#5704

Getting Started with CFD

N S J FAWCETT, DPHIL, BSc
BHR Group Ltd, Cranfield, Bedford, UK

Synopsis

The recent developments in low cost, high power computing have made the use of computational techniques for the solution of complex industrial flow problems a reality. The market for general purpose Computational Fluid Dynamics (CFD) programs has grown as a result and is extremely competitive. CFD still requires a large investment (both money and time) and therefore it is essential that newcomers familiarize themselves with the methods, the limitations and the jargon used if they are to avoid a costly mistake. This paper describes in outline the above issues and gives some guidelines to choosing hardware and software.

1. Introduction

Computational Fluid Dynamics (CFD), the simulation of fluid flow and heat transfer processes, is one of the emerging numerical techniques which is set to revolutionize design engineering.

The use of computational techniques for design are becoming more commonplace in most sectors of industry. CFD has lagged behind other computer aided design (CAD) methods, such as mechanical stress analysis packages, due to the complexity of the underlying equations that describe fluid flow. Large supercomputers, which were (and still are) well out of the reach of most companies budgets, were required to solve 'industrial' problems. But with the recent advances in low cost, high power computing, and developments in numerical methods, the use of CFD as a practical and cost-effective design tool is becoming a reality.

Over the past ten years, the market for 'off-the-shelf' general purpose CFD packages has become fiercely competitive. To the inexperienced the codes seem to be much-of-a-muchness. Moreover, like other branches of science, CFD has its own jargon. This too can be confusing to the potential buyer of CFD packages.

Although the cost of CFD, in terms of hardware and software, has come down significantly over the past few years, it still requires a large investment. Potential buyers must make the right decision about which CFD software to buy first time.

CFD methods are useful across a broad spectrum of industries and disciplines and have much to offer the engineer in design, production or R&D, but the subject is still evolving and is by no means comprehensive as yet. To understand its capabilities and limitations, the potential user of CFD must appreciate the methods it uses, and become familiar with the jargon. Only then can the right questions be asked after the sales patter has ended.

2. Methodology

CFD programs use numerical methods to solve the basic equations describing the conservation of mass, momentum and heat in fluids.

First, the volume of interest (such as a room) is divided up into a large number of small cells (known as 'the grid'). The generation of the grid is perhaps the most important stage of setting up a CFD simulation taking typically 80% of the total effort. This is because the number and distribution of grid cells can effect whether a solution is obtained, the speed at which it is obtained and the accuracy of the simulation. With the aid of fluid physical property information and the boundary conditions (eg. inlet/outlets, flow rates etc.) the conservation equations are solved yielding typically three components of the velocity, pressure and temperature for each cell in the grid.

The solution is found by iteration. That is, the values of all the variables (velocity, pressure, chemical species etc.) are guessed. These values are then updated by feeding them into the very equations which you are trying to solve. If the updated values are the same as the previous values (to a desired tolerance) the solution is said to have 'converged'.

3. Applying the results

Once a solution is obtained, the simulation can be put to use. At this point, a separate module of the program is required that can transform the flow calculation into data readily usable by the engineer; in CFD jargon this is the 'post-processor'. Graphical display of the large volume of data is the most effective way to get a first impression of what it all means. The use of colour adds to the effect.

The usefulness of flow calculations extends from the aerospace and automotive industries to the process and utilities and the environment. Many of the applications of CFD would have been impossible or prohibitively expensive to carry out in any other way (i.e. experimentation) due to health hazards or scale. For example, fire and hazard analysis in rooms and buildings or environmental flows such as the dispersal of pollutants into tidal river estuaries.

4. Benefits

As with all computational techniques, CFD has many advantages over more traditional, experimentally based, design methods. Firstly, CFD is a more fundamental approach. The actual equations that describe a fluid flow process are solved. This leads to a better understanding of the problem and ultimately more confidence in a design. Empirical approaches to design, such as rules of thumb and correlations, tell the designer nothing about the physics of the problem and are usually of limited applicability.

It provides a complete picture of the flow, unlike an experimental programme which only provides measurements at a few pre-selected measuring positions. Experimental programmes are expensive both in terms of costs and time. CFD is both cheaper and quicker and therefore can reduce development costs and the time taken from design to production, installation etc.

Perhaps the greatest virtue of CFD is the potential for optimisation. Once a CFD model is up and running it is easy to make small changes in geometry (bigger, smaller etc.) or changes of flow rates or inlet/outlet positions. For example, suppose that you were studying the ventilation in a chemical production plant. Having set-up the basic model it would be a relatively simple matter to examine what would happen if a chemical was spilt, i.e how the fumes (which could be toxic) would disperse around the plant. Furthermore, the positions of the ventilation system (fans etc.) could be altered to see whether rapid removal of the fumes could be improved. Moreover, alternative designs can be investigated, reducing the reliance on standard designs. Small changes to a physical model are both time consuming and expensive, complete changes in design are virtually out of the question. This ability to carry out 'what if?' calculations and investigate alternative scenarios make CFD a powerful and flexible design tool.

5. Limitations

The technique itself has some limitations in terms of how accurately it can simulate what is happening in the real world. These limitations are, in general, caused by either or both of two things: the grid, and the mathematical models.

The grid has to play two roles in a CFD simulation. Firstly, it has to accurately represent the geometry of interest. Most commercially available CFD programs use what are known as 'body-fitted coordinates' (BFC). This allows the grid to conform exactly to the shape of the volume of interest. However, quite often the external shape (such as the walls of a room) can be body fitted without too much difficulty but, having done this, it can be difficult to represent accurately complex internal structures (eg. the contents of the room). Furthermore, it is up to the user to decide to what level of detail the grid needs to represent. For example, the ventilation in a room is not going to be affected if a cup is placed on a table, and so it would be ridiculous to refine a grid in order to pick up such a detail. Clearly, the finer the detail required, the more cells are needed in the grid, and the more expensive the solution is to compute.

The second function of the grid is to pick up details of the flow. If too few cells are used in a grid, the fine details of the flow, such as recirculation regions (which are the cause of most flow distribution problems) would not be seen by the calculation. In general, the more cells used in a grid, the more accurate the solution will be, but again the simulation will be more expensive. In many cases the fine details of the flow which are important are much smaller than the size of the geometry of interest, and thus if a grid of uniform cell size were to be used, an enormous number of cells would be required to represent the whole volume. 'Grid refinement' is therefore required. That is, placing more cells in regions where rapid changes in the flow are expected and fewer cells where the changes in the flow are small. This requires considerable user skill. A poor grid will give poor answers, and in many cases, no answer at all.

The solutions produced by a CFD code are only as good as the mathematical models they are based on. If the physics is poorly understood, as is the case for combusting flows for example, then the simulations cannot be expected to be any better. Almost all CFD solutions contain an element on numerically induced flow. Checks for the 'grid-independence' of the solution are worthwhile to prevent this.

As long as the limitations of the technique are appreciated by the user, CFD can still be used to great effect. Often, only qualitative results are needed. Here CFD comes into its own for the study of trends (i.e. evaluating the effect of changing a parameter, such as a length, on the measure of performance of interest, such as the heat transfer coefficient). There is always the option of validating predictions using experiments, if added confidence is required. And this does not defeat the purpose of using CFD in the first place. CFD can be used to complement an experimental programme, reducing the number of experiments needed, and making sure that only the important experiments are carried out.

6. Resources

The resources required for a CFD capability are: people, hardware, and software.

6.1 People

Ideally, the CFD user should have a good grounding in both fluid mechanics and numerical methods. A sound knowledge of fluid flow is required because a great deal of intuition is called for in the setting up of a problem, e.g. where to refine a grid in order to pick up all the features of a flow, and in the interpretation of the results, i.e. deciding whether they can be believed. Familiarity with the numerical methods used is not essential in order to get started. However, many industrial applications of CFD will force the user to change some of the model parameters from their default values, or use different models altogether. It is important for the user to understand the implications of making these changes on the accuracy of the solution, the rate of convergence and so on.

The learning curve to become a skilled CFD user is long, six months of continuous use of a particular code would be a fair estimate. It is relatively easy to set up simple problems without much experience, such as two dimensional laminar flow, but for multi-phase turbulent flow with swirl and heat transfer, the pitfalls are everywhere. Unfortunately, most real problems have the latter level of complexity. Newcomers to the field, and in particular their managers, should be prepared to wait a while until some useful results emerge.

6.2 Hardware

At the bottom of the scale are PCs. CFD programs now exist that can be run on a 386 PC (for example, EASYFLOW). However, the size of the memory of a PC limits the size and complexity of the problems that can be handled. By size what is meant is the number of cells that make up the grid, and by complexity the features of the problem such as body fitted grids, turbulence, temperature calculations, multi-phase, etc. For example, figures released by EASYFLOW show that for a 1,000 cell turbulent problem with a body fitted grid the time taken for 100 iterations was about one hour on a standard 386 machine. It has been estimated that the maximum size of problem that can be handled by a PC is about 10,000 cells. 10,000 cells is a small problem in industrial CFD terms. Also, for turbulent problems 100 iterations is not usually enough to obtain a converged solution.

Faster, and more expensive, are the UNIX work-stations such as those offered by Sun, Silicon Graphics, Apollo etc. These machines dominate the CFD market at present due to their relatively low cost and high performance. Typically, work-stations can handle problems up to around 100,000 cells, although the run times for these problems would be of the order of 3-5 days. For example, on B.H.R's 21 MIP work-station (24 MB RAM and 780 MB hard disk) a 30,000 cell turbulent problem with a body fitted grid takes 12 hours. In order to avoid run times longer than a day, the maximum problem size for a work-station of this size would be around 45,000 cells.

6.3 Software

Firstly, and most importantly, what problems will the program be applied to ? This will allow you to decide what features of the available CFD codes are important to you and which are not. The program you eventually choose must be able to solve the majority, if not all, of the problems you set it.

Knowing the kind of problems that you will want to solve can help you decide, for example, whether a package that uses the finite volume (FVM) or finite element method (FEM) is appropriate. FVM is more common at the present (e.g. PHOENICS, FLUENT, HARWELL-FLOW3D, ASTEC, STAR-CD) and is generally faster to compute than FEM (e.g. FIDAP). However, complicated geometries are easier to represent using FEM.

PHOENICS, FLUENT and HARWELL-FLOW3D are capable of handling the widest range of physics (e.g. accurate turbulence models, multi-phase flow, compressible flow, combustion) whereas FIDAP, ASTEC and STAR-CD are better suited to complex geometries.

The range of mathematical models used within the programs are worthy of close scrutiny. In turbulent flows, good codes give a choice of turbulence models, for example. The basic k-ε model is appropriate to simple flows but more accurate models are needed to capture the details of swirl and strong recirculation. Use of the former approach can give results that are simply wrong. The treatment of multi-phase flows (gas-liquid, for example) is possible with only a few programs (eg FLOW3D, PHOENICS, FLUENT). Methods are based on either tracking individual particles or groups of particles, or treating the phases as interacting continua. Both approaches involve considerable simplifications and empirical data input. Other features to consider are combustion, non-Newtonian flow, solution algorithms and so on. Consider all of them carefully and ascertain whether they are useful to you. Take with a pinch of salt the list of 'features that are being developed'. If a CFD code can not solve the problem right now, it is not worth considering.

User-friendliness is an important issue. Beware though, user-friendliness doesn't necessarily mean a code is good. It can be a hindrance as quite often user-friendly codes are inflexible, i.e. to do something slightly unusual with the code can be difficult or tedious. On this point, it is worth checking how easy it is to 'get into' the code in order to change parameters or to specify a problem not covered in the standard commands. For example, you may wish to specify physical properties of the fluid which depend on temperature (eg flows involving molten plastics). These changes are usually accomplished through FORTRAN subroutines. If you feel you will need to use this kind of facility, a demonstration is strongly advised.

Many codes now supply interfaces with CAD programs such as IDEAS or PATRAN -the user then only needs to be expert with one grid generator and can use any CFD package. Some suppliers have also developed intelligent grid generators, which guide the user's actions. Grid generation using only a mouse (no tedious typing) helps productivity further. For example, FLOW3D have recently launched a new grid generation package and post-processor which are both mouse and menu driven. CFDS, the authors of FLOW3D, estimate that these new packages will more than halve the time taken to set-up a simulation.

Lastly, check how easy it is to get at the numbers rather than the results being displayed in graphical form. This is useful if you wish to compare the results of your simulations with experimental data.

7. Still interested ?

Having identified your needs invite the suppliers to give you a demonstration. You will not only get an impression of the program but also of the people behind it (what kind of support will they give?). Once you've made a short-list, specify an example problem, making it as real as possible, and ask them to simulate it. Most of the code suppliers offer consultancy services and will be happy to do this if they think it will lead to a sale. Without a real-live application, you run the risk of plumping for a code, but finding out very quickly that it lacks that one little feature you badly need.

If you don't think the cost of an in-house CFD service is justified, you could use consultants as needed. A number of independent organisations such as engineering consultants and R & D organisations (e.g. BHR Group Ltd.) offer consultancy in all aspects of CFD. These organisations will generally have a number of CFD codes in-house and, being independent, choose the most appropriate one for your problem.

8. Outlook

In the near term, it is likely that CFD will become commonplace in the design office, as well as the research and development laboratory. As the programs become more user-friendly, incorporating artificial intelligence (and experimental data), parallel computing and graphical user interfaces (GUIs), they become more suited to engineers in high-pressure environments with short deadlines. In the longer term, CFD will have its greatest effectiveness when eventually absorbed into umbrella computer systems which handles sourcing, specification, simulation, selection, scheduling and all other relevant aspects of the development, design and construction process, i.e. a 'total systems' environment.