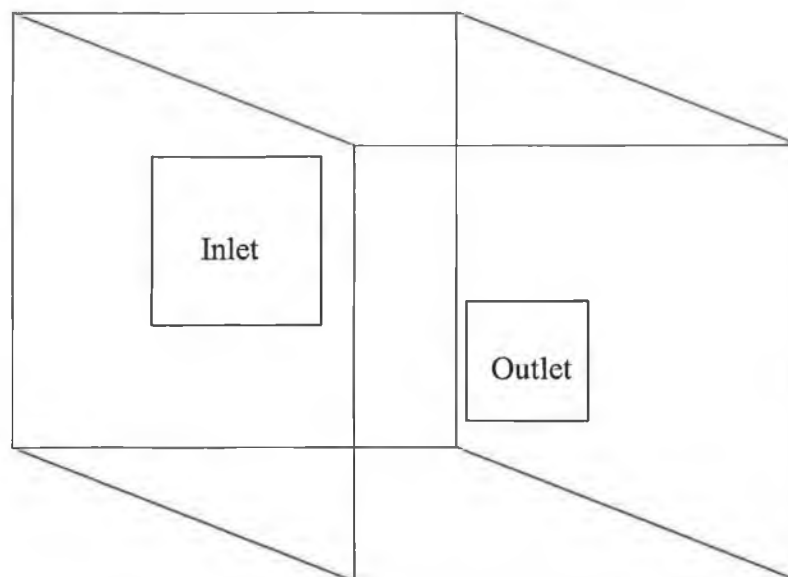

Figure 11.38. Schematic diagram of the geometry for the second test case for the CELS3D solver.

Testing of the behaviour of CELS3D for the second test case was similar to that for the first case. Again only the results for one value of the Reynolds number are presented in detail and, for a value of 7000, the parameters controlling the convergence (relaxation, residual reduction, etc.) were tuned for a good convergence time and then they were varied in turn about these near-optimum values to obtain information about their effect on the progress of the solution.

The first factors examined were the numbers of sweeps and cycles within the calculation procedure for the turbulence equations. First the number of sweeps of the tri-diagonal solver for the turbulence equations was addressed. Figures 11.39 and 11.40 show the results of varying this factor and the best convergence time and minimum coefficient updates were given by a value of 2.

259

Figure 11.41 shows the variation of convergence time with the number of cycles through the turbulence parameter equations and figure 11.42 shows the number of coefficient updates required for each case presented in figure 11.41. A value of 2 was selected as the optimum for further testing.

Figures 11.43 and 11.44 show the result of varying the relaxation factor applied to the air velocity component equations. A value of 0.98 minimised the convergence time and the number of coefficient updates and the solver was found to be divergent above values of 0.99. This behaviour contrasts with that noted for the first test case where it was possible to operate at all values up to 1 (no relaxation). Using this value of 0.98, the relaxation of the turbulence equations was varied and again a threshold was found above which the solver diverged. This value was 0.98 and it can be seen from figures 11.45 and 11.46 that the value for best convergence time and minimum number of coefficient updates was 0.95. The behaviour here contrasts with that for the first test case where the convergence time and the number of coefficient updates decreased steadily as the under-relaxation factor approached a value of 0.97, above which the CELS3D procedure diverged.
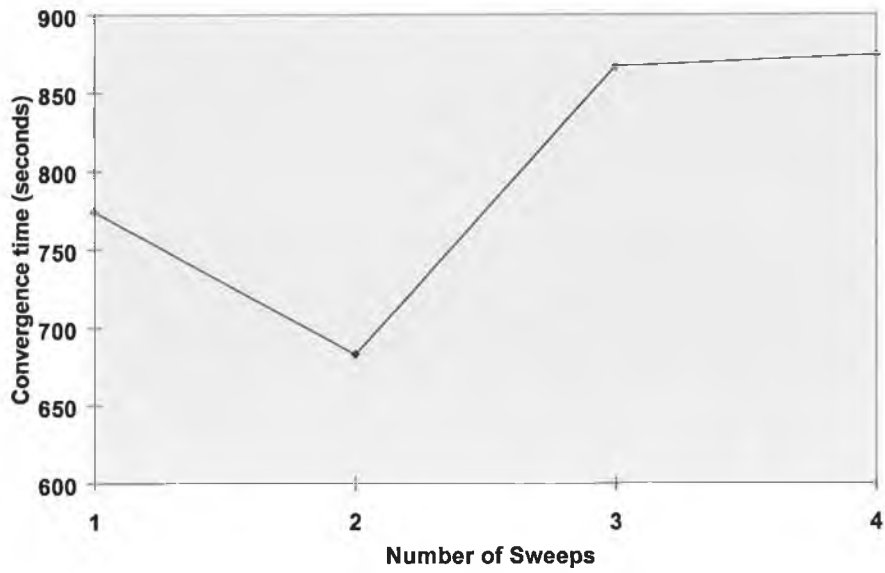
Figure 11.39. Effect of the number of sweeps of the turbulence equations on the convergence time.
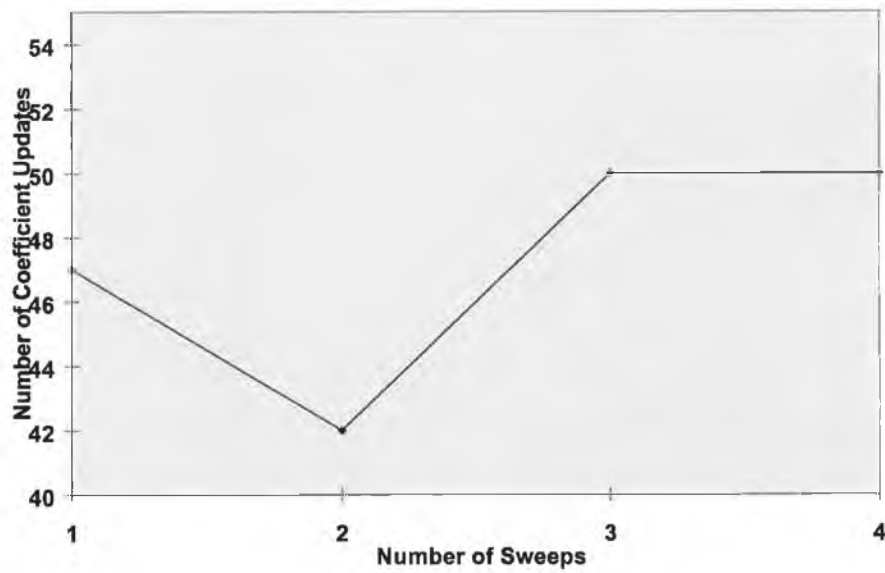

Figure 11.40. Effect of the number of sweeps of turbulence equations on the number of coefficient updates.

Figure 11.41. **Effect of the number of turbulence calculation cycles on the convergence time.**
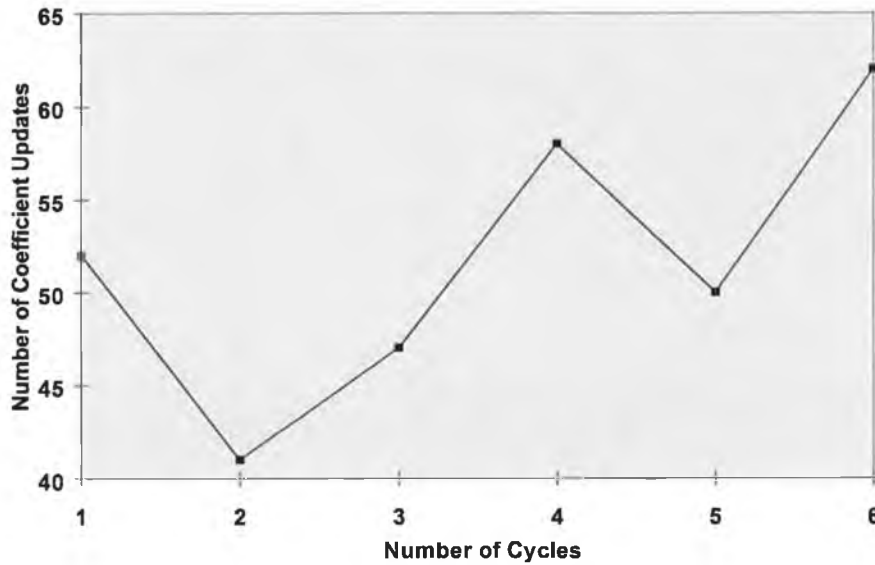
Figure 11.42. **Effect of the number of cycles of turbulence calculation cycles on the number of coefficient updates.**
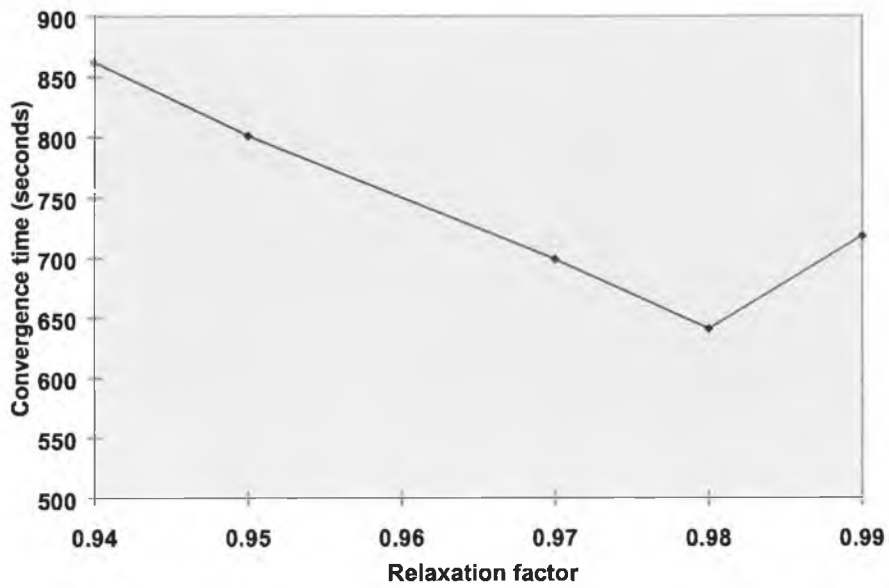
Figure 11.43. Effect of the under-relaxation of the velocity equations on the convergence time.
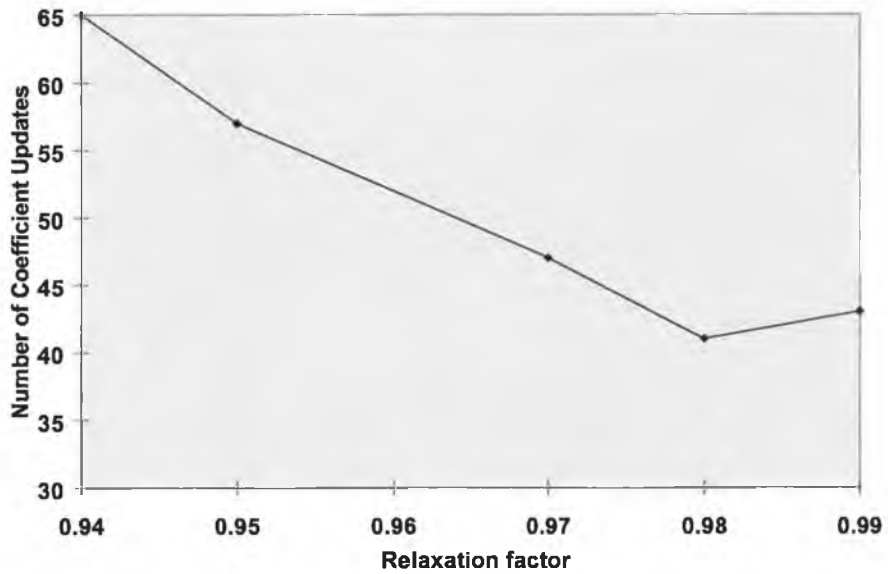
Figure 11.44. Effect of under-relaxation of the velocity equations on the number of coefficient updates.
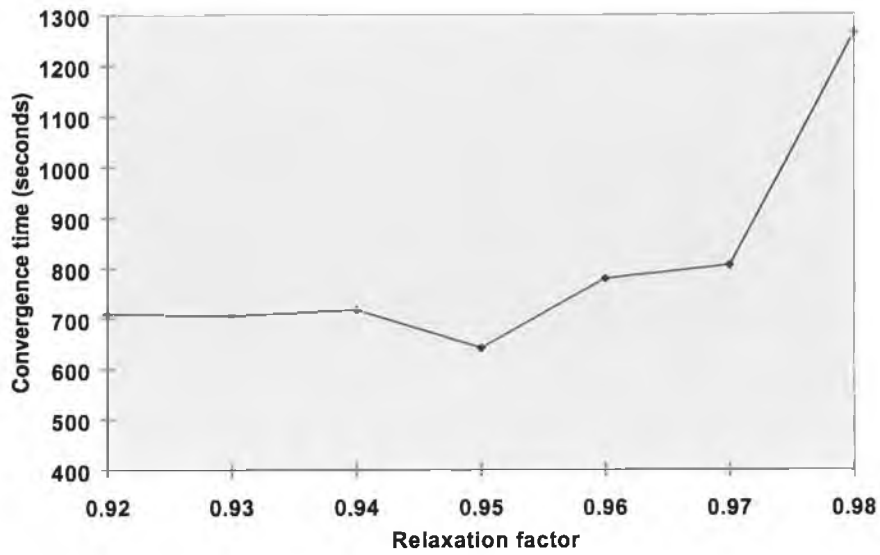
Figure 11.45. Effect of the under-relaxation of turbulence equations on the convergence time.
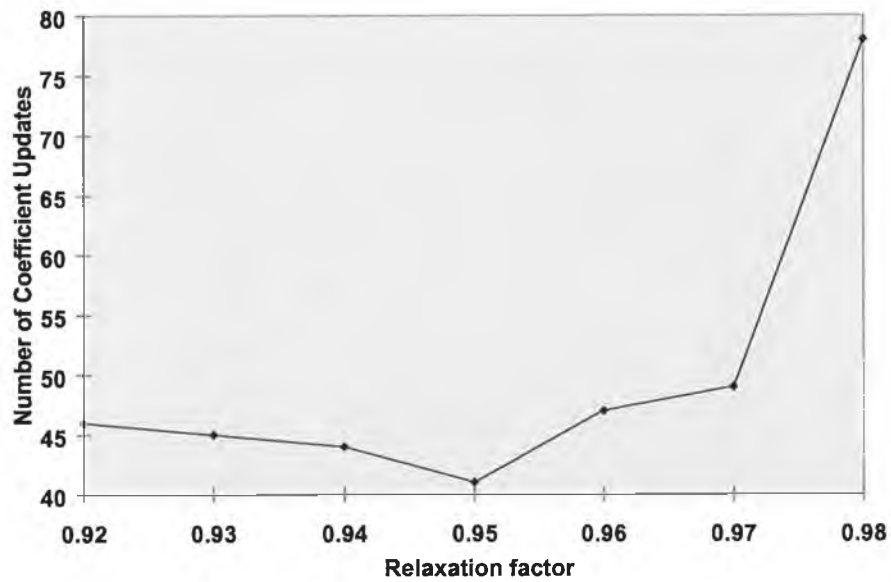
Figure 11.46. Effect of the under-relaxation of turbulence equations on the number of coefficient updates.
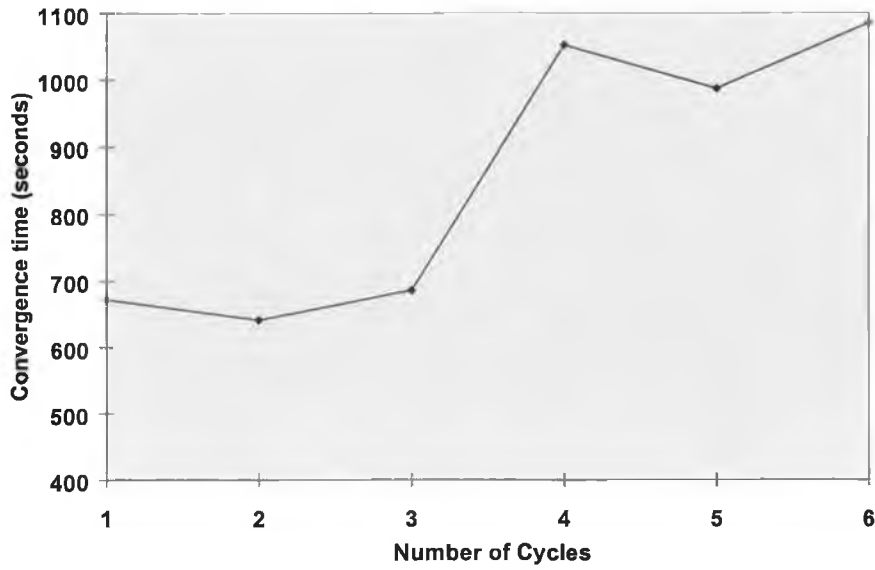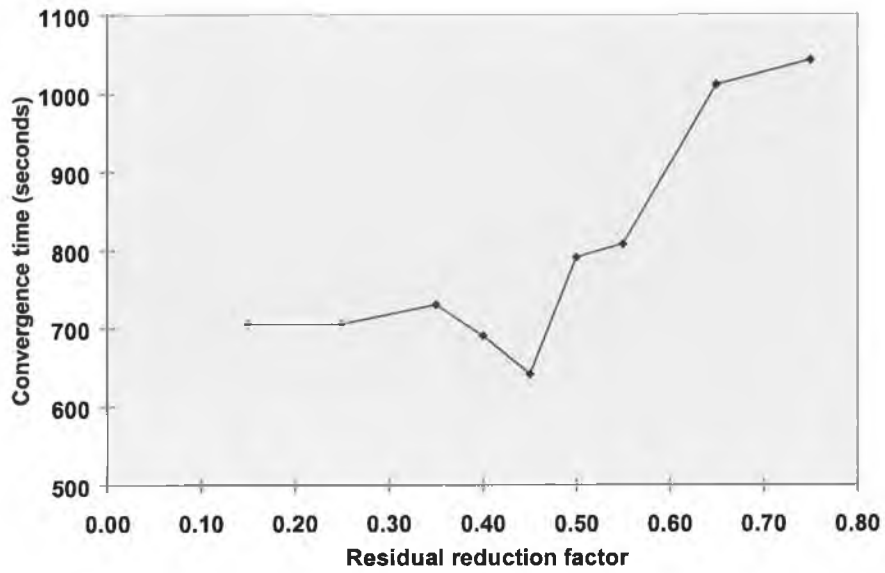
Figure 11.47. Effect of the residual reduction factor on the convergence time for CELS3D.
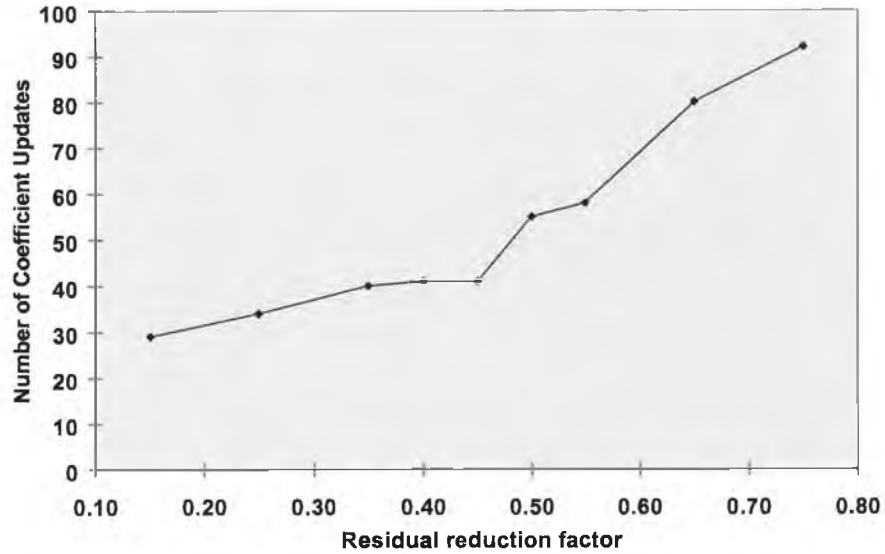
Figure 11.48. Effect of the residual reduction factor on the number of coefficient updates for CELS3D.

Figure 11.49. Effect of the under-relaxation factor on the convergence time for CELS3D.

Figure 11.50. Effect of the under-relaxation factor on the number of coefficient updates for CELS3D.

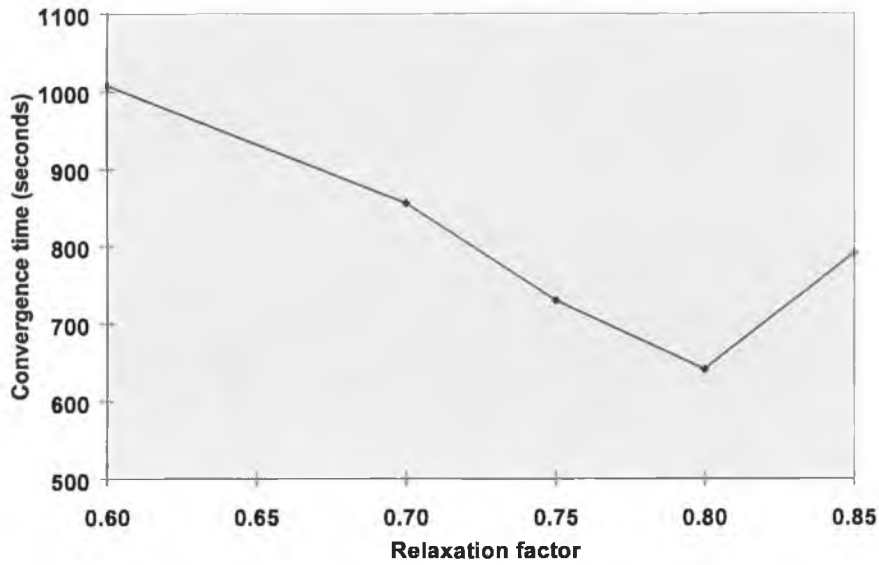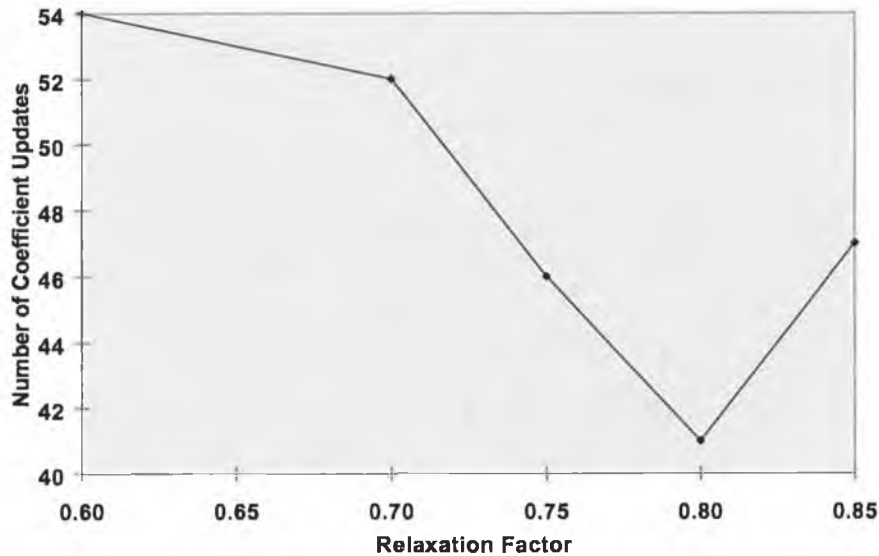With values determined for all of the external controlling parameters, those factors internal to CELS3D were examined. The residual reduction factor was varied (figures 11.47 and 11.48) and, while a value was found that minimised the convergence time,

the number of coefficient updates continued to decrease as the value of the residual reduction factor was reduced. This was not unexpected as a greater computational effort per coefficient update would be expected, at least over some range, to reduce the number of updates required. The value of 0.45 that resulted in a minimum convergence time was chosen for further testing.

Clear minima were produced for both the convergence time and the number of coefficient updates when the value of the CELS3D relaxation factor was varied (figures 11.49 and 11.50). The minimum can be seen to correspond to a value of 0.8.

This type of testing was repeated for the three values of the Reynolds number that were used for the first test case above and these results are presented in table 11.10. Unlike the first test case, there is almost no variation required in the values of the convergence parameters and the only change occurs for $Re = 10,000$ when the number of cycles of the turbulence parameter increases from 2 to 3. Only the solution for $Re = 10,000$ was found to be exhibit appreciable turbulence.

The path to convergence for each of the cases tabulated is presented in figures 11.51 and 11.52. Figure 11.51 shows the results for the SIMPLE solver. It can be seen that the residuals for all three cases follows very much the same trajectory at first and then, approaching half way through the time to convergence, they deviate significantly. For the CELS3D solver it is easier to distinguish between the three cases from the beginning. The saving in computational effort at $Re = 3,500$ is not as good as for the first test case but in that case higher relaxation factor values were in use and it was likely that a general application of the solver would demand a set of relaxation factors of lower value.

The savings in computational effort using CELS3D, while useful, were not dramatic and might not justify its use unless other advantages were found in its application or if a multigrid technique could produce further large savings. The application of such a technique is the subject of Chapter 12.

Table 11.10. Convergence parameters for three values of the Reynolds number

	Reynolds number		
	Re = 3,500	Re = 7,000	Re = 10,000
vel. relax.	0.98	0.98	0.98
Turb. relax.	0.95	0.95	0.95
Cycles	2	2	3
Sweeps	2	2	2
Resid. reduction	0.45	0.45	0.45
CELS relax.	0.8	0.8	0.8
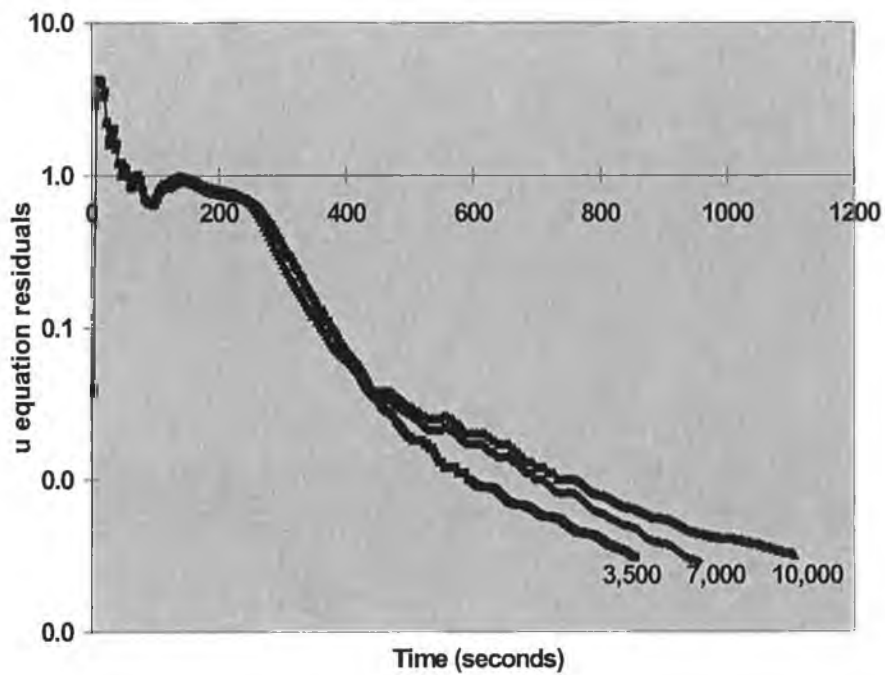
Figure 51. Reduction of the *u*-equation residuals using SIMPLE for three values of the Reynolds number.
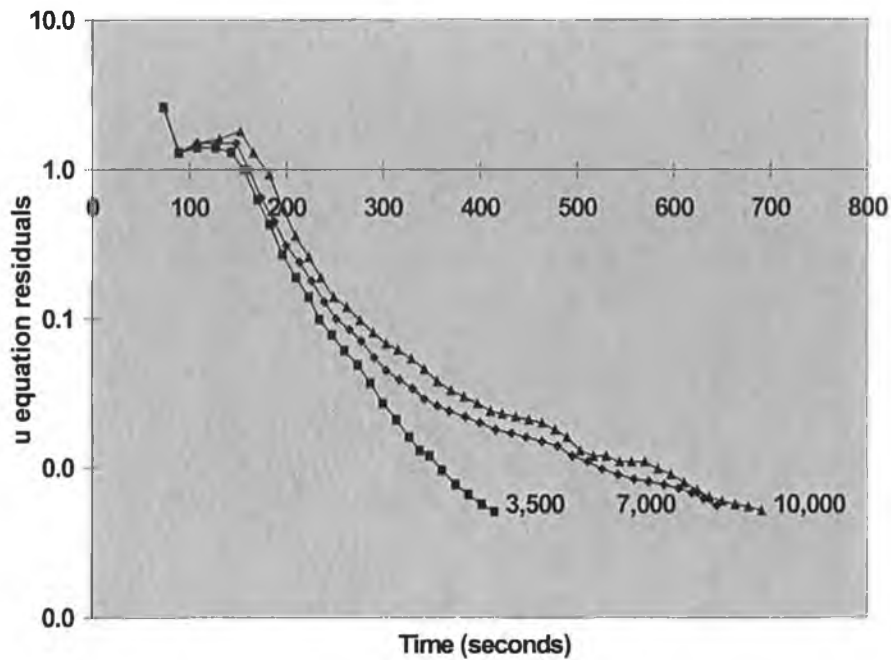
Figure 52. Reduction of the *u*-equation residuals using CELS3D for three values of the Reynolds number.

Table 11.11 Performance of SIMPLE and CELS3D

Reynolds number	SIMPLE (seconds)	CELS3D (seconds)	% SAVING
10,000	691	1106	38%
7,000	644	956	33%
3,500	415	856	49%

It is of interest to note that, for both test cases, the increase in time taken to achieve convergence as the Reynolds number increases was more significant for CELS3D than for the SIMPLE solver. The relationship for the second test case between solution time t and the Reynolds number Re was of the form $t \propto \text{Re}^{\beta}$ where β was

approximately 0.23 for SIMPLE and 0.5 for CELS3D but, again, this latter value would be dependent on the tuning of the solver convergence parameters. The values of β for the SIMPLE results were very similar in both tests (0.25 and 0.23) while the relationship for CELS3D was quite different. The careful optimising the factors controlling convergence for each Reynolds number might account for some of this. If the CELS3D solver were to be used only with the configuration required for the high Reynolds number then this would erode some of what appears to be a greater problem dependence for CELS3D.

Chapter 12

Application of ACM3D

12.1 The application of ACM in three dimensions

The multigrid method ACM3D produced no useful savings when applied with any of the CELS3D solution procedures that were developed in Chapter 10. Reduction in the number of cycles required for a given residual reduction factor were found using ACM3D but no overall savings resulted. All testing was carried out for laminar flows with the hybrid differencing scheme in order to reduce the number of possible sources of error. The grid densities used were not high ($34 \times 34 \times 34$ and $47 \times 47 \times 19$) but little evidence was found to create an expectation of greatly improved performance with increasing grid density. Any gains due to the ACM3D method required operation at CELS3D relaxation factors below those that gave best fine grid performance. Convergence times at these values of the relaxation factors were more than twice the best fine grid times.

In the course of its development, the method was first applied to variant 3 on a $34 \times 34 \times 34$ grid but no convergent solution procedure was obtained. It was expected that transfers to coarse grids using ACM3D should generate a set of u, v, w and p corrections that would improve residual reduction. This was not the case. The calculation of corrections on the coarse grid itself proceeded well and little effort was required for several decades of residual reduction.

No obviously anomalous values were encountered in the correction fields and considerable effort was expended in checking the code and testing it numerically. The code had already been through a comprehensive de-bugging process and any other problems encountered during this testing had no effect on the overall performance of the solver.

The investigation of the failure of ACM3D led to the generation of the procedure variants for the fine grid operation. Variant 4 was the first of the alternatives examined because efforts to understand the reasons for the poor operation of ACM3D led to the application of coarse grid operations that were geared to accelerate CELS3D in a pseudo-two-dimensional fashion. Since it was known that the two-dimensional form of ACM was effective it was hoped that trying to isolate a given plane set for solution would throw light on the reasons for failure.

12.2 PseudoTwo-Dimensional Operation of ACM3D

In applying the method to generate corrections for a given velocity component pair, no corrections were calculated for the third velocity component and its values were set to zero, i.e. it was assumed that the planes were being solved for the case of a perfect solution in the third direction. The technique produced an acceleration in the reduction of residuals on the fine grid when used for any given set of planes and the acceleration worked equally well for each set of planes.

Figure 12.1 shows the reduction of the residuals for one set of planes for a fine grid operation only and the corresponding reduction in the residuals when the accelerator was applied at one coarse level. Transfers to the coarse grid were applied from the first fine grid sweep. The CELS3D relaxation was set at the previously determined optimum value of 0.7 but operating at 0.6 produced no significant difference. Note that low values for the CELS3D relaxation factor were required for all situations that showed successful action of ACM3D on fine grid cycles.
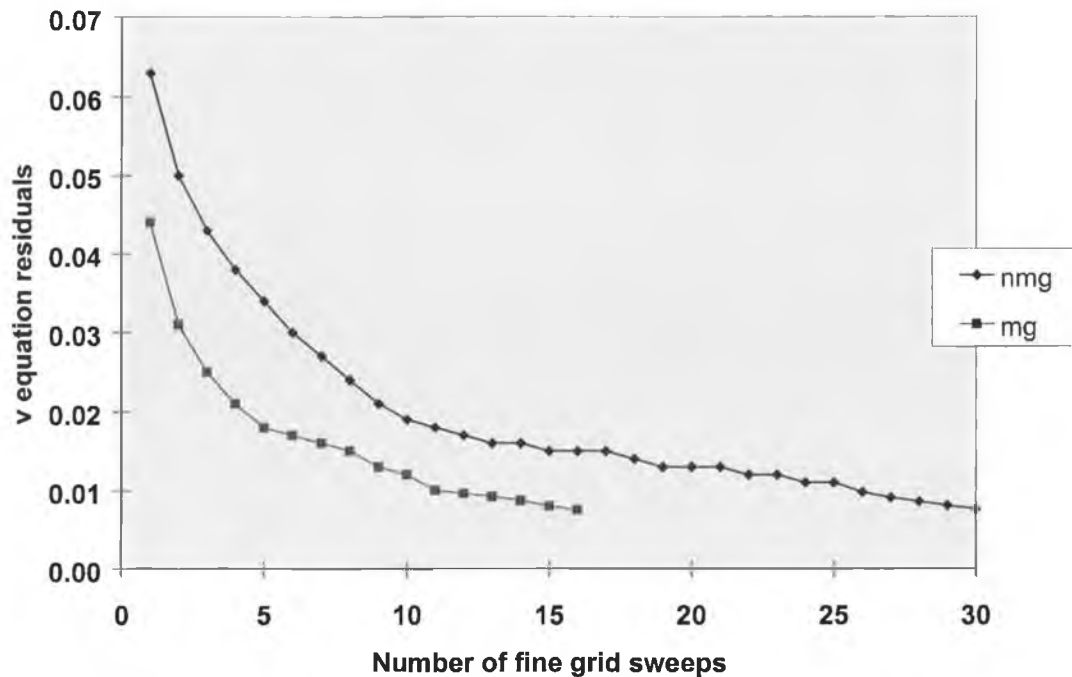
Figure 12.1. The variation in residuals as a function of the number of fine grid sweeps with ACM3D applied to vw planes (mg: multigrid, nmg: no multigrid).

The multigrid version (mg) required only 50% of the number of sweeps required by the fine grid solver alone. For this particular example, i.e. one coefficient update, with $\alpha=0.8$ on the coarse grid (effectively one sweep), the multigrid solver produced a time saving of 16% relative to the fine grid alone. In general, it was found that reductions in the number of sweeps for a range of values of α on the fine grid were considerable.

However, when the method was applied to the solution of a flow (in this case the test problem from section 10.1), it proved difficult to achieve an overall saving in computational effort, i.e. a reduction in the time to convergence.

The effect of the convergence control parameters on the application of ACM3D to variant 4 was investigated. Comprehensive experiments were carried out involving all the tuning factors for the solver, degree of convergence on the coarse grid and control of conditions for invoking the coarse grid operation, i.e. fine grid residual reduction

rate. The ACM accelerator worked well for each set of planes individually but these gains were not reflected in improvements at the coefficient updates. The net result was that, while the effort expended on any given set of coefficients was faster using ACM3D, there were extra coefficient updates that increased the total time to convergence.

It was possible to show a beneficial action of the ACM3D accelerator on variant 4 by reducing the relaxation factor on the air velocity equations (which caused the convergence behaviour of the CELS3D solver to increase markedly). When the ACM3D method was applied there were small savings (approximately 5%) in the time to convergence relative to the case with no acceleration. Such results were of no practical significance as solution times were comparable to those for SIMPLE. They simply served to test the coding and coarse grid transfers at all stages of convergence. It had been considered that the effectiveness of the method might improve as the solver converged because of the insufficient smoothing of the fine grid errors in the initial coefficient updates. Experiments with delaying the introduction of ACM produced no benefit and no evidence of significantly improved operation in the final updates was found in the reduction rates of the residuals.

Attention was directed again towards the solution by variant 3 but this time attempts were made to apply ACM3D in a manner that could make use of the pseudo-two-dimensional procedure, i.e. seeking to accelerate the solution on one set of planes to give overall savings in the time to convergence.

With one pass per plane set, a coarse grid solution in this mode could be applied either before or after the fine grid pass. Applying ACM by either method to a given set of planes alone improved the residual reduction for that plane but the reduction for the others slowed and, ultimately, the solution process became less efficient than CELS3D alone. The CELS3D relaxation was at low values (0.5 - 0.7) whenever a positive effect of ACM3D was noted.

Applying ACM3D in a similar manner to each plane set in turn during the solution produced an initial improvement but the gains were gradually eroded until ultimately

convergence was considerably less efficient than CELS3D alone. Using too great a value of the reduction factor α resulted in a divergent procedure. Reducing the CELS3D under-relaxation improved the action of the accelerator but this did not result in any overall improvements in the solution time compared to the best results from CELS3D.

Other unsuccessful experiments included variant 3 with two passes per plane set and the coarse grid solution carried out between passes. No overall savings were made.

12.3 Complete Correction Sets

Having spent some time working with corrections in pairs, attention was again directed towards the calculation of a complete set (u,v,w and p) of corrections after sweeping the entire flow domain. Residuals calculated before and after the coarse grid transfers showed consistent improvements but, as before, these failed to translate into improvements in the procedure. Variant 3 consistently produced divergent solutions when calculating complete correction sets.

Consideration of possible reasons for failure led to the idea that variant 3 did not optimise the interactions between sets of planes in the sense that holding one velocity component fixed while solving for the others placed a limit on the computational effort that should be expended on one set of planes before moving on to the next. There is no point in expending too much computational effort while the fixed values in the equations are Variants 1 and 2 were developed and ACM3D was applied, solving the complete set of corrections, i.e. u, v, w and p corrections. These were calculated after a complete pass through all planes.

A degree of smoothing on the fine grid was required and both variants showed reductions in the number of inner iterations when the computational effort for a given coefficient set was prolonged using low values of α. These methods produced no time savings in the reduction of residuals on the fine grid. By under-relaxing on the fine

grid it was possible to show convergence when applying ACM3D but the convergence was not competitive with the solver without acceleration. No significant difference was found between variants 1 and 2.

Figure 12.2 shows the application of one coarse grid during the third coefficient update. A low residual reduction factor was used to prolong the calculation for this set of coefficients. The fine grid solution (nmg) took eight sweeps more of the calculation domain than the multigrid application. The multigrid procedure was applied in two ways. The first (mg1) is operating from the first sweep. In this case the residuals initially exceed those for the fine grid and later reduce more quickly. If the application of the coarse grid was deferred until the fifth sweep (mg5) then the residuals were again briefly greater than the fine grid but they reduced quickly and finished with the same number of sweeps as the first multigrid case.

There seems to be a case for deferring the application of the coarse grid. This may reflect insufficient smoothing of the error components on the fine grid until after a number of sweeps have taken place. The total computational effort was lower where the multigrid application was deferred because both methods required the same number of fine grid sweeps but mg5 had no coarse grid computations at first. These results also illustrate the difficulty that arises in trying to apply ACM3D where only a small number of sweeps are required per coefficient update. In such situations application of the coarse grid can induce extra sweeps due to initial slowing of residual reduction.
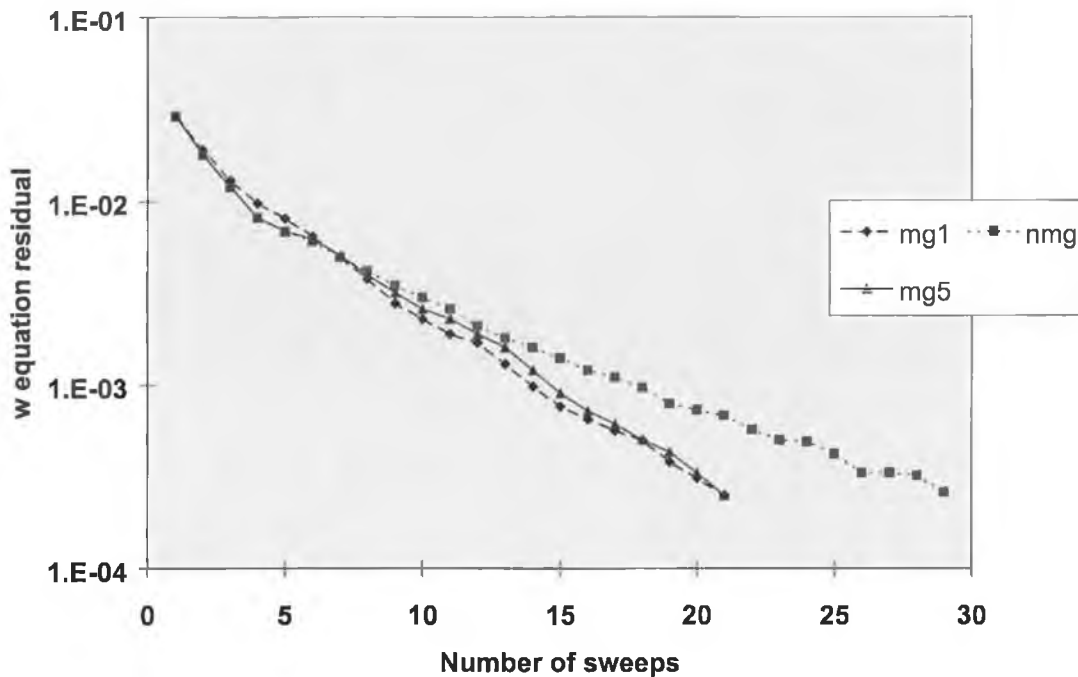
Figure 12.2. Fine grid residual reduction with no multigrid (nmg), multigrid after 1 fine grid sweep (mg1) and multigrid after 5 fine grid sweeps (mg5).

Since ACM3D does return corrections that show a reduction in the residuals calculated immediately afterwards, it was applied as the last step in a coefficient update cycle, i.e. the solution was reduced to a particular degree and then corrected immediately before recalculating the coefficients. Once again some benefit could be shown by altering the overall solver tuning to extend solution times but this technique produced no useful gains in efficiency.

12.4 Coarse to Fine Grid Ratios

The possibility of improvement by using coarser grids was examined. It is a widespread experience that a factor of 2 gives the best results when defining the relationship between coarse and fine grid spacings (Briggs, 1987 and Brandt, 1987)

In view of the fact that the accelerator does save iterations in some circumstances, but at the cost of increased convergence time, the use of still coarser grids could offer an advantage. The corrections that would be calculated would correspond to a lower frequency error component than a factor of two provides and iterations on a coarser grid are more economical.

Fortran code extensions for the ACM3D routines were prepared that moved coefficients directly to coarser grids at factors of 2, 4 and 8 for a $34 \times 34 \times 34$ grid in a square flow domain. These correspond to cell groupings of 8, 64 and 512. The behaviour of the residuals on the fine grid was monitored.

It is clear from figure 12.3 that the most effective coarse grid for reducing the residuals during inner iterations is the one with spacing a factor of 2 greater than the fine grid. The usefulness of the corrections generated decreases as the spacing factor increases. The important factor to be taken into account, however, is the time taken to converge these inner iterations. None of the ratios produced a saving in convergence time.
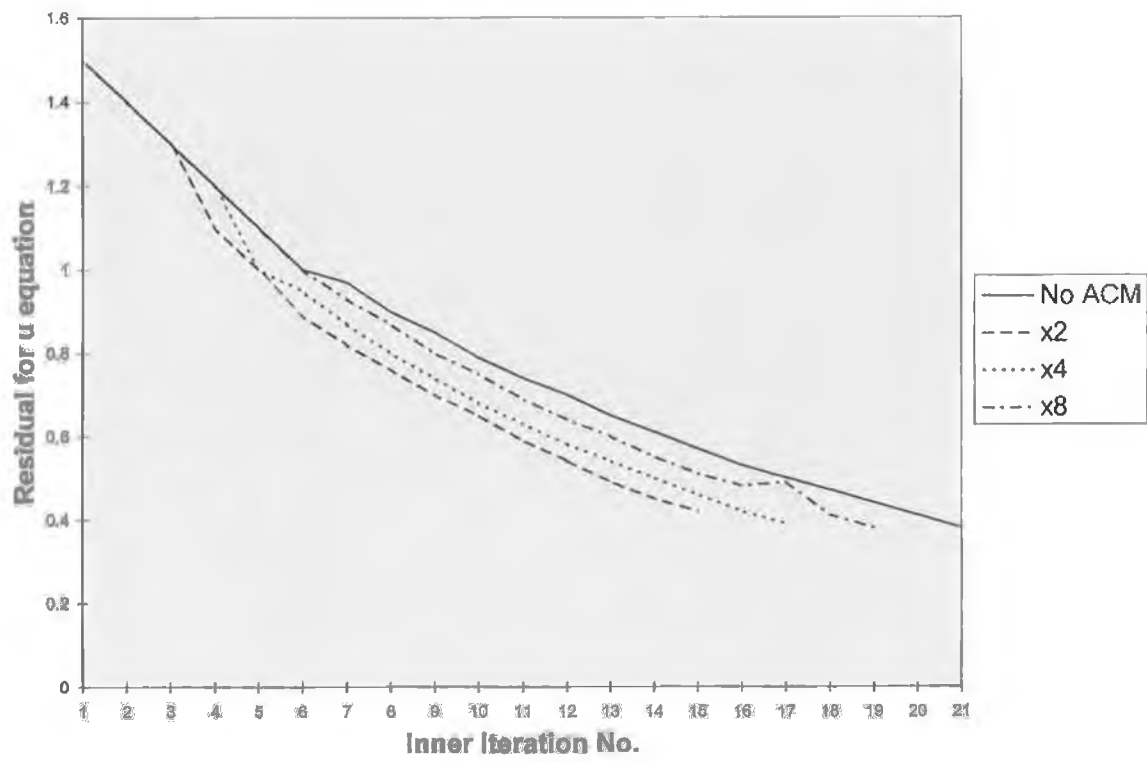
Figure 12.3. Fine grid residual reduction using a number of coarse/fine grid ratios.

12.5 Operation On a Denser Grid

Having failed to produce any benefit on the $34 \times 34 \times 34$ grid, it was decided to experiment with a denser grid. Since the $34 \times 34 \times 34$ grid was near the 32Mbyte limit of the compiler being used, the dimensions had to be altered to stay within the current memory constraints.   Therefore a $47 \times 47 \times 19$ grid was set up and the laminar test problem was set up. Variant 2 of the CELS3D solver was used. At the tuning set that was developed for the square grid laminar case the solver proved to be efficient and gave a short convergence time. The residual reduction factor was set to 0.6 and the CELS3D under-relaxation was set to 0.95 and there was no under-relaxation on the air velocity equations.

The first coarse grid was added to the solution procedure and calculations were carried out for a number of values of the CELS under-relaxation factor.

The addition of the corrections calculated on the coarse grid caused the solver to diverge at both 0.95 and 0.9. These values gave the fastest times for the fine grid solver. The convergence times are set out in table 12.1. Note that multigrid convergence was achieved, but not savings, at a value of 0.8, which is higher than was possible on the $34 \times 34 \times 34$ grid.

Table 12.1. Addition of one coarse grid to the 47x47x19 problem.

CELS3D Relaxation	Fine Grid		Fine Grid + 1 Coarse Grid	
	time (secs)	coefficient updates	time (secs)	coefficient updates
0.5	874	15	968	16
0.6	689	15	764	16
0.7	524	14	618	15
0.8	421	14	533	13
0.95	316			

With this tuning set, it was found that there was no benefit in terms of the overall convergence time.

During the course of experimentation with the convergence tuning factors, a saving was produced at a low CELS3D relaxation factor. This case was for a situation of varying the relaxation on the air velocity equations, i.e. the solution procedure began with a factor of 0.98 and proceeded to 0.999 over the course of 10 iterations. As noted above, this approach had no merit for a fine grid solution and the results are reproduced in table 12.2 as the only example found of significant savings (relative to the fine grid operation with the same tuning). The residual reduction factor was 0.5.

As in the previous results, the solver diverged at residual reduction factors of 0.9 and 0.95 when the corrections were applied.

Table 12.2. Times to convergence and number of coefficient updates for a range of CELS3D relaxation factors both with and without the use of mulitigrids.

CELS3D Relaxation	Fine Grid		Fine Grid + 1 Coarse Grid	
	time (secs)	coefficient updates	time (secs)	coefficient updates
0.5	1210	16	1179	16
0.6	895	15	921	16
0.7	727	15	753	15
0.8	584	14	1129	13
0.9	529	14		
0.95	516	14		

Table 12.3 contains the numbers of fine grid iterations for each of the coefficient updates for the settings set that produced a saving on convergence time. The action of the coarse grid corrections is clearly seen in reductions for the multi-level case. The normalised residual shown is the maximum of those for the three air velocity equations at each coefficient update. It should be noted that the test for convergence is made after the coefficients have been updated and before the solution for that set. This is why the multigrid case appears not to meet the convergence criterion. The information in the table was stored at the end of the last solution cycle.

Table 12.3. Number of inner iterations and the residuals of the velocity component equations both with and without the use of multigrids.

Update number	Fine Grid		Fine Grid + 1 Coarse Grid	
	No. of inner iterations	Normalised residual	No. of inner iterations	Normalised residual
1	20	1.3	20	1.3
2	3	3.3	3	3.3
3	5	1.8	5	1.8
4	6	0.99	8	0.99
5	13	0.61	10	0.62
6	21	0.44	14	0.43
7	30	0.31	20	0.31
8	41	0.21	26	0.22
9	56	0.16	37	0.16
10	71	0.11	55	0.12
11	52	0.075	37	0.085
12	34	0.04	30	0.047
13	33	0.021	28	0.027
14	29	0.013	26	0.017
15	28	0.008	25	0.011
16	28	0.005	24	0.0069

The residual reduction factor on the coarse grid was 0.5. This resulted in only one iteration on the coarse grid before applying corrections and returning to the fine grid. There was therefore no scope for adding further coarse grid levels. Also, while lower values of α on the coarse grid increased the computational effort and eroded the savings, no benefit from the increased accuracy of the solution was seen in terms of the fine grid residuals or the reductions of the numbers of fine grid inner iterations.

Thus, one coarse grid cycle produced an acceptable degree of convergence in the calculation of corrections on that grid.

Unfortunately, the gap between the convergence times generated by these savings and the best performance from the fine grid procedures alone is very large. This is taken as indicating that the likely benefit from operating on higher grid densities was poor, except possibly for very dense situations that would result in convergence times that put the testing of these outside the time scale of this work.

The most important factor in this decision is the lack of any benefit as CELS3D relaxation factors increase. This could indicate a structural problem in that the fine grid solver is not suitable for the use of ACM3D. Therefore no action was taken to extend the memory or change the compiler for the purpose of increasing grid density.

Also, it may be that there is some element of the pseudo two-dimensional efficiency that was noted above. If this was the case then the savings might not carry through on a, say, 47×47×47 grid.

The QUICK scheme was implemented as a final check to ensure that improved accuracy would not affect the behaviour of the ACM routines. Apart from an increase in the time to convergence there was no change in the relative performance both with and without ACM3D.

Chapter 13

The Application of CELS3D to Mushroom Growing Structures

13.1 Applications

As this work was completed before detailed measurements of air flows in mushroom tunnels were available, it wasn't possible to evaluate the accuracy of the results or to examine the suitability of the approximations made in describing the inlet conditions and the tunnel walls. Nevertheless, it was possible to examine the application of the CELS3D solver to the type of problem to be solved for the mushroom-growing application. Therefore, this chapter deals with the general principles of the application and shows that difficulties can arise in some cases.

13.2 Curved walls

As in the extension of the SIMPLE code to three dimensions in Chapter 8, CELS3D was used to calculate air flows in a curved wall tunnel and the flow configuration details were similar to the test case of Section 10.1. The air inlet was a narrow slot running along the tunnel and the air exhaust was at a low level at one end of it. The flow was thus fully three-dimensional in nature.

The wall shape, and hence the flow domain internal to the calculation domain was modelled in exactly the same way as for the SIMPLE case, i.e. using source term

modifications to set the air velocity components outside the flow domain to zero and setting appropriate boundary conditions at the internal wall.

The application of CELS3D for this flow produced no serious difficulties. Figure 13.1 shows a flow calculated for such a wall shape at a Reynolds number of 3500. It was necessary to reduce the values of the under-relaxation factors for the solution of both the air velocity component and the turbulence model equations. Whereas these values were 0.99 and 0.97, respectively, for the first test flows in the square room, in this case it was found that they had to be reduced to 0.95 and 0.85, respectively, in order to achieve convergence and to make it as efficient as possible. For this flow, CELS3D converged in 33% of the time taken using SIMPLE. On increasing the Reynolds number to 10,000 the air velocity component relaxation factor for the CELS3D solver had to be reduced to 0.85 but the solution was still obtained with just 37% of the computational effort required for the SIMPLE-based solver.
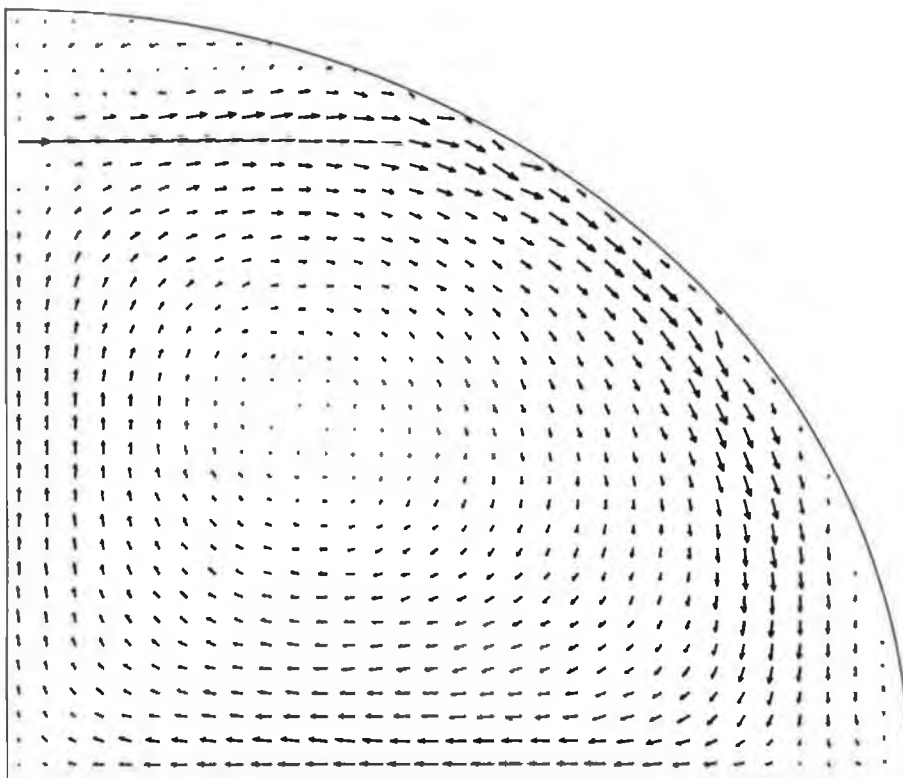

Figure 13.1. Three-dimensional flow field (vertical plane k = 15) incorporating a curved wall (Re=3500).

13.3 Flow around shelves

Complex flows in mushroom growing that are of considerable current interest are those where a shelf system is employed, as described in Section 1.2. Research work in Holland (Loeffen, 1992), where growing is exclusively in this type of system, has empirically examined the air distribution within these growing rooms. Air ducts were designed to provide a uniform delivery of air into the space but no means of addressing the distribution of air within the space was found. In industry this situation is accepted as the best compromise that can be made between the effort that would be required to improve the control of air flow and the cropping effects due to the current non-uniformity of that flow. In these growing rooms, with five or six levels of shelving, it is very difficult to achieve a uniform air flow but the development of a hybrid system in Ireland with three levels of shelving (limited by the shape of the existing tunnels) makes the problem more approachable. An experimental study of these flows is difficult and time-consuming and a computational approach could allow flexibility in experimenting with methods of air distribution. It would be an important feature of any solver that it should be capable of handling the internal boundary conditions that would be required to model these flows.

13.4 Internal boundary conditions

Internal boundary conditions are easily inserted under the segregated SIMPLE-type solvers because of the facility afforded by the use of large numbers in the source terms as explained in Section 4.9. As in the use of CELS in Chapter 9, and CELS3D in Section 13.2, this type of boundary was applied in the crude modelling of the curved surface of the tunnel wall and convergence was achieved without complication in that instance. The same boundary conditions have been applied in both approaches. There was excellent agreement between the results from both solvers and it seemed that the technique would be useful for the modelling of the internal boundaries that would be used to model the shelving used for mushroom growing.

However, this curved wall turned out to be a special case. Difficulties arose in trying to insert a block in the flow field as the first attempt to model the inclusion of shelving. When the appropriate air velocities were set to be 'turned off' using the source terms then the CELS solver diverged, or began to oscillate, at an early stage in the process.

The main difficulty is due to the simultaneous nature of the solution procedure in CELS. When the air velocity component equations were independently solved, in turn, by the tri-diagonal routines in the SIMPLE solver then there was no inter-equation interaction possible between the modified source terms in the air velocity component equations. This situation has changed in CELS. When large numbers are used in the source terms then complications can arise in the solution routines for a number of reasons, namely,

i. Attempting to calculate the pressure within an internal boundary produces excessively large values if the source terms are not also modified for the neighbouring grid cells on the line just calculated. It was this point that fortuitously allowed the insertion of the curved wall as in the calculations described above. In that case, every point that had modified source terms in the associated equations was also bordered by a point with the same treatment where calculations had just been performed.

ii. Solving the air velocities along a line using CELS3D is more involved than the SIMPLE case because, for the penta-diagonal equations involved in the procedure, the influences on the calculations extend further than the adjacent finite volumes.

iii. Attempting to turn off (i.e. set to zero) the v air velocity components, say, associated with a given grid point while leaving the associated u air velocity component as part of the flow field does not allow a correct calculation of the u air velocity components as the source values for the v component equation 'leak' into the u line solution.

Firstly, then, in order to make any progress, it is now necessary (presuming that no set of complicated conditions is to be included in the code) to think of the air velocity component and pressure fields as a mutually dependent set with a structure that should be treated as a unit for the purposes of inserting internal boundary conditions. For any given scalar grid point, the associated staggered grid velocity components must both be set to zero in order to solve the flow around internal boundaries using CELS3D. In order to model a given geometry then, the shape in question has be set up to consist of sets of these units, i.e. if the shape to be inserted is to encompass the position of a given u velocity component then it must also include the associated v component and scalar grid point. This unit is shown in figure 13.2 which consists of the scalar grid point and the associated (sharing the same grid reference) air velocity components on their staggered grids. In three dimensions the staggered w velocity would also be included in the set.

When inserting internal boundary conditions, the object should be modelled as consisting of a combination of these units. Breaking these units by setting one of the air velocity components to zero causes difficulties that leads to a requirement for special additions to the coding in the solver routines.

Operating with these calculation units it was possible to trace other difficulties back to the operation of the continuity equation. In the construction of the coefficients for a line solution then this equation provides a link between the current line and the one just calculated. It is not possible to insert and modify a source term in this relation to remove the problems that arise. It is not possible to prescribe no-flow continuity boundaries for the block to be inserted because the products of the cell face area and the fluid density that are on the south and west faces of the continuity volumes are divisors for other terms in the calculation of the coefficients for the CELS equations.

However, by setting the north and west continuity faces to zero then the immediate difficulty with convergence was overcome. The convergence time of CELS was poor, relative to SIMPLE, because the air velocity equations needed to be set at a low under-relaxation factor of 0.85. Unfortunately, the flows calculated at the top of the shelf were not realistic and did not agree with those generated by the SIMPLE solver.

Setting the continuity coefficients to zero achieved the correct effect of containing the large source terms but did not offer a solution to the problem. There were also complications at the corners of the block due to changing continuity coefficients which meant that SIMPLE and CELS3D would produce different results for these areas.

Although it seemed that some modification of the north and east boundary conditions for the block might improve the situation, no method of overcoming this problem in a straightforward manner was found.
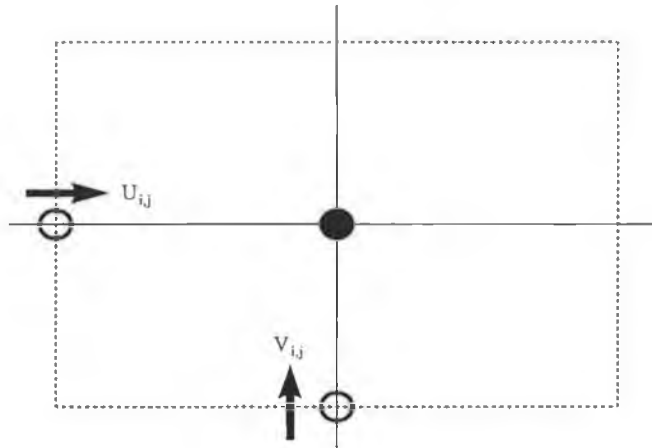
Figure 13.2. Combination of variables to be considered as a unit for the modelling purposes.

13.5 An alternative method

With the breakdown of the method of large source terms then an alternative method had to be found. No convenient adjustment to the equation coefficients was found. It was noted that the introduction of a shelf as an internal boundary and the consequent fixing of the air velocity values on and within the boundary at zero values was the same situation that produced the need for a special procedure at the edge of the computational domain after a line sweep. The general CELS method breaks down at

the bottom of a line by the line sweep and requires the introduction of the procedure, as outlined in Section 5.3, which applies mass conservation to calculate the air velocity component on the last line of the sweep.

Wherever there are fixed values within the flow field, mass conservation must be applied locally to satisfy the equations and produce a physically realistic solution. Referring to figure 13.3, which shows a block representing a rectangular section of a shelf, the requirements for the solver can be explained. This figure is an extended version of the one for the last line procedure in Section 5.3 and it is to make clear the similarity between calculations for a block in the calculation domain and those for the last line of a sweep. The continuous lines and dotted intersections represent the scalar grid and the dotted lines are the boundaries of the finite volumes for this grid. The arrows represent the staggered grid air velocity components and therefore they lie on the dotted lines between the scalar grid points.
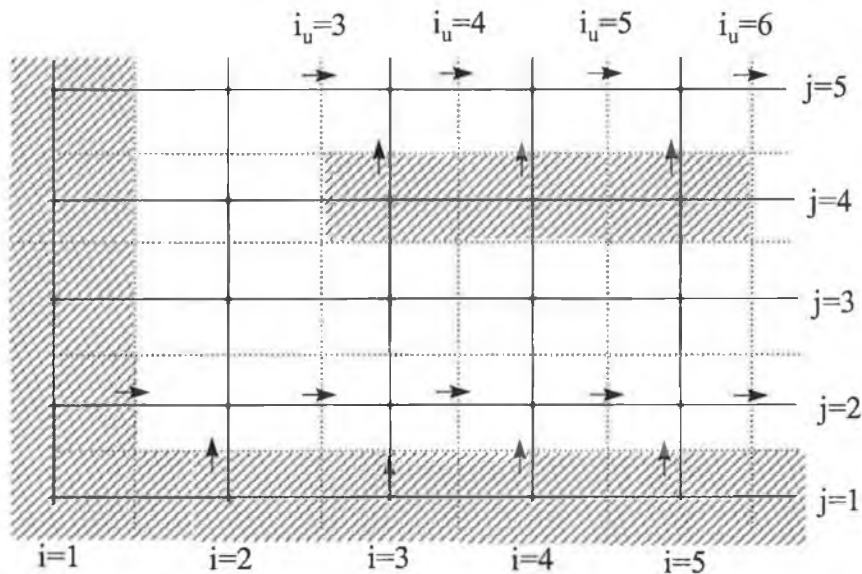

Figure 13.3. A solid block inserted in the flow domain.

When the sweep for the horizontal air velocity component reaches the line $j = 5$ in the figure, the vertical velocity component associated with that line (between lines $j = 5$ and $j=4$) are no longer defined by the equation for that component. They are fixed at zero, as at a wall boundary, and the calculation procedure breaks down in exactly the same way as for the last line solution in a sweep of the entire calculation domain. It becomes necessary to switch from the line solution for the horizontal air velocity components to a mass conservation procedure. The approach required is that the outer components ($i_u = 3$ and $i_u = 6$) are determined by the line solver and between these (along the surface of the block) the other components (in this case only those at $i_u = 4$ and $i_u = 5$) are calculated by mass conservation.

Unfortunately, the application of this idea is not straightforward. The presence of a block across a number of lines splits those lines into three sections. Before the block surface is reached from the west side the CELS line solver applies and this is also true between the end of the block and the east wall. This means that the solver can no longer sweep through the entire solution domain, irrespective of the flow to be calculated. It is necessary to build information about the length of the line sweeps into the solver routine for each block to be examined. When solving for the vertical component in this application there are a number of blocks to be calculated and the line is thus split into even more sections with the line solver applying between blocks. Below the block, as at line $j=3$ in figure 13.4, the line solution sweep continues as usual.

By applying this technique it is possible to calculate a two-dimensional flow, such as that shown in figure 13.4. This figure shows a square flow domain with a large block inserted. The result is realistic and agrees well with the same flow calculated using TEACH. It was found that convergence was only possible when the relaxation factors of the velocity component equations had values below 0.965. This resulted in the CELS solver requiring 32% more time to converge than the SIMPLE-based approach.
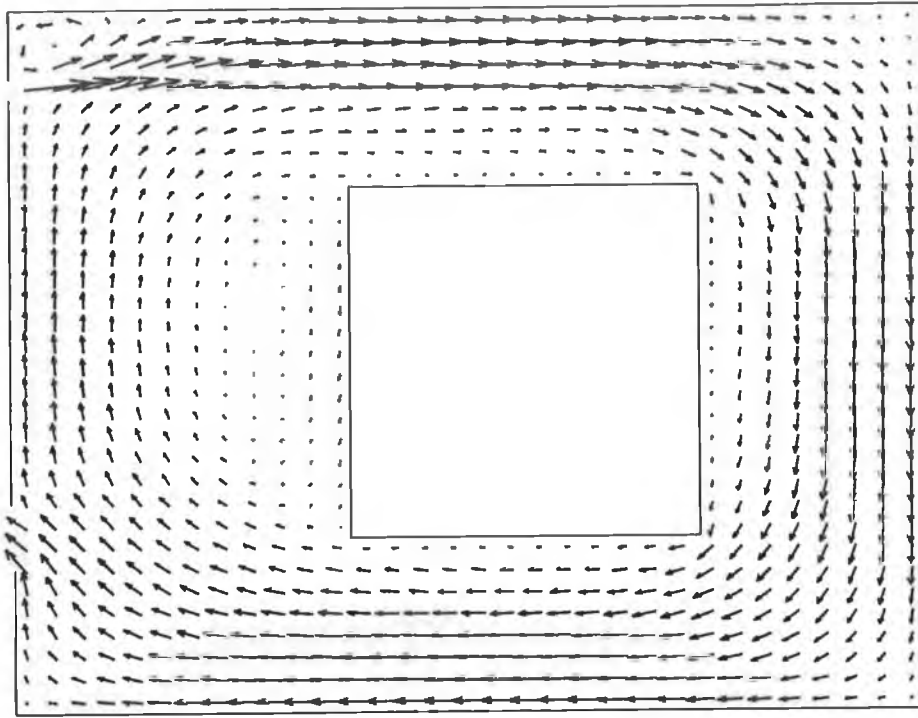
Figure 13.4. Flow around a block calculated using CELS with mass conservation applied to the upper and east faces of the block.
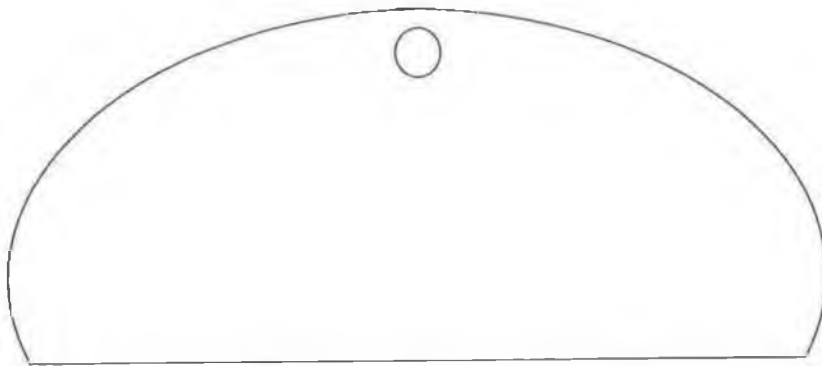
Figure 13.5. Outline drawing of a mushroom tunnel with in-curving walls.
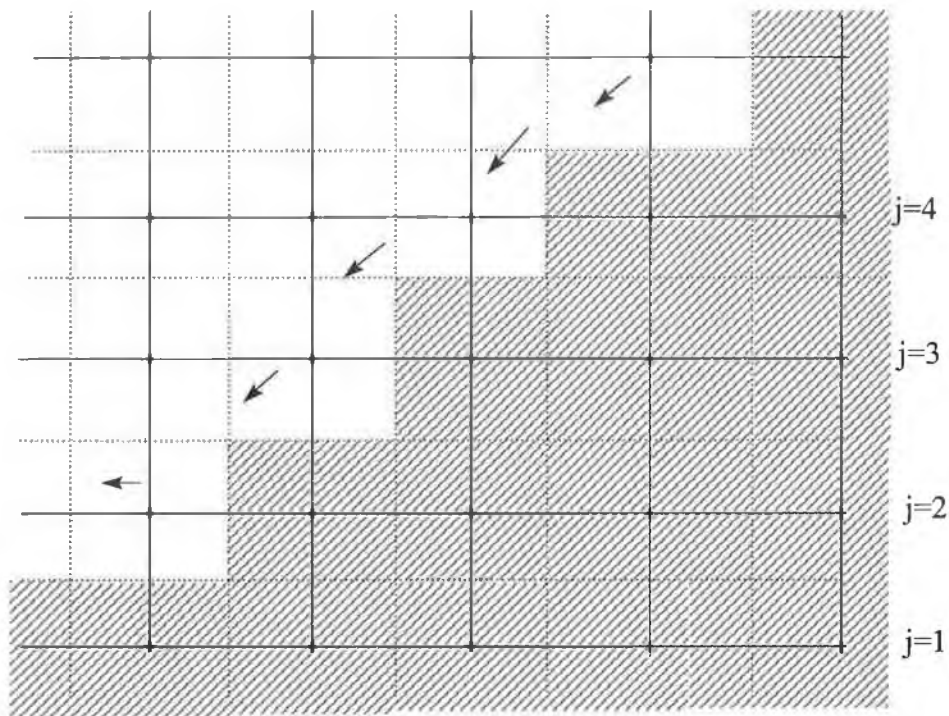
Figure13. 6. Coarse modelling of the base of an in-curving tunnel wall.

While it has been shown to be possible to arrange for the local mass balance calculation whenever an internal boundary is encountered the programming involved is extremely inconvenient. The boundaries encountered are a straightforward rectangular block that are relatively convenient for the operation of the line solvers in a rectangular grid. Consider the wall shape shown in figure 13.5. This could be modelled in a grid cell approximation by the sort of block layout shown in figure 13.6. In order to avoid the difficulty with the CELS solver, it would be necessary to implement the mass balance calculation at each step as the line solver moves down

through the lines j = 4 to j = 2. While possible this would require considerable programming effort in setting up new solver code.

A simple insertion technique, like that in the segregated solution technique which leaves the solver routines unchanged, is not possible and this is a serious drawback for the CELS technique.

13.6 Application of ACM

Further difficulties for the block insertion arise when applying the ACM technique. In the curved wall calculations, the presence of large numbers in the source terms causes problems when fine grid cells are combined to produce the coarse grid cells. The calculations presented in section 9.8 required a test to determine if the large numbers were present in any of the source terms that were to be combined. The source term for the coarse cell was then set to a large value.

This means that the original wall shape becomes 'blurred' as the solver telescopes down through a series of coarse grids. Information on the fine grid flow near the wall is being lost and this would lead to inaccuracies in the calculation of the corrections and therefore diminish their usefulness.

While the residual reduction factor, α, was set to a high value for the curved wall calculations in Section 9.8 and this effectively limited the usefulness of ACM to one coarse grid, it is probable that its usefulness is also limited by the increasing coarseness of the wall shape on coarser grids.

Assuming that block insertion becomes possible under CELS, the same approach to coarse grid transfers could be used by monitoring the presence of large values in the source terms for volumes to be included in the description of shelves. Coarse grid cells containing the modified cells on the fine grid would need to be modified but this

must result in a large shift in the geometry and loss of information as the grids telescope down to the coarsest to be used.

Thus CELS presents problems in the use of this convenient technique which are not straightforwardly overcome. Further work on this would be justified by the success of the extension of these methods to three-dimensional calculations.

13.7 Three-dimensional flows

A number of flow examples are included here as illustrations of the type of flow that would be required for a useful model of some of the mushroom growing structures but all of which contain internal boundaries that mean that CELS3D cannot, practically, be applied to their solution. The Reynolds number in each case was 3500 and turbulence was weak in the calculated flows. The overall geometry is the same in each case with the internal boundaries being different. The dimensions were set to approximate the actual dimensions of a mushroom growing room, i.e. 3 m square. The computational grid was $34 \times 34 \times 34$. The inlet slot was 1 grid unit high and 30 grid units in length. It was located on the west face of the flow domain at $j = 28$. The outlet was a square aperture on the back face of the flow domain. It was 3 grid units wide and 3 grid units high. It was located from $i = 30$ to $i = 32$ and from $j = 5$ to $j = 7$.

Such an example flow field resulting from a calculation using SIMPLE is shown in figure 13.7 where the shelves were inserted using the source term modification technique of section 4.5. QUICK differencing was used for the velocity equations. In spite of the coarse inlet approximation (as outlined in section 10.1), the general features of this flow have been observed qualitatively in mushroom growing tunnels. There is strong flow along the wall to the floor and then underneath the lowest shelf. Air speeds are relatively low above the lowest and middle shelves while the edge of the uppermost shelf intercepts some of the downward flow and gives a positive flow across the top of the shelves.
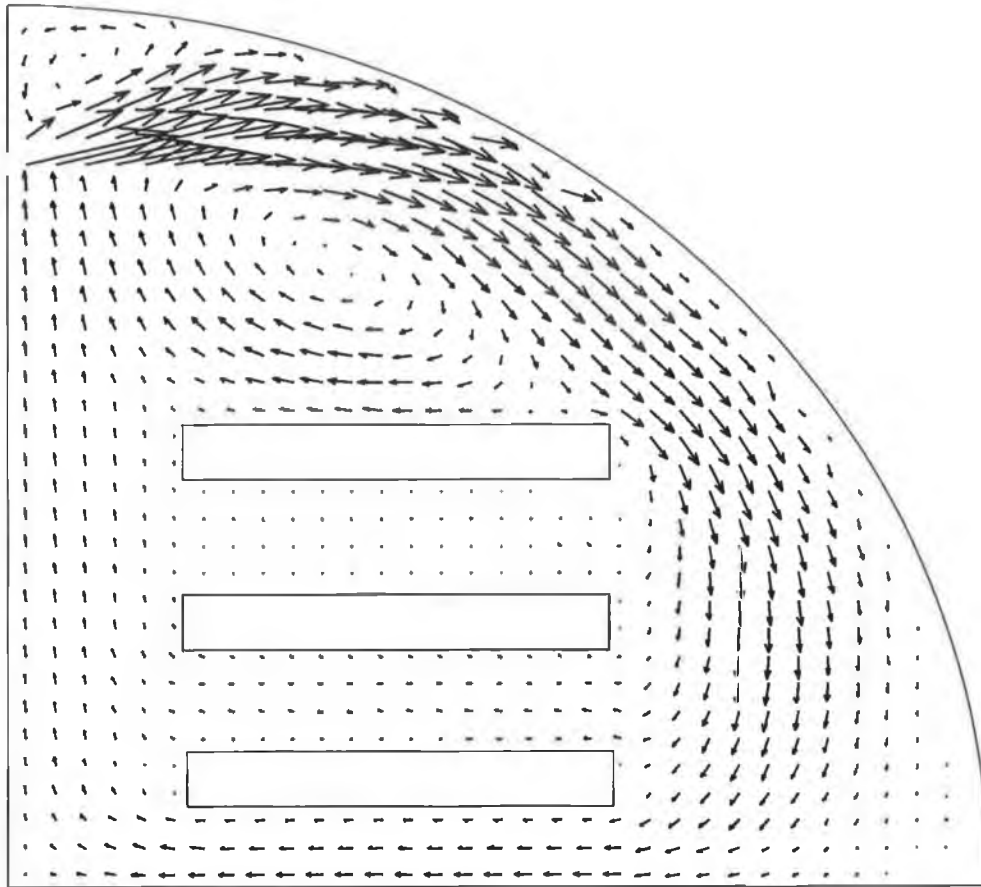
Figure 13.7. Air flow around shelf-like obstructions

With the air flow pattern shown in figure 13.7 there is a non-uniform flow across the three shelves. This can be improved by inserting deflectors to modify the flow pattern. For example, the effect of such a deflector mounted at a position between the middle and upper shelves is illustrated in figure 13.8. Comparing this to figure 13.7 it can be seen that the air speeds across the middle shelf have increased due to a portion of the flow being driven across it. The deflector serves to change the flow path so that the middle shelf can intercept a portion of the air stream.
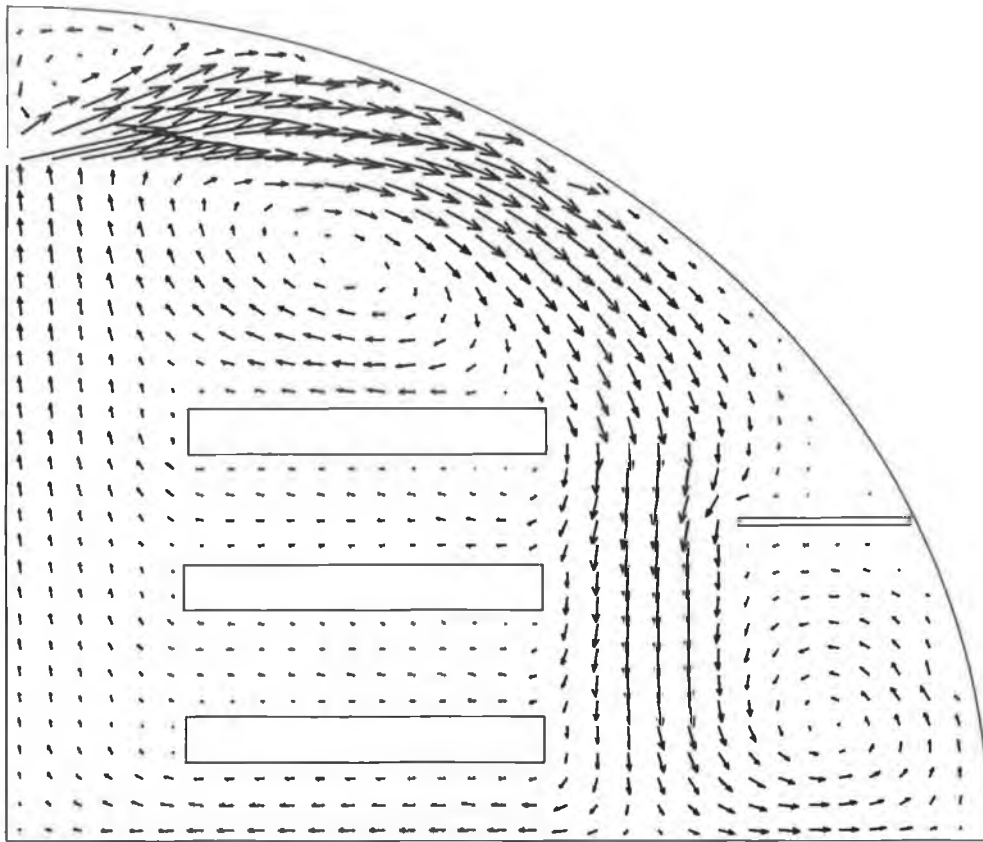
Figure 13.8. Air flow across shelves with a deflector (Re=3500).

Modelling of this type of airflow correction is an attractive alternative to experimentation because the measurements required (multiple averages of the air speed over periods of at least one minute, for each level, for a large number of deflector positions, dimensions and angles relative to the wall) would be tedious, particularly if multiple deflectors are required.

While the system shown in figure 13.8 is a very coarse approximation to the tunnel air flow and the shelf/deflector obstructions, and considerable grid refinement and improved modelling of the air inlets are likely to be required for accurate results, it demonstrates that the concept could work and that it would be worth pursuing the ideas involved.
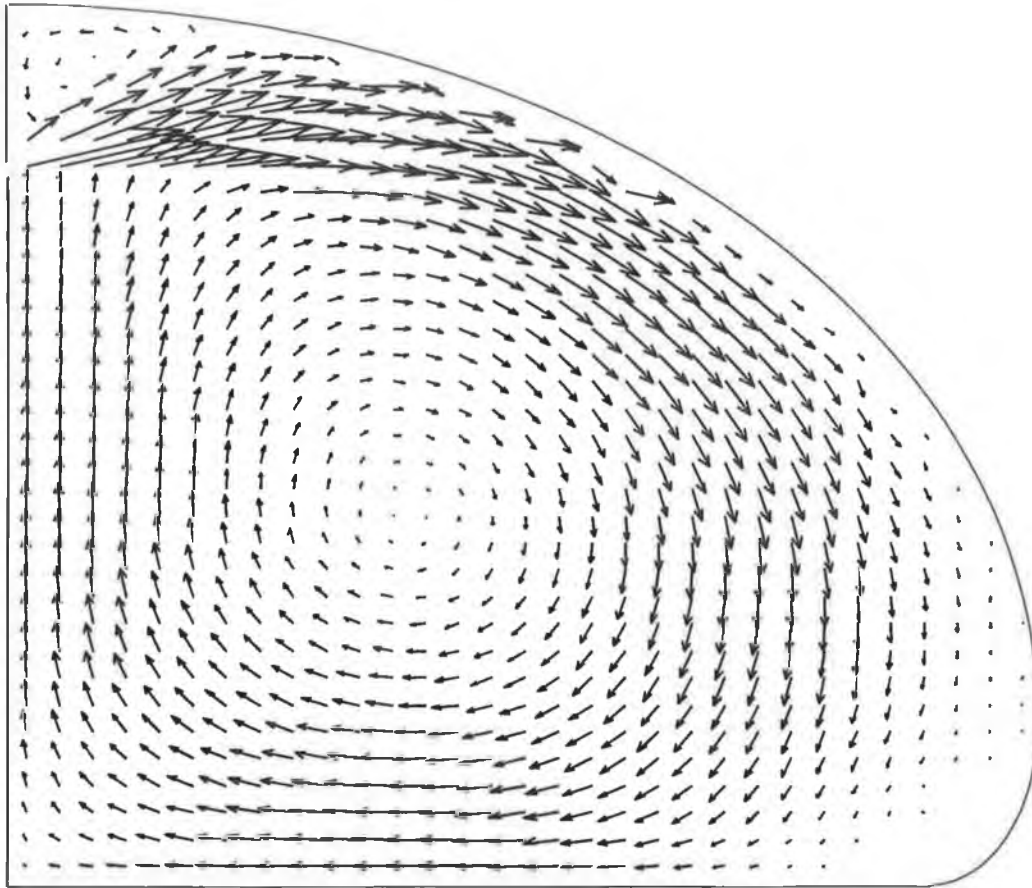
Figure 13.9. Air flow pattern in a structure with an in-curving wall base.

Figure 13.9 shows the result of a three-dimensional calculation for a tunnel with an in-curving base. The centre of the tunnel was set as an axis of symmetry and the wall shape was modified at the base as shown in figure 13.6. The calculation was carried out using the SIMPLE technique and the zero flow control volumes 'turned off' by modifying the source terms in the air velocity equations. Again, because of the position of the control volumes to be turned off it would have been impractical to use the CELS3D solver.

Chapter 14

Conclusions

The author's experimental investigation (see Chapter 2) of the airflow in mushroom growing tunnels has shown that there is much potential for improvement of the airflow at the cropping surface.

For the simplest case of single layer growing (bags of compost on the floor) the grower should be able to select a range of airspeeds with which the development of the crop can be manipulated. However, achieving good control will depend on having a fan speed controller with a smooth operating characteristic free from sudden surges in output over small portions of the range of operation. While this is lacking on most growing units the problem is easily solved with modern electronics. The more demanding aspect for the grower is the identification of the exact relationship between the output from the air distribution duct and the resultant cropping surface air speed. This will certainly have to be done for a given growing unit where all tunnels would have similar equipment and most likely for each individual tunnel, particularly if the relationship is to be between the cropping surface air speed and the fan speed controller rather than the duct exit air speed.

This can be done for isothermal conditions but it is very difficult, in setting a particular air speed, to take into account the effect of heating with its seasonal and, indeed, day/night variations. Ideally, the new two-duct system (see Section 2.6) should be used because it can eliminate to a great extent the air speed reductions due to heating. The retrofit solution developed in this work is easy to install but there are two requirements for its successful use.

The first is a continuation of the measurements that were begun in this work to identify the relationship between the output speeds from the distribution ducts and the resultant cropping surface air speed that would then be controlled by a combination of the output speeds. There should also be an examination of the possibilities for a re-design of the distribution ducts so that the high-speed contribution would come from the additional re-circulation duct with perhaps a fixed, high-volume flow rate through the duct for conditioned air.

The second requirement for the application of the two-duct system is that it requires dynamic control of the duct output speeds so that the system can react during heating but the manufacturers of control equipment for the industry do not currently provide this. The incorporation of a sensor for the low-speed cropping surface air flows in a fan speed control loop would be ideal but there would still be potential for improved airflow with open loop control when the operating characteristics of the two duct system are better understood.

The use of deflectors for three level shelf growing (see Section 2.7) have shown that straightforward, practical approaches can yield useful benefits but there is much more work to be done before a general recommendation for their use would be possible. The influence of the geometry of a given installation on their effectiveness has to be fully investigated (e.g. the shape of the tunnel and the distance between the wall and the shelves) as well as their interaction with other modifications to the air distribution duct or ducts.

These last factors in particular and the number of new types of multi-level growing systems make the development of an airflow model highly desirable. It will not be possible to move forward without the continuing use of empirical methods but the use of a model could shorten considerably the development time for a flow solution for each system. The empirical methods will be required for the validation/verification of the flow model and then to evaluate those options that modelling might show to be useful.

The initial development of the flow modelling was successful in that the CELS method was extended to three-dimensional flows and its application produced savings in the solution times for the test cases that were considered. However, while the two-dimensional results seemed encouraging for the extension of ACM3D to three-dimensional solution techniques for CELS3D, no savings in computational effort were found in its application. The combination of methods is required in order to make an attractive solver because the savings due to CELS3D alone, while significant, are not enough to justify moving away from the use of established solvers. The savings in computational effort could be attractive at low values of the Reynolds number but as this increases the savings are diminished. The second test problem highlighted the problem dependence of the savings. The alternative segregated methods that are available can be operated with FAS/FMG multigrid schemes (see Chapter 6) with great success and can produce substantial savings. The application of a multigrid scheme that operates only on inner iterations has attractions from the programming point of view and it is unfortunate that no such scheme was found to be effective.

On the basis of the results described in Sections 12.4 and 12.5 and the effort made in debugging and testing the code, it was concluded that syntactical and logical errors in the coding of ACM3D were unlikely and that the poor performance in giving savings relative to fine grid convergence times was the inability of the ACM3D technique as applied here to contribute to the efficiency of the variants of the three-dimensional solution techniques that were developed for CELS3D.

The testing was carried out for both laminar and turbulent flows with no particular differences observed. Although the turbulent solver is efficient at low values of CELS3D relaxation, the small number of iterations per coefficient update limits the potential of ACM as in inner solver for turbulent flows in three dimensions. There is insufficient effort per coefficient update cycle to allow savings to reach the same level as the two-dimensional case.

The problem with the application of ACM to the methods developed for CELS3D may be an analogue of the problem already identified for segregated solution methods. The efficient CELS3D methods developed do not require a large effort per plane set

302

before moving on to the next set. Also, the solution of, say, the *uv* set acts to some extent to offset the gains made in previous *vw* and *uw* calculations. A prompt return to the other sets is required in order to maximise the benefit of the updates in the variable values without driving them too far off the solution previously calculated for those planes. This is reflected in the relative inefficiency of variant 4 of the solution procedure (see Section 10.1). Good gains for a given plane set are not necessarily beneficial overall.

The low relaxation factor required for the use of ACM3D may be indicative of another aspect of the problems encountered. Galpin et al. (1985) attributed the success of this relaxation to its effect on the effective radius of influence of pressure on momentum (Zedan and Schneider, 1983). When significant under-relaxation is present on the equations the radius is small and the pressure influence is local. A solver in this situation is less sensitive to the approximation involved in fixing off-line or off-plane variables at their best estimates. The effective radius of influence increases as under-relaxation is removed (as it is completely for the best fine grid laminar solutions) and the iterative solver becomes very sensitive to off-line and off-plane estimates.

The failure of ACM3D may be due to the solution procedures adopted for three dimensions but there may also be more fundamental underlying causes of poor performance. For example, Brandt and Yavneh (1990) discuss difficulties in the application of multigrid techniques with upwind difference schemes for some re-circulating flows. They proved the inadequacy of first-order upwind differencing for some high Reynolds number two-dimensional flows. However, they freely acknowledge that such analyses are difficult to extend to more complex, coupled non-linear systems and that, for such cases, the problem itself can be difficult to identify. It is often stated that the failure of a multigrid method exposes fundamental problems with the manner in which the flow problem is posed on the fine grid.

It may be possible to develop a further alternative solver, based on the line solution procedure of the CELS method. While convergence was not achieved for another procedure in this thesis no point of principle was identified as being responsible for

303

this and it may be that the greater interactions between the variables during the solution would result in more potential for ACM3D applications.

It had been anticipated that the CELS3D method would be tested on a wider range of problems and grid densities but the time required for the investigation of and experimentation with ACM3D precluded this. This investigation was particularly time-consuming because of the amount of de-bugging and testing that was required for each variation in the application of the multigrid method. It would be very easy to miss the potential savings by using the wrong tuning for the solver and each variation required a tedious set of calculations over a range of values of all the tuning factors. This is not an unusual problem in computing. When a technique is working it is relatively easy to define the process but when there is a problem then a great many things could be wrong and it can take some time to apply logic and eliminate possibilities in trying to solve the problem for the problem.

On the use of the CELS method for the immediate modelling requirements of mushroom growing systems the failure to be able to apply the straightforward source term technique of Section 4.5 to insert obstructions in the flow puts CELS3D at a serious disadvantage relative to SIMPLE, in spite of useful savings for the flows in empty spaces. CELS3D should have a significant advantage in the modelling of buoyant flows in the Irish system relative to SIMPLE but the use of SIMPLEC with its competitive solution times is an alternative. SIMPLEC shares the advantage of SIMPLE in facilitating the modelling of shelves. Any major benefits of the use of CELS3D would require the successful implementation of another solution acceleration technique like ACM3D.

The accuracy of the CELS3D solver itself was investigated for a range of grid sizes and values of the Reynolds number. It was found in all cases that the accuracy increased as the grid was refined but further work is required for the three-dimensional case because the range of grid sizes was limited. No indication of difficulty with the resolution of the inlet flow was found in three dimensions but this may emerge as the dominant source of error at higher grid densities. This might be expected because the two-dimensional flows showed that, on the coarser grids, other

regions of the flow were more problematic but, as the grid density increased and the flow became better defined overall, then the largest differences from between the solutions on one fine grid and the next were associated with the inlet flow.

Included in the evaluation of solution accuracy was a comparison of difference schemes where the use of the QUICK scheme was found to be essential for low grid density modelling of re-circulating flow because it allows almost grid-independent solutions at lower density than the hybrid scheme and also provides an accuracy that can allow the testing of other aspects of flow modelling, free from the errors due to false/numerical diffusion that are associated with the use of the hybrid differencing scheme. It is reported (e.g. Leonard, 1988) that QUICK's shortcoming is that it can produce, under highly convective conditions, overshoots and possibly oscillations on each side of what should be a step in the dependent variable when convected at an angle oblique to the grid. While not implemented in this work, Leonard (1988) did develop a technique (SHARP) to overcome this behaviour without compromising the advantages of QUICK and with overall characteristics that were virtually insensitive to flow-to-grid angle and this could be incorporated in the code used in this work.

Development and validation of the model for air flows in mushroom tunnels will now resume. Grids are likely to be much finer than the densities used in this work when possible requirements for local grid refinement are taken into account. The use of CELS for modelling growing shelves will be re-assessed but it is likely that SIMPLEC will provide the most practical route to progress in this area.

References

Armfield, S.W., 1991. Finite difference solutions of the Navier-Stokes equations on staggered and non-staggered grids, Computers and Fluids, vol. 20, 1, 1-17.

Atkins, F.C., 1965a. How many air changes needed? The Grower, October 16, 609

Atkins, F.C., 1965b. Warmed air intake solves oxygen problem, The Grower, December 18, 973

Bailey, B.J., 1982. The design of film-plastic ventilation ducts, National Institute of Agricultural Engineering Report R. 39, Silsoe, Bedford, UK

Baker., C.J., 1994. Aerodynamics of poultry transporters - implications for environmental control, Worlds poultry science journal, vol. 50:1, 62-63

Baker, A.J., Roy, S., Kelso, R.M., 1994. CFD experiment characterisation of airborne contaminant transport for 2 practical 3-D room air flow fields, Building and Environment, vol. 29, no. 3, 253-259

Barber, E.M. and Ogilvie, J.R., 1982. Incomplete mixing in ventilated airspaces. Part I. Theoretical considerations. Can. Agric. Eng., vol. 24, 25-29

Barber, E.M. and Ogilvie, J.R., 1984. Incomplete mixing in ventilated airspaces. Part II. Scale model study, Can. Agric. Eng., vol. 26, 189-196

Barber, W.H. and Summerfield, L.R., 1989. Environmental control of bacterial blotch on Pennsylvania shelf farms, Mushroom Journal, vol. 201, 285-296

Barcus, M., Peric, M., Scheuerer, G., 1987. A control-volume based full multigrid procedure for the prediction of two-dimensional laminar, incompressible flows, Proc. 7th GAMM conference on numerical methods in fluid dynamics, Sept. 9-11, 1987, Louvain-la-Neuve, Belgium, 9-16

Barton, I.E., 1995. A numerical study of flow over a confined, backward-facing step. International Journal for Numerical Methods in Fluids, vol. 21, 653-665

Berckmans, D., Randall, J.M., Van Thielen, D., Goedseels, V., 1993. Validity of the Archimedes number in ventilating commercial livestock buildings, J. Agric. Engng. Res., 56, 239-251

Bishop, C.F.H., 1979. Air in the growing room, Mushroom Journal, no. 73, 464-466

Boulard, T., Roy, J.C., Lamrani, M.A., Haxaire, R., 1997. Characterising and modelling

the air flow and temperature profiles in a closed greenhouse in diurnal conditions, Mathematical and Control Applications in Agriculture and Horticulture (IFAC Workshop, Hannover, Germany), 37-42

Bowman, G.E., 1987. Air circulation in mushroom houses. Mushroom Journal, no. 173, 151-169

Bowman, G.E., 1991. Application of multiple jets to the supply and distribution of conditioned air in mushroom cropping houses, J. Agric. Engng. Res., vol. 49, 151-159

Bozeman, J.D. and Dalton, C. 1973. Numerical study of viscous flow in a cavity, Journal of Computational Physics, vol. 12, 348-362

Brandt, A. 1984. Multigrid techniques: 1984 guide with applications to fluid dynamics. Weizman Institute of Science, Rehovot, Israel

Brandt, A. 1987. Multi-level computations: review and recent developments, Third Copper Mountain Conference on Multigrid Methods, University of Colorado, Denver

Brandt, A., and Yavneh, I. 1990. Inadequacy of first order upwind difference schemes for some recirculating flows. Internal report, The Weizmann Insitute of Science, Rehovot, Israel

Briggs, William L., 1987. A multigrid tutorial, Society for Industrial and Applied Mathematics, Philadelphia, Pennsylvania.

Burggraf, O.R., 1966. Analytical and numerical studies of the structure of steady, separated flows, Jnl. of Fluid Mechanics, vol. 24, 113-128

Burrage, S.W., Varley, M.J., Noble, R., Perrin, P.S.G., 1988. A microcomputer-based environmental control system for mushroom cropping tunnels, Computers and Electronics in Agriculture, vol. 2, 193-207.

Caretto, L.S., Gosman, A.D., Patankar, S.V., Spalding, D.B. 1972. Two calculation procedures for steady, three-dimensional flows with re-circulation, Proc. 3rd Int. Conf. Num. Methods in Fluid Dynamics, Paris, vol. 2, 60-68

Caretto, L.S., Curr, R.M., Spalding. D.B. 1972a. Two numerical methods for three-dimensional boundary layers, Computer Methods in Applied Mechanics and Engineering, vol. 1, 39-57

Chapman, M, 1981. FRAM-nonlinear damping algorithms for the continuity equation, J. Comp. Phys., vol. 44, 84-103

Chen, Y. and Falconer, R.A., 1992. Advection-diffusion modelling using the modified QUICK scheme, International Journal for Numerical Methods in Fluids, vol. 15, 1171-1196

Cheng, S.I. and Shubin, G., 1978. Computational accuracy and mesh Reynolds number, Jnl. Comp. Physics, vol. 28, 315-326

Choi, H.L., Albright, L.D., Timmons, M.B., Warhaft, W. 1988. An application of the k-ε turbulence model to predict air distribution in a slot-ventilated enclosure, Transactions of the ASAE, vol. 31, no. 6, 1804-1813

Choi, H.L., Albright, L.D., Timmons, M.B. 1990. An application of the k-ε turbulence model to predict how a rectangular obstacle in a slot-ventilated enclosure affects air flow, Transactions of the ASAE, vol. 33, no. 1, 274-281

Choi, H.L 1991. Private communication

Connell, S.D. and Stow, P., 1986. The pressure correction method, Computers and Fluids, vol. 14, no. 1, 1-10

Connolly, M., 1999. Private communication

Courant, R., Isaacson, E., Rees, M., 1952. On the solution of non-linear hyperbolic differential equations by finite differences, Comm. Pure Appl. Math., vol. 5, 243

Davidson, L. 1996. A pressure correction method for unstructured meshes with arbitrary control volumes, International Journal for numerical methods in fluids, vol. 22, 265-281

Demirdzic, I, Gosman, A.D., Issa, R.I., Peric, M., 1987. A calculation procedure for turbulent flow in complex geometries, Computers and Fluids, vol. 15, no. 3, 251-273

Demuren. A.O., 1985. False Diffusion in three-dimensional flow calculations, Computers and Fluids, vol. 13, no. 4, 411-419

De Praetere, K., and Van Der Biest, W., 1990. Airflow patterns in piggeries with fully slatted floors and their effect on ammonia distribution, J. Agric. Engng. Res., vol. 46, 31-44.

Deurloo, J.A., Feddes, J.J.R., Leonard, J.J., Darby, D.E., 1990. Effect of reciculated air on air speeds at animal level: Commercial-scale swine barn tests, Can. Agric. Eng., vol. 33, 179-183.

De Vahl Davis, G. and Mallinson, G.D., 1976. An evaluation of upwind and central difference approximations by a study of re-circulating flow, Computers and Fluids, vol. 4, 29-43

Van Doormal, J.P. and Raithby, G.D. 1984. Enhancements of the SIMPLE method for predicting incompressible fluid flows, Numerical Heat Transfer, vol. 7, 147-163

Edwards, R., 1973a. Mushroom house ventilation in theory and practice (part 1), Mushroom Journal, no. 3, 118-131

Edwards, R., 1973b. Mushroom house ventilation in theory and practice (part 2), Mushroom Journal, no. 4, 166-174

Farrell, P.A., Hegarty, A.F., Miller, J.J.H., O'Riordan, E., Shishkin, G.I. 2000. Robust

computational techniques for boundary layers, Applied Mathematics 16, Chapman & Hall/CRC

Flegg, P.B., 1974. The measurement of evaporative loss in relation to water management during cropping of A. bisporus, Mushroom Science IX (Part 1), Proceedings of the Ninth International Scientific Congress on the Cultivation of Edible Fungi, Tokyo, 285-292

Freitas, C.J., Street, R.L., Finidikakis, A. N., Koseff, J. R., 1985. Numerical simulation of three-dimensional flow in a cavity, International Journal for Numerical Methods in Fluids, vol. 5, 561-575

Galpin, P.F., Van Doormal, J.P., Raithby, G.D. 1985 Solution of the incompressible mass and momentum equations by application of a coupled equation line solver, International Journal for Numerical Methods in Fluids, vol. 5, 615-625

Galpin, P.F. and Raithby, G.D. 1986. Numerical solution of problems in incompressible fluid flow: treatment of the temperature-velocity coupling, Numerial Heat Transfer, vol. 10, 105-129.

Gaskell, P.H., Lau, A.K.C., Wright, N.G. 1988. Comparison of two strategies for use with higher-order discretization schemes in fluid flow simulation, International Journal for Numerical Methods in Fluids, vol. 8, 1203-1215.

Ghia, U, Ghia, K.N., Shin, C.T., 1982. High-Re solutions for incompressible flow using the navier-stokes equations and a multigrid method, Journal of Computational Physics, vol. 48, no. 3, 387-411

Gosman, A. D. and Ideriah, F.J.K. 1976, Teach -2E: a general computer program for two-dimensional, turbulent, re-circulating flows, Department of Mechanical Engineering, Imperial College, London.

Grant, J. 1991. A mathematical model of the environment in mushroom tunnels, Mushroom Science XIII, Science and Cultivation of Edible Fungi, vol. 1, 299-305

Grant, J. 1995. Air conditions in mushroom tunnels - the effects of heating and cooling systems, Proceedings of the 11[th] National Mushroom Conference, 27-36, Teagasc, Dublin

Grant, J. and Williams, D. 1996. Modelling air flow in Irish mushroom tunnels, Proc. Mathematical and Control Applications in Agriculture and Horticulture, Acta Horticulturae 406, 399-405

Gresho, P.M. and Lee, R.L., 1981. Don't suppress the wiggles - they're telling you something! Computers and Fluids, vol. 9, 223-253

Han, T., Humphrey, J.A.C., Launder, B.E., 1981. A comparison of hybrid and quadratic upstream differencing in high Reynolds number elliptic flows, Computer Methods in Applied Mechanics and Engineering, vol. 29, 81-95

Harral, B.B. and Boon, C.R., 1997. Comparison of predicted and measured airflow patterns in a mechanically ventilated building without animals. Journal of Agricultural Engineering, vol. 66, no. 3, 221-228

Harlow, F.H. and Welch, J.E., 1965. Numerical calculation of time-dependent viscous incompressible flow of fluid with free surface, The Physics of Fluids, vol. 8, 2183-2189

Hayase, T., Humphrey, J.A.C., Greif, R., 1992. A consistently formulated QUICK scheme for fast and stable convergence using finite-volume iterative calculation procedures, Journal of Computational Physics, vol. 98, 108-118

Hayes, R. 1991. The development of a computer simulation model of a mushroom-growing tunnel to investigate environmental control, Mushroom Science XIII, Science and Cultivation of Edible Fungi, vol. 1, 307-313

Heber, A.J., Boon, C.R., Peugh, M.W., 1996. Air patterns and turbulence in an experimental livestock building, Journal of Agricultural Engineering Research, 64, 3, 209-226

Hjertager, B.H. and Magnussen, B.F., 1981. Calculation of turbulent three-dimensional jet induced flow in rectangular enclosures, Computers and Fluids, vol. 9, 395-407

Hoff, S.J., Janni, K.A., Jacobson, L.D., 1992. 3-Dimensional buoyant turbulent flows in a scaled model, slot-ventilated, livestock confinement facility, Transactions of the ASAE, vol. 35, no. 2, 671-686.

Hoff, S.J., 1995. A simplified turbulence mode for describing airflow in ceiling slot-ventilated enclosures, Transactions of the ASAE, vol. 38, no. 6, 1853-1862

Hoff, S.J., Janni, K.A., Jacobson, L.D., 1995. Evaluating the performance of a low Reynolds number turbulence model for describing mixed-flow air speed and temperature distributions, Transactions of the ASAE, vol. 38, no. 5, 1533-1541.

Huang, P.G., Launder, B.E., Leschziner, M.A., 1985. Discretisation of non-linear convection processes, a broad-range comparison of four schemes, Computer Methods in Applied Mechanics and Engineering, vol. 48, 1-24

Huh, K.Y., Golay, M.W., Manno, V.P., 1986. A method for reduction of numerical diffusion in the donor cell treatment of convection, Journal of Computational Physics, vol. 63, no. 1, 201-221

Hutchinson, B.R., Galpin, P.F., Raithby, G.D. 1988. Application of additive correction multigrid to the coupled fluid flow equations, Numerical Heat Transfer, vol. 13, 133-147

Hutchinson, B.R., and Raithby, G.D. 1986. A multigrid method based on the additive correction strategy. Numerical Heat Transfer, vol. 9, 511-537

Ikahogi, T., Shin, B.R., Daiguji, H., 1992. Application of an implicit time-marching scheme to a three-dimensional incompressible flow problem in curvilinear co-ordinate systems, Computers and Fluids, vol. 21, no. 2, 163-175

Issa, R.I., 1985. Solution of the implicitly discretised fluid flow equations by operator-splitting, Journal of Computational Physics, vol. 62, no. 1, 40-65

Khosla,.P.K. and Rubin, S.G., 1974. A diagonally dominant second-order accurate implicit scheme, Computers and fluids, vol. 2, 207-209

Latimer, B.R. and Pollard, A., 1985. Comparison of pressure-velocity coupling solution algorithms, Numerical Heat Transfer, vol. 8, 635-652.

Launder, B.E. and Spalding, D.B. 1972. Mathematical models of turbulence, Academic Press, London

Launder, B.E. and Spalding, D.B. 1974, The numerical computation of turbulent flows. Computer Methods in Applied Mechanics and Engineering, vol. 3, 269-289

Le Feuvre, R.F. 1978. The prediction of 2-dimensional re-circulating flows using a simple finite difference grid for non-rectangular flow fields. Computers and Fluids, vol. 6, 203-218

Leonard, B.P., 1979. A stable and accurate convective modelling procedure based on quadratic upstream interpolation, Computer Methods in Applied Mechanics and Engineering, vol. 19, 59-98

Leonard, B.P., 1988. Simple, high accuracy resolution program for convective modelling of discontinuities, International Journal for Numerical Methods in Fluids, vol. 8, 1291-1318

Leschziner, M.A., 1980. Practical evaluation of three finite difference schemes for the computation of steady-state recirculating flows, Computer Methods in Applied Mechanics and Engineering, vol. 23, 293-312

Leschziner, M.A. and Rodi, W., 1981. Calculation of annular and twin parallel jets using various discretization schemes and turbulence model variations. J. Fluids Engineering, vol. 103, 352-360

Leschziner, M.A. and Rodi, W., 1984. Computation of strongly swirling axisymmetric free jets. AAIA Journal, vol. 22, 1742-1747

Li, Y. and Baldacchino, L., 1995. Implementation of some higher-order convection schemes on non-uniform grids. International Journal for Numerical Methods in Fluids, vol. 21, 1201-1220

Li, Y. and Rudman, M., 1995. Assessment of higher-order upwind schemes incorporating FCT for linear and non-linear convection-dominated problems. Numerical Heat Transfer, B, vol. 27, 1-21

Liu, Q., Hoff, S.J., Maxwell, G.M., Bundy, D.S., 1996, Comparison of three k-ε turbulence models for predicting ventilation air jets. Transactions of the ASAE, vol. 39, no. 2, 689-698

Loeffen, H., 1991. Private communication

Loeffen, H., 1992. Air distribution in mushroom growing rooms. Mushroom Journal, no. 505, 13-14

Loeffen, H., 1993. Extraneous ventilation of growing room because of leakage. Mushroom Journal, no. 517, 15-16

Loeffen, H., 1994. Carbon dioxide production rate as a measure of growth of mushroom crops. Mushroom Journal, no. 540, 28-29

Loeffen, H., 1995. The influence of compost temperature on the activity of mushroom substrate. Mushroom Journal, no. 548, 28-29

Lomax, K., 1993. Private communication

Lomax, K., Gottfried, S., Lavelle, H., 1995. Air flow indicators for mushroom farms. Journal of Agricultural Engineering, vol. 60, 43-48

Lomax, K., Beyer, D., Rhodes, T., 1996. Air flow effects on mushroom production. Proceedings of the 2nd International Conference on Mushroom Biology and Mushroom Products, Penn State University. 233-240

Maghirang, R.G. and Manbeck, H.B. 1993, Modelling particle transport in slot-ventilated air-spaces. Transactions of the ASAE, vol. 36, no. 5, 1449-1459

Maghirang, R.G., Manbeck, H.B., Puri, V.M., 1994. Numerical simulation of particle transport in slot-ventilated airspaces, Transactions of the ASAE, vol. 37, no. 5, 1607-1612.

McCormick, S.F. 1989, Multilevel adaptive methods for partial differential equations. Society for Industrial and Applied Mathematics, Philadelphia

McGuirk, J.J. and Palma, J.M.L.M. 1993. The efficiency of alternative pressure-correction formulations for incompressible turbulent flow problems. Computers and Fluids, vol. 22, no. 1, 77-87

McGuirk, J.J. and Rodi, W. 1978. A depth-averaged mathematical model for near field of side discharges into open-channel fluid flow, J. Fluid Mechanics, vol. 68, 761-781

Meath, F.P. 1993 Environmental control in mushroom growing tunnels. MSc Thesis, Dublin Institute of Technology, Ireland

Miller, T.F. and Schmidt, F.W., 1988. Evaluation of multilevel technique applied to the Poisson and Navier-Stokes equations, Numerical Heat Transfer, vol. 13, 1-26

Mulder, W.A., 1985. Multigrid relaxation for the euler equations. Journal of Computational Physics, vol. 60, no. 2, 235-252

Murray, F. 1995. Modelling and control of the environmental conditions in a muhroom tunnel. M. Eng. Sc. Thesis, University College Galway, Ireland

Murray, F., Corcoran, P., Nolan, P.J. and Grant, J., 1995. Design of an environmental control system for use in a mushroom tunnel complex, Proc. 12[th] Conf. Irish Manufacturing Committee, 974-981

Nallasamy, M., 1987. Turbulence models and their applications to the prediction of internal flows: a review, Computers and Fluids, vol. 15, no. 2, 151-194

O'Flaherty, T., 1988. Private communication

O'Flaherty, T., 1990. Plastic structures. Crop production under plastic, Teagasc, Kinsealy Research Centre, Ireland

Pan, F. and Acrivos, A., 1967, Steady flows in rectangular cavities, Jnl. of Fluid Mechanics, vol. 28, 643-664

Parsons, D.J., 1991. Modelling gas flow in a silage clamp after opening. J. Agric. Engng. Res., vol. 50, 209-218.

Patankar. S.V. and Spalding, D.B. 1972. A calculation procedure for heat, mass and momentum transfer in three-dimensional parabolic flows. International Journal of Heat and Mass Transfer, vol. 15, 1787-1806

Patankar, S.V. 1975. Numerical prediction of three-dimensional flows. Studies in convection: theory, measurement and applications, vol. 1 (ed. B.E. Launder), Academic Press, New York

Patankar, S.V., Basu, D.K., Alpay, S.A. 1977. Prediction of the three-dimensional velocity field of a deflected turbulent jet. Transactions of the ASME, December, 758-762

Patankar, S.V., Liu, C.H., Sparrow, E.M. 1977. Fully developed flow and heat transfer in ducts having steamwise-periodic variations of cross-sectional area. Transactions of the ASME, vol. 99, 180-186

Patankar, S.V., Ramadhyani, S., Sparrow, E.M. 1978. Effect of circumferentially non-uniform heating on laminar combined convection in a horizontal tube. Journal of Heat Transfer, vol. 100, 63-70

Patankar, S.V., Sparrow, E.M., Ivanovic, M. 1978. Thermal interactions among the confining walls of a turbulent recirculating flow. International Journal of Heat and Mass Transfer, vol. 21, 269-274

Patankar, S.V., Ivanovic, M., Sparrow, E.M. 1979. Analysis of turbulent flow and heat transfer in internally finned tubes and annuli. Journal of Heat Transfer, vol. 101, 29-37

Patankar, S.V. 1980. Numerical heat transfer and fluid flow. McGraw-Hill, London

Patankar, S.V., 1988. Recent developments in computational heat transfer, Journal of heat transfer, vol. 110, 1037-1045.

Patel, V.C., Rodi, W., Scheuerer, G., 1985. Turbulence models for near-wall and low Reynolds number flows: a review. AIAA Journal, vol. 23, no. 9, 1308-1319

Patel, M.K. and Markatos, N.C., 1986. An evaluation of eight discretization schemes for two-dimensional convection-diffusion equations. International Journal for Numerical Methods in Fluids, vol. 6, 129-154

Patel, M., Cross, M., Markatos, N.C., Mace, A.C.H., 1987. An evaluation of eleven discretization schemes for predicting elliptic flow and heat transfer in supersonic jets. International Journal of Heat and Mass Transfer, vol. 30, no. 9, 1907-1925

Peric. M., Kessler, R., Scheurer, G. 1988. Comparison of finite-volume numerical methods with staggered and co-located grids. Computers and Fluids, vol. 16, no. 4, 389-403

Philips, R.E. and Schmidt, F.W., 1984. Multigrid techniques for the numerical solution of the diffusion equation, Numerical Heat Transfer, vol. 7, 251-268.

Philips, R.E. and Schmidt, F.W., 1985a. Multigrid techniques for the solution of the passive scalar advection-diffusion equation, Numerical Heat transfer, vol. 8, 25-43.

Phillips, R.E. and Schmidt, F.W. 1985b. A multi-level multigrid technique for recirculating flows. Numerical Heat Transfer, vol. 8, 573-594

Pollard, A and Siu, A. L.-W., 1982. The calculation of some laminar flows using various discretization schemes. Computer Methods in Applied Mechanics and Engineering, vol. 35, 293-313

Rahnema, M., Yaghoubi, M., Sabzervari, A., 1996. Turbulent wind flow patterns and pressure fields across a tall building. Wind Engineering, vol. 20, no. 4, 241-257.

Raithby, G.D., 1976. Skew upstream differencing schemes for problems involving fluid flow. Computer Methods in Applied Mechanics and Engineering, vol. 9, 151-162

Raithby, G.D. and Schneider, G.E., 1979. Numerical solution of problems in incompressible fluid flow: treatment of the velocity-pressure coupling, Numerical Heat Transfer, vol. 2, 417-440

Raithby, G.D. and Torrance, K.E. 1974. Upstream-weighted differencing schemes and their application to elliptic problems involving fluid flow. Computers and Fluids, vol. 2, 191-206

Randall, J., and Battams, V.A., 1979. Stability criteria for air flow patterns in livestock buildings. Journal of Agricultural Engineering Research, vol. 24, 361-374

Randall, J.M. and Moulsley, L.J., 1990a. Propellor fans for ventilating livestock

buildings: 1. Measurement of performance. Journal of Agricultural Engineering Research, vol. 47, 89-99.

Randall, J.M. and Moulsley, L.J., 1990b. Propellor fans for ventilating livestock buildings: 3. Performance criteria, Journal of Agricultural Engineering Research, vol. 47, 101-113.

Randall, J.M. and Moulsley, L.J., 1990c. Propellor fans for ventilating livestock buildings: 3. Selection for least cost, Journal of Agricultural Engineering Research, vol. 47, 114-122.

Reynolds, A.M., 1997. A model for predicting airborne dust concentrations within a ventilated air space, Journal of Agricultural Engineering Research, vol. 66, no. 2, 103-109.

Rodi, W. 1980, Turbulence models and their application in hydraulics, International Association for Hydraulic Research, Delft, Netherlands

Settari, A. and Aziz, K. 1973. A generalization of the additive correction methods for the iterative solution of matrix equations. SIAM Journal of Numerical Analysis, vol. 10, 506-521

Schroeder, M.E., 1968. Forced ventilation with proportional control for the standard double mushroom house. Mushroom Science VII, 421-428

Schroeder, M.E., Schisler, L.C., Snetsinger, R., Crowley, V.E., Barr, W.L., 1974. Automatic control of mushroom ventilation after casing and through production by sampling CO_2. Mushroom Science IX, 269-278

Shyy, W., 1985. A study of finite difference approximations to steady-state, convection-dominated flow problems. Journal of Computational Physics, vol. 57, no. 3, 415-438

Shyy, W., Tong. S.S., Correa, S.M., 1985. Numerical recirculating flow calculation using body-fitted coordinate system. Numerical Heat transfer, vol. 8, 99-113.

Shyy, W. and Sun, C-S, 1993, Development of a pressure-correction staggered-grid based multigrid scheme for incompressible re-circulating flows. Computers and Fluids, vol. 22, no. 1, 51-76

Smith, D.M. 1990. Fluid dynamic calculations using the multigrid approach. PhD Thesis, University of London

Spalding, D.B. 1972. A novel finite difference formulation for differential expressions involving both first and second derivatives, Int. Journal of Numerical Methods in Engineering, vol. 4, 551.

Sparrow, E.M., Patankar, S.V., Ramadhyani, S. 1977. Analysis of melting in the presence of natural convection in the melt region. Transactions of the ASME, vol. 99, 520-526

Storey, I.F., 1968. Recent trends in crop production methods as affected by environmental conditions. Mushroom Science VIII, 95-102

Stubley, G.D., Raithby, G.D., Strong, A.B., 1980. Proposal for a new discrete method based on an assessment of discretization errors, Numerical Heat Transfer, vol. 3, 411-428.

Stubley, G.D., Raithby, G.D., Strong, A.B., Woolner, K.A., 1982. Simulation of convection and diffusion processes by standard finite difference schemes and by infuence schemes. Computer methods in applied mechanics and engineering, vol. 35, 153-168

Tennekes, H. and Lumley, J.L., 1972. A first course in turbulence. The MIT Press, Cambridge, MA

Tritton, D.J. 1977. Physical fluid dynamics, Van Nostrand Reinhold, Wokingham, England.

Tschierpe, H.J., 1973a. Environmental factors and mushroom growing, part1. Mushroom Journal, no. 1, 30-45

Tschierpe, H.J., 1973b. Environmental factors and mushroom growing, part1. Mushroom Journal, no. 2, 77-94

Vanka, S.P., 1986a. Block-implicit multigrid solution of navier-stokes equations in primitive variables. Journal of Computational Physics, vol. 65, no. 1, 138-158

Vanka, S.P., 1986b. A calculation procedure for three-dimensional steady recirculating flows using multigrid methods, Computer methods in applied mechanics and engineering, vol. 55, 321-338.

Vanklooster, C.E., Bontsema, J., Salomons, L., 1995. Dynamic model to tune a climate control algorithm in pig houses with natural ventilation, Transactions of the ASAE, vol. 38, no. 3, 911-918.

van Soest, G.J.A., 1979. Possibilties of ventilation and cooling in mushroom growing. Mushroom Journal, vol. 77, 189-191

Wanik, A. and Schnell, U., 1989. Some remarks on the PISO and SIMPLE algorithms for the steady turbulent flow problems. Computers and Fluids, vol. 17, no. 4, 555-570

Worley, M.S. and Manbeck, H.B., 1995. Modelling particle transport and air flow in ceiling inlet ventilation systems, Transactions of the ASAE, vol. 38, 1, 231-239.

Zhang, Y., Barber, E.M., Sokhansanj, S., 1992. A model of the dynamic thermal environment in livestock buildings, J. Agric. Engng. Res. vol. 53, 103-122.

Zhang, G., Morsing, S., Strom, J.S., 1996. Modelling jet drop distances for control of a non-isothermal, flap-adjusted ventilation jet, Transactions of the ASAE, vol. 39, no. 4, 1421-1431.

Zhu, J. and Rodi, W., 1991. Zonal finite-volume computations of incompressible flows. Computers and Fluids, vol. 20, no.4, 411-420