Session 1566

Integration of Simulation into the Undergraduate Fluid Mechanics Curriculum using FLUENT

Rajesh Bhaskaran, Lance Collins
Cornell University
Ithaca, New York

Abstract

The objective of this effort is to integrate simulation technology into the intermediate-level fluid mechanics course in the undergraduate mechanical engineering curriculum at Cornell University. This is achieved using FLUENT, an industry-standard computational fluid dynamics (CFD) package. We seek to expose students to the intelligent use of CFD as well as use FLUENT as a virtual lab environment for hands-on exploration of flow physics and reinforcement of fundamental concepts. Prior to introducing students to FLUENT, we illustrate the underlying numerical concepts such as discretization, grid and iterative convergence, stability, etc. on a simple one-dimensional equation. The classroom examples we have developed are: laminar and turbulent flow in a circular pipe; compressible flow in a nozzle; and flow past an airfoil. In the pipe flow exercise, students are taken through the steps in simulating steady, incompressible, viscous, developing flow in a pipe at low and moderate Reynolds numbers. The concept of turbulence modeling is introduced. Results at the pipe exit are compared with classical results for developed flow (laminar and turbulent) taught in our introductory course in fluid mechanics. The nozzle flow example simulates the high-speed, inviscid airflow through an axisymmetric converging-diverging nozzle. Results for the isentropic case are compared with the classical quasi-one-dimensional results. A non-isentropic case with a shock wave in the diverging section is also presented. In the airfoil example, students simulate inviscid as well as turbulent flow over an airfoil. The lift-curve is compared with thin-airfoil theory. The emphasis of the examples is on the understanding of the solution procedure, and the analysis and justification of results. Our experience demonstrates that FLUENT can be a valuable tool in teaching the proper use of CFD as well as important physical principles at the undergraduate level. The use of hands-on simulations and a rich visual environment facilitates learning of abstract concepts and stimulates student interest.

Introduction

Within the last fifteen years, computer-based simulation has become an integral part of design, analysis and research in fluid dynamics. As in other fields of engineering, the increasingly widespread use of computation has been driven by the dramatic reduction in the cost of computing hardware and the maturing of off-the-shelf, commercial software packages. Despite the prevalence of computational fluid dynamics (CFD) software in industry and research, their use in our undergraduate curriculum has been slight to non-existent.

An informal survey of eight peer universities indicated that their situation was similar. Individual students might use CFD software for projects but it is not part of the curriculum. The use of CFD tends to be restricted to courses that focus on teaching numerical methods. Usually, the basics of the finite-difference method are taught in detail and students develop computer codes to solve a few simple problems. The emphasis is on understanding and programming numerical methods. This is the approach used in the introductory CFD textbooks by Anderson [1] and Tannehill *et al* [2], for example.

There are several reasons why general-purpose CFD software has not penetrated undergraduate fluid dynamics courses to any significant extent. The software used in universities tend to be specialized research codes with steep learning curves. These are unsuitable for use in courses. Commercial CFD packages with user-friendly interfaces that are relatively easy to learn are available but most academics are not familiar with them and unsure about their quality. The cost of licensing commercial software is also a concern. Simulations can have high turn-around times on the hardware found in the typical university lab or classroom.

In light of the important role of simulation in the current practice of fluid dynamics, we feel that the intelligent and proper use of CFD to solve engineering problems is an important skill that needs to be imparted at the undergraduate level. Students need to know how to assess the validity of the numerical results obtained from simulation packages and to understand the limitations of the underlying physical and numerical models. At the same time, hands-on simulations can be used to make theory concrete and can serve as surrogate physical experiments. CFD simulations can be used to make the connection between fundamental science and the solution of real-life engineering problems. Students can be exposed to a larger array of concepts than is possible otherwise.

This project seeks to integrate CFD simulations into the intermediate fluid dynamics course in the mechanical and aerospace engineering department at Cornell University. The software used is FLUENT from Fluent Inc. The software is available at a nominal cost for educational purposes. Our goals are two-fold. First, we'd like to give students a solid foundation in the use of CFD simulations so that they can be informed and critical users. Second, we use FLUENT as a virtual-lab environment in which to explore physical and numerical concepts and to reinforce fundamental concepts in fluid mechanics. This is done through a series of examples or case studies in FLUENT. The focus is not on teaching the use of FLUENT but on teaching general principles through FLUENT. The classroom examples we have developed are: (1)Laminar and turbulent flow in a circular pipe; (2)Compressible flow in a nozzle; (3)Flow past an airfoil. Each of these is discussed in greater detail below.

We consider it important that students have a good understanding of the numerical solution process before being turned loose on a CFD package. Prior to the first FLUENT exercise, we illustrate the CFD process on a simple one-dimensional example:

$$\frac{du}{dx} + u^m = 0; \quad 0 \leq x \leq 1; \quad u(0) = 1$$

We first consider the $m = 1$ case and later the $m = 2$ case, to illustrate how nonlinearity is treated. The discretization is done using the finite-difference method on a four-point grid. The advantage of using a simple example like this is that one can write down the discrete system and solve it easily, making the underlying concepts such as discretization, grid and iterative convergence, stability, etc. concrete. The extension of these ideas to the numerical solution of the Navier-Stokes equations is discussed very briefly. Throughout the CFD exercises, we refer back to the one-dimensional example while discussing numerical concepts. Restricting the discussion of the numerics to a simple example is necessary since there is not enough time available to treat the Navier-Stokes equations. Fortunately, the shorter discussion is sufficient for understanding the CFD solution procedure in our introductory CFD examples and perhaps, even improves clarity.

Laminar Pipe Flow

In the laminar pipe flow example, students are taken through the steps in simulating steady, incompressible, viscous, developing flow in a circular pipe at a Reynolds number Re of 100. Before going through this calculation, we review the analysis of laminar fully-developed pipe flow considered in our introductory fluid dynamics course. We calculate the wall skin friction coefficient and maximum velocity from theory for later comparison with numerical results. The Reynolds number is based on the pipe diameter and the average velocity at the inlet. The ratio of length L to diameter D of the pipe is 40. The inlet velocity is taken to be constant over the cross-section and the exit pressure is set to atmospheric. Students create the geometry and grid using the preprocessor GAMBIT which comes bundled with FLUENT. Alternately, the geometry and grid can be provided by the instructor. The problem is modeled as axisymmetric. So the geometry created in GAMBIT is a rectangle of length L and width $D/2$. This simple geometry provides a good example to introduce students to the process of geometry and grid creation using a general purpose CFD package and to the concept of labeling boundaries in the geometry in order to assign boundary conditions. Constant mesh spacing is used in both axial and radial directions. Three different grids are used with 100 divisions in the axial direction and 5, 10 or 20 divisions in the radial direction.

In FLUENT, residuals are reported for each conservation equation. While it is not necessary for students to understand the details of how these residuals are calculated, they need to have a general idea about it. They need to appreciate that the residual for a conservation equation is the remainder, appropriately scaled, that is obtained when the current solution is substituted back into the discretized form of the equation. We relate these residuals to that in the earlier 1D example in order to aid in understanding their significance. The residuals converge smoothly to a level smaller than 10^{-6} in just 46 iterations on the coarse grid. The convergence is exponential on all grids as one would expect from a good solver. Note that the double precision version of the solver needs to be used to get this convergence behavior. With the single precision version, the continuity residual flattens out and doesn't fall below a level of $5\text{x}10^{-4}$. We mention to students that it is a good idea to monitor a relevant physical quantity or two for convergence in addition to the default residuals provided by a CFD code. Students are asked to monitor the change in the wall skin friction coefficient and centerline velocity in this case.

The variation of the wall skin friction coefficient C_f in the axial direction obtained on the three grids is shown in figure 1. The skin friction coefficient is defined as

$$C_f = \frac{\tau_w}{\frac{1}{2}\rho u_{ref}}$$

where τ_w is the shear stress at the wall and u_{ref} is some reference velocity (taken to be the average velocity at the inlet). The figure shows that we get developing flow near the entrance which becomes fully developed after a certain length. The fully-developed flow analysis gives C_f=0.16 for $Re = 100$, this value of C_f being indicated by the horizontal dotted line in figure 1. As expected, the C_f in the fully-developed region tends towards this value as the grid is refined. The difference in the C_f curves for the 100x10 and 100x20 grids is small indicating that the grid resolution in the radial direction is adequate for the finest grid. Similar trends are observed for the axial velocity at the center of the pipe. It is pointed out to students that a grid resolution study is also necessary in the axial direction.

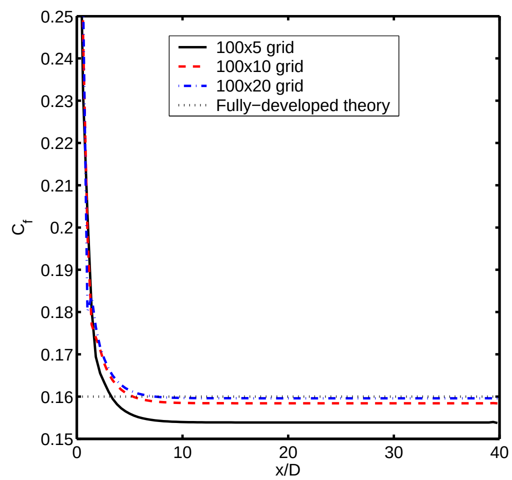

Figure 1: Skin friction coefficient at the wall for laminar pipe flow.

The homework assignment following this example requires students to repeat the simulations at Reynolds numbers of 200 and 500. The variation in the outlet velocity profile and the development length are studied. The development length increases with the Reynolds number with the $Re = 500$ case not being fully-developed even at the exit. Students are asked to justify this trend on physical grounds using boundary layer theory. This gets them to think about the relation between boundary-layer growth and Reynolds number. We ask students to evaluate the order of accuracy of the numerical scheme by plotting the error in C_f in the fully-developed region versus the grid spacing. This exercise leads them to consider the meaning of the order of accuracy of a numerical scheme a little more deeply.

Turbulent Pipe Flow

We next repeat the pipe flow calculation at a higher Reynolds number for which the flow is turbulent. This example is preceded by a review of the structure of a turbulent boundary layer and turbulent fully-developed flow analysis presented in introductory textbooks such as White [4]. Our students have seen these topics in an earlier course and the review serves to remind them of the key ideas which they will use in the analysis of the numerical results. A brief discussion of the need for turbulence modeling and the k-ϵ turbulence model is also undertaken. These concepts are treated in more detail later on in the course. Our hope is that the actual use of turbulence models in this example will provide additional motivation for the later discussion.

We use the k-ϵ turbulence model, with the "enhanced wall treatment" option for near-wall modeling which combines a two-layer model with wall functions. While this option is aimed at making the results less sensitive to the near-wall mesh resolution, it turns out that one still has to be careful about the mesh resolution near the wall. The software vendor recommends that the wall-adjacent cells shouldn't be in the buffer region. To avoid the buffer region, the y^+ value for the wall-adjacent cell, denoted here as y_w^+, should be $\leq$4 or >30, with y_w^+ defined as

$$y_w^+ = \frac{\rho u_\tau \Delta y_w}{\mu} \tag{1}$$

where $u_\tau = \sqrt{\tau_w/\rho_w}$ is the friction velocity, Δy_w is the distance of the center of the wall-adjacent cell from the wall, ρ is the fluid density and μ is the fluid viscosity. Students need to go away with a general idea of the necessity for wall functions in turbulent flow calculations and the restrictions that wall functions place on the near-wall mesh resolution. Since the wall shear stress τ_w in (1) comes from the solution, it is not known *a priori*. In order to determine the mesh spacing Δy_w at the wall such that the wall-adjacent cells are not in the buffer region, we first solve the problem on a coarse grid and plot y_w^+. We then assume that Δy_w scales linearly with y_w^+ and select a Δy_w value such that the corresponding y_w^+ is not in the buffer region.

We consider the case where the Reynolds number is 10,000. At the inlet, we specify a constant velocity as before and a turbulence intensity of 1%. We choose second-order upwind differencing for the momentum, k and ϵ equations. The residuals converge smoothly to a level below 10^{-6} in 350 iterations. It is pointed out to students that the convergence is considerably slower than the corresponding laminar calculation because of the highly nonlinear and stiff nature of the k and ϵ equations. Also, each iteration takes longer than the corresponding laminar case because of the additional equations from the turbulence model.

The coarse grid used is the 100x5 grid from the laminar pipe flow case. The finer grids used are of size 100x30 and 100x60 with Δy_w=1x10^{-3} and grid stretching in the radial direction. The wall-adjacent cell-centers on all three grids lie outside the buffer region. This is verified by plotting y_w^+ in FLUENT. The variation of C_f in the axial direction on the three grids is shown in figure 2. The results on the two finest grids are almost identical indicating sufficient grid resolution on these grids. The classical semi-empirical analysis of fully-developed pipe flow yields C_f=0.0077 for Re=10,000. This value of C_f is shown by the dotted horizontal line in the figure. All three grids yield very nearly the same value of C_f in the fully-developed region. For the finest grid, C_f=0.0086 in the fully-developed region which is off from the classical fully-developed value by 12%. In the developing region, the coarse grid result differs qualitatively from that on the finer grids. This is an instructive lesson to students that they can end up with very wrong results with inadequate grid resolution. The velocity ratio u/u_{av} at the pipe center differs from the classical fully-developed result by 3%. This indicates that the k-ϵ turbulence model gives a very good estimate of maximum velocity but not as good an estimate of the skin friction. This is not surprising since the velocity gradients are mild near the centerline whereas they are steep near the wall. The development length is around 18 pipe diameters. In contrast, the laminar flow is not fully developed at the much lower Reynolds number of 500. This observation provides an opportunity to get students to think about the very disparate growth rates of laminar and turbulent boundary layers.

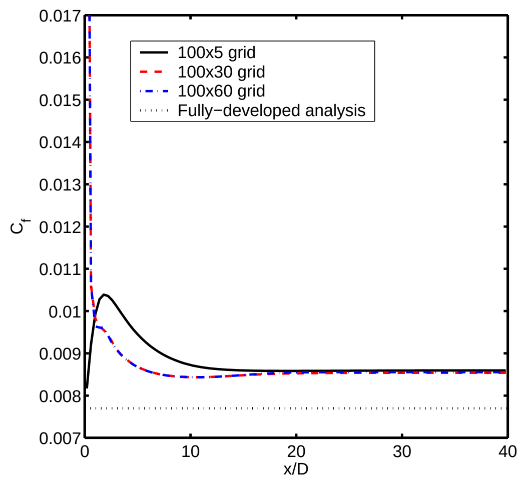

Figure 2: Skin friction coefficient for the turbulent pipe flow case at a Reynolds number of 10,000.

After going through this example, there are a number of useful studies that students can be asked to carry out either in class or as homework assignment. The calculations can be repeated with the k-ω turbulence model. It turns out that the k-ω model doesn't predict fully-developed flow away from the inlet on the finest grid used here. One would have to refine the grid further and check if the flow becomes fully-developed. This exercise drives home the important point that different turbulence models can yield very different results and much care must be taken while using turbulence models to avoid spurious results. A second useful study would be to rerun the calculations at a higher Reynolds number and determine the effect on the outlet velocity profile and development length. It turns out that the agreement with classical fully-developed results gets better on increasing the Reynolds number. This exposes students to the pitfalls involved in using turbulence models intended for moderate to high Reynolds numbers at lower Reynolds numbers. A third study would involve changing the wall roughness at the same Reynolds number and comparing the skin friction values obtained with that obtained from the Moody chart. The effect of wall roughness on the velocity profiles and development length can be examined and the results justified from physical principles.

Nozzle Flow

The nozzle flow example simulates the high-speed, compressible airflow through an axisymmetric converging-diverging nozzle. The Reynolds number for this high-speed flow is large.

So we expect viscous effects to be confined to a small region close to the wall. Hence, it is reasonable to model the flow as inviscid. Switching from an incompressible, viscous model in the pipe flow example to a compressible, inviscid model here involves just a few mouse clicks in FLUENT. This convenience is one of the advantages offered by general-purpose CFD packages. The pedagogical effort can be concentrated on helping students understand underlying principles such as picking the appropriate physical model. This example is preceded by a discussion of the classical, quasi-one-dimensional compressible flow analysis for a nozzle found in introductory-level textbooks such as White [4]. The numerical results for the Mach number, pressure and temperature are compared with those from the simplified one-dimensional analysis.

The geometry and grid are generated in GAMBIT. Setting up the calculation in FLUENT is similar to the pipe flow problem with a few important differences. The fluid is selected to be an ideal gas. When this is done, FLUENT automatically turns on the energy equation. We discuss why this is so whereas it was not necessary to solve the energy equation in the incompressible pipe flow problem. At the inlet, we set stagnation pressure and temperature. We contrast this with the incompressible pipe flow case where the velocity is specified at the inlet. The static pressure p_e is specified at the outlet. Students are asked to predict the flow regime expected using the quasi-1D analysis. The analysis predicts that the flow would be choked at the throat and supersonic in the diverging section from the analysis.

The Mach number variation at the axis and the wall in the axial direction are plotted in figure 3. The corresponding result from the one-dimensional analysis is also shown. The flow is subsonic in the converging section and supersonic in the diverging section as expected for the pressure ratio specified. The Mach number at the axis compares well with the one-dimensional result. The Mach number variation at the wall deviates from the one-dimensional result near the outlet. Students are asked to think about possible reasons for this. A contour plot of the Mach number reveals that the flow is strongly two-dimensional in the wall region near the outlet due to the geometry of the nozzle and this accounts for the deviation from the one-dimensional result in this region. For the homework assignment, students need to study the grid convergence of the results. They are also required to plot the axial variation of the pressure and temperature (static as well as stagnation values) and compare them with the one-dimensional result. This example is used to reinforce understanding of static and stagnation pressure and temperature concepts in a compressible flow. The stagnation temperature obtained is constant indicating that the adiabatic condition is satisfied. There is a slight loss in the stagnation pressure due to numerical dissipation. We briefly discuss the source of this numerical dissipation.

The simulation is repeated at a much higher exit pressure. The quasi-1D analysis predicts a shock in the diverging section of the nozzle. The numerical results show an abrupt jump in the static pressure and an abrupt decrease in the stagnation pressure at a location downstream of the throat. This indicates that there is a normal shock in the diverging section of the nozzle as expected from the quasi-1D analysis. A contour plot of the static pressure reveals the normal shock through bunching of the contours at the shock location. We discuss why the static pressure increases and stagnation pressure decreases across the shock.

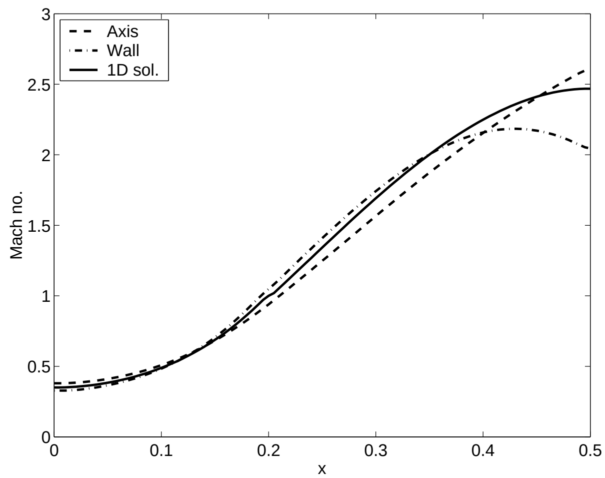

Figure 3: Comparison of the calculated Mach number variation along the axis and wall of the nozzle with the 1D isentropic result.

Airfoil Flow

In the airfoil example, students simulate inviscid as well turbulent, incompressible flow over the NACA 5412 airfoil. The Reynolds number based on the chord length is 3.4 million. This exercise would be a good follow-on to a discussion of thin-airfoil theory. Grid generation for this geometry is more involved than in the pipe and nozzle cases. So the grids are provided to students. However, those interested can go through a step-by-step tutorial for generating the grid that is provided online. A C-grid is used with the coarser grid containing 12,000 quadrilateral cells and the finer grid 24,000 cells. Since this is an external flow in contrast to the preceding examples, we discuss setting the boundary conditions at the farfield boundaries. The velocity is specified at the inflow boundary and the pressure at the outflow boundary. The k-ϵ turbulence model, which was introduced in the pipe flow example, is used in the viscous flow simulations.

Figure 4 shows the variation of the pressure coefficient C_p on the airfoil surface for the inviscid case at 5° and 10° angles of attack. There is a sharp suction peak near the leading edge at the higher angle of attack. Since the area under the C_p curve is approximately equal to the lift coefficient, the figure shows that most of the increase in lift with angle of attack is coming from the suction peak at the leading edge. Such visual reinforcement of concepts is an advantage provided by using hands-on CFD simulations.

Students carry out simulations at three different angles of attack for the inviscid and viscous cases as homework in order to calculate the lift-curve slope. FLUENT provides the capability to calculate the net pressure and friction forces acting on a body in specified directions and also the corresponding force coefficients. This capability is used to calculate the lift and drag coefficients at each angle of attack. The lift-curve slope obtained from the simulations for the inviscid case is 1.92π per radian which is within 3% of the 2π per radian value from thin-airfoil theory. The lift-curve slope for the viscous case is 1.44π per radian which is 25% lower than for the inviscid case. We do expect a reduction in the slope due to viscous effects. However, a reduction of 25% is too large. This could be due to an inadequate resolution of the boundary layer near the wall and deficiencies in the

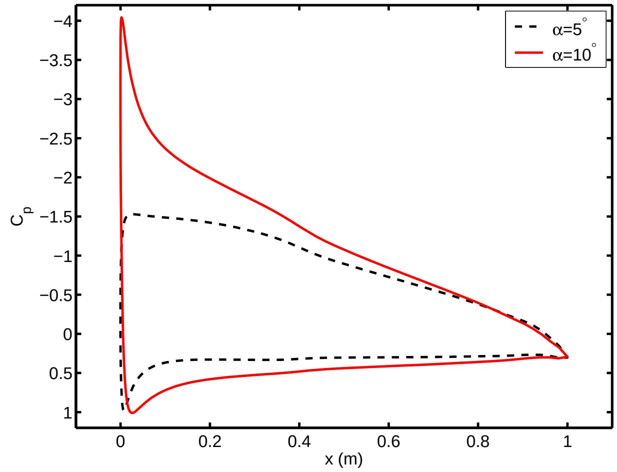

Figure 4: The variation of the pressure coefficient along the airfoil surface.

turbulence model. We have to address both these issues to get a more reliable result. The drag coefficient c_d for the inviscid case is 0.023 when c_l=1.69 at 10° angle of attack. This c_d value is about two orders of magnitude larger than expected for a good numerical solution of the inviscid airfoil problem. Most of this drag is pressure drag. This relatively large drag is due to a significant amount of numerical dissipation. We are investigating ways to reduce the drag obtained for the airfoil with help from the software vendor. We also present a simulation at a high angle of attack when the lift decreases due to flow separation at the airfoil surface. We point out that the separated flow is most likely unsteady in reality and so the steady calculation is not to be trusted. However, the hands-on simulations are very helpful in making students appreciate how boundary layer separation occurs when the angle of attack is increased, causing the airfoil to stall.

For their class projects, some students investigated the effect of airfoil camber on the maximum lift coefficient at low speeds and the drag divergence Mach number at high speeds. For the maximum lift coefficient calculation, they modeled the flow as incompressible and turbulent and increased the angle of attack in steps to determine the angle of attack corresponding to maximum lift. To determine the drag divergence Mach number, they modeled the flow as inviscid and compressible and increased the freestream Mach number until the drag increased sharply. They were able to see that this drag increase is caused by a shock on the airfoil. This exercise gave the students a feel for how to pick different physical and numerical models in different situations.

Conclusion

The above examples indicate that CFD simulations using FLUENT can be of great pedagogical value in teaching important physical and numerical concepts as well as the proper use of CFD. We relate the CFD examples to topics such as fully-developed pipe flow and quasi-one-dimensional nozzle flow presented in introductory fluid dynamics. We also connect the simulations to topics considered in greater detail later in the course such as discretization, stability, turbulence modeling etc. Basic numerical concepts and the CFD process are introduced prior to the CFD exercises by illustrating them on a simple one-dimensional example. We use the one-dimensional example as the foundation for explaining numerical concepts as they come up in the simulations. Other examples being considered are laminar and turbulent boundary layers on a flat plate for teaching fundamental concepts in boundary layers; laminar separation in a diverging channel for teaching the basics of separation and its relationship to the pressure gradient; and flow in a heated pipe for teaching heat transfer basics.

One of the challenges faced is to prevent the cavalier use of CFD and blind trust in results obtained from the software. We emphasize the importance of validating results. Students are required to validate the results presented in their homework assignments and projects. Instances where incorrect results are obtained from simulations are included throughout to alert students to the dangers of the improper use of simulations. Students are encouraged to think about the problem, and make approximate estimates of key parameters before beginning a simulation.

The run times for the simulations used range from less than a minute to an hour on the Pentium III computers in our computer labs. These times are acceptable for use in our course. Longer simulations are assigned as homework. The menu-driven software interface of FLUENT allows calculations to be set up in a matter of minutes. The combination of inexpensive, fast computers and user-friendly CFD software has reduced the turn-around times for simulations to levels suitable for use in the curriculum.

We are developing WWW-based tutorials that provide step-by-step instructions on setting up and running the simulations discussed above. The motivation for this is that it allows classroom time to be spent on hands-on exploration of physical and numerical concepts rather than on teaching the mechanics of using the software. There is evidence suggesting that delivering codified materials through the WWW to liberate classroom time for higher value-added activities such as mentoring and experiential activities can be more effective than a lecture-based presentation of the material [3].

References

[1] J. D. Anderson. *Computational Fluid Dynamics: The Basics with Applications.* McGraw Hill, 1995.

[2] D.A. Anderson J. C. Tannehill and R. H. Pletcher. *Computational Fluid Mechanics and Heat Transfer.* Taylor & Francis, 1997.

[3] D. R. Wallace and S. T. Weiner. How might classroom time be used given WWW-based lectures. *Journal of Engineering Education*, pages 237–248, 1998.

[4] F. M. White. *Fluid Mechanics.* McGraw Hill, 1999.