Int. J. Engng Ed. Vol. 25, No. 6, pp. 1129–1144, 2009
Printed in Great Britain.

0949-149X/91 $3.00+0.00

Computational Fluid Exploration as an Engineering Teaching Tool*

MARC K. SMITH
George W. Woodruff School of Mechanical Engineering, Georgia Institute of Technology, Atlanta, GA 30332-0405, USA. Email: marc.smith@me.gatech.edu

A course in computational fluid dynamics (CFD) at the senior or first-year-graduate levels has traditionally emphasized an understanding of the numerical techniques involved, i.e., finite difference, finite volume, or finite elements, followed by a project in which the student writes his or her own Navier–Stokes solver for a simple flow geometry. The educational pedagogy of this format is that the only way one could truly learn and appreciate CFD was to work through the underlying nuts-and-bolts of these respective methods. The evolution of CFD software over the last twenty years has brought us to the point where a challenge to this traditional pedagogy is in order. In this paper, a CFD course given during the Spring 2007 term at the Georgia Institute of Technology will be described. The course was based on the idea that a tool to successfully solve the Navier–Stokes and continuity equations is available, called COMSOL Multiphysics. The course involved the exploration of a number of fluid flows with the aim of developing a deep understanding of the underlying fluid mechanical mechanisms involved in the flow. Along the way, the student learned about the finite-element method used in the software, how to properly pose the underlying mathematical model for the fluid flow, and about the limitations of the modeling process itself. Specific examples from the course that illustrate these ideas are presented and discussed.

Keywords: undergraduate fluid mechanics; computational fluid dynamics; CFD; numerical flow laboratory; fluid mechanics education

INTRODUCTION

THE HISTORY OF HUMANKIND is full of examples in which technology produces a new way of performing a task that saves time, effort, or money or all of these. The old way is forgotten or just ignored by most people. For example, 150 years ago farmers would plow their fields using a horse-drawn plow. This was back-breaking work for the farmer, as well as the horse, and it took many hours to plow even a small field. When tractors fitted with plowing tools and internal combustion engines were designed for this task, it was easy to see the benefit of this new technology. Today, these tools are a universally accepted part of farming life.

Technological progress has also been made in the area of mathematics. The recording and manipulation of numbers has gone through many stages, from counting fingers and toes to making marks on clay tablets, the abacus, mechanical adding machines, and finally electronic computers. Today, computers are used to automate many routine mathematical computations without much thought by the user. We call this type of activity computer-assisted mathematics (CAM). For example, consider the following MATLAB [1] command to plot the sine function for four periods:

```
>> fplot(@sin, [0 8*pi -2 2]);
```

The result, shown in Fig. 1, allows one to easily visualize the sine function and thus to note some its well-known characteristics.

In another example, consider the problem of finding the length of rope needed to run from the top of a tree a distance H above the ground out to a point on the ground such that the rope makes an angle α with the horizontal as shown in Fig. 2.

Any student of trigonometry knows that the solution is given by the formula

$$L = H/\sin(\alpha)$$

So, if $H = 15$ m and $\alpha = 20^\circ$, most people today would pull out a calculator and compute the rope length to be 43.86 m.

The point of these two examples is to recognize that the computation of the sine function is done by a computer using an algorithm that you trust to be accurate. At first you may construct tests for your computer to make sure you are getting correct answers. However, later on, when you have developed confidence in this tool, you concentrate on the more immediate problem, such as how much rope should you purchase? This is computer-assisted mathematics.

Moving up a level, consider the engineering problem of finding the temperature distribution along a pin fin 10 cm in length whose base is held at a temperature of 1200 K and that radiates to an enclosure with a temperature of 300 K. This conduction/radiation heat transfer problem is

* Accepted 24 August 2009.

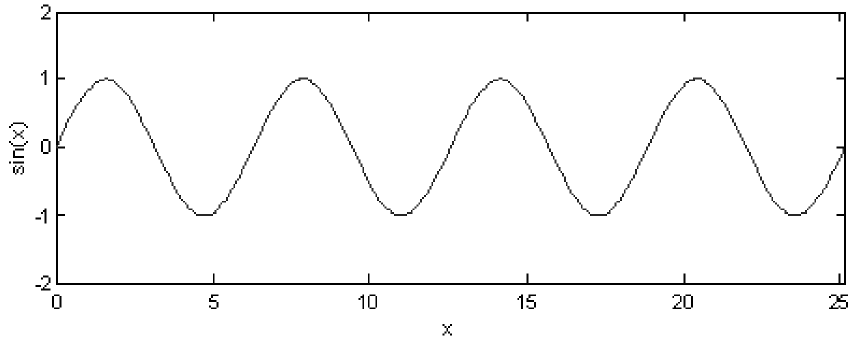

Fig. 1. A MATLAB plot of the sine function.

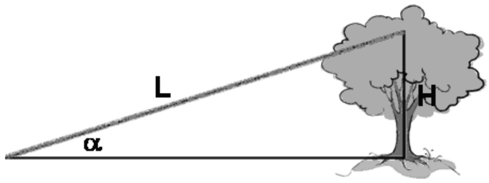

Fig. 2. The rope problem.

governed by the following nonlinear differential equation and boundary conditions

Here, T, L, P, A_c, k, and ϵ are the pin temperature, length, perimeter, cross-sectional area, thermal conductivity, and emisivity, σ is the Stefan–Boltzmann constant, and T_{amb} is the temperature of the enclosure.

$$-(kT_x)_x + \frac{P}{A_c}\epsilon\sigma(T^4 - T_{amb}^4) = 0,$$
$$T(0) = 1200\ \mathrm{K},\ -kT_x(L) = \epsilon\sigma(T^4 - T_{amb}^4),$$
$$T_{amb} = 300\ \mathrm{K}.$$

The solution to this problem is easily given by the following MATLAB command [2],

```
>> sol = bvp4c(@radpinode, @radpinbc,
  solinit);
```

where radpinode and `radpinbc` are MATLAB functions that define the ordinary differential equation and boundary conditions and `solinit` is a structure containing the initial guess for the solution. For a silicon pin 5 mm in diameter with an emissivity of one, the solution is plotted in Fig. 3.

An engineer given this problem would make sure that he or she understands how to use the MATLAB function `bvp4c` [2], how to generate the three arguments to the function, and that the arguments are constructed correctly. After a few checks to make sure this is the case, she would concentrate on the solution and do the studies necessary to inform the radiator design under consideration. The engineer has done CAM.

One concept involved in CAM is known as *abstraction*. This is an important concept in computer science and it is defined as hiding the lower-level details of an algorithm in order to emphasize the algorithm's higher-level structure. The lower-level details are not forgotten, but they are not needed to understand the algorithm at the higher level. This is exactly what is going on in CAM. It is not that the computational details of the ordinary differential equation in the above example are not important. Rather, they are not needed in order to appreciate the behavior of the system modeled by the differential equation. The engineer needs to acquire knowledge about the

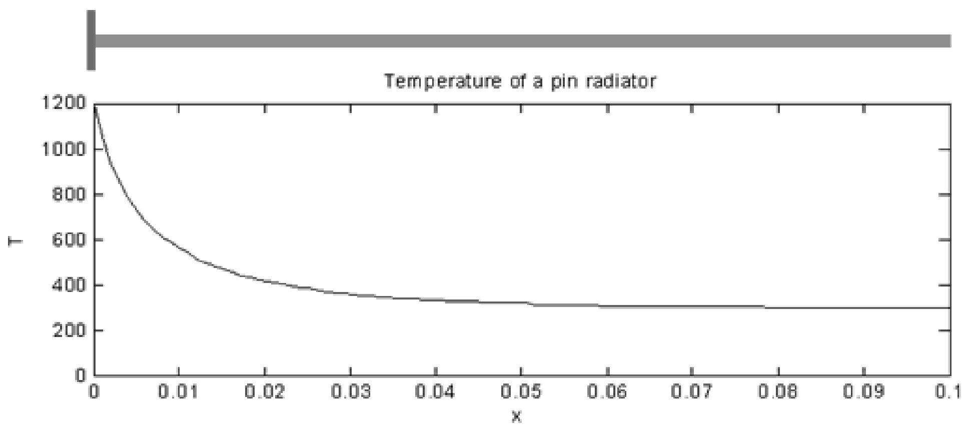

Fig. 3. The temperature distribution of a pin-fin radiator

heat transfer behavior of the system in order to address his design problem. Detailed knowledge about how the differential equation was solved is not necessary for this task.

Abstraction is often confused with the concept of a black box. A black box is a process (possibly unknown) that produces an output based on a given input. It is not a good thing when an engineer is described as ' . . . using a computer like a black box . . .' Typically, this implies that the computer output was accepted without question. This may be a sign of incompetence or at least of laziness or negligence. With abstraction, you concentrate on the details of the work that are most important. For example, in the pin-fin problem it is important to ask whether the base temperature of the fin is truly fixed or whether radiation is the major or only mechanism of heat transfer. With the exception of accuracy, the question of how the first or second derivative of the temperature is computed will not produce any knowledge leading to a better understanding of the physics or to an improved design.

It seems fair to say that for systems of linear and/or nonlinear algebraic equations or for systems of ordinary differential equations computer-assisted mathematics is here. Once the system is defined, software like MATLAB [1] or Mathematica [3] can solve the system to a given accuracy using very sophisticated built-in algorithms that are executed with a single command. For systems of partial differential equations, CAM is close. Systems of two-dimensional PDEs are easily solved and visualized. Time-dependent and three-dimensional PDEs are still a challenge, although this seems to be mostly an issue of execution time rather than algorithm development.

A senior-level elective course in computational fluid dynamics (CFD) was given in the School of Mechanical Engineering at Georgia Tech in Spring 1998. The prevailing educational pedagogy during the development of this course was that the only way to really learn how to do CFD was to write the code yourself. If you didn't, you were just using someone else's code as a black box and your results would be viewed with the suspicion of being inaccurate or completely untrustworthy. Because of this view, part of this CFD course concentrated on the finite volume numerical technique. Lectures were devoted to describing the basic method, ways of meshing the domain, pressure–velocity coupling issues, various kinds of finite differencing, artificial diffusion, etc. Students were given the task of writing parts of a finite-volume Navier–Stokes solver in order to thoroughly learn this numerical method. Afterwards, the students used the commercial finite-volume code FLUENT [4] to solve some simple fluid mechanics problems.

By spring 2007, the field of computational fluid dynamics had matured significantly and computer processing speeds were more than thirty times faster. Using the ideas of CAM and abstraction, the mechanical engineering CFD course was redeveloped to abstract *some* of the numerical issues involved with the computations and to use the time saved to explore and understand fluid flows in more complex geometries and with more complex physics than those seen in elementary fluid dynamics courses. One of the primary reasons for the success of this new version of the course was the use of the software developed by COMSOL Multiphysics [5]. This integrated software made the generation, solution, and visualization of systems of coupled PDEs so easy that the student could easily concentrate on the flow rather than the numerical methods.

The remainder of this paper is organized as follows. In the next section, the Georgia Tech CFD course is described in detail. The core of the course is seven teaching labs performed individually by the students. Examples of these labs are presented and discussed in the next five sections. Students also did individual CFD projects and a few examples of these projects are described in the next section. In the final section, the CFD course is summarized from the perspectives of the student and the professor. Suggestions for future course development are also given.

COURSE DESCRIPTION

The course ME 4342 Computational Fluid Dynamics [6] at Georgia Tech's School of Mechanical Engineering is a senior- or first-year graduate-level elective. An integral part of the course is the use of the Comsol Multiphysics software for the solution of the governing partial differential equations of fluid mechanics. The integrated nature and ease of use of this software allowed the course development to abstract the algorithmic details of the numerical solution and to concentrate instead on the use of the software to explore the behavior and physics of the underlying fluid dynamics. CAM was employed as much as possible, but its limits were explored. It also became possible with this approach to explore the differences between results from model problems and real-world behavior. This was an interesting part of the course for the students because it allowed them to see the differences between various modeling assumptions. Thus, disagreement between numerical and experimental results could be partially explained without resorting to unsupported statements about numerical or experimental error. One final goal of the course was to instill an appreciation that even with numerical abstraction the software could not be treated as a black box. The underlying equations are nonlinear and multiple solutions are possible. Students were able to see this kind of behavior directly from their own solutions.

The course was designed as a three-credit-hour one-semester course. Each week had either two classroom lectures and one hour in the computer laboratory working with the software or one lecture and two lab hours. Students first worked on a COMSOL tutorial and then on seven CFD labs each lasting from one to two weeks. There was also a month-long student project and about two weeks of lectures on the fundamentals of the finite-element method. Students were expected to work an additional ten to fifteen hours a week on the CFD labs or their project.

The seven CFD labs are the core of the course. They are described briefly in the following list. More detailed discussions of five of these labs are located in the five sections below.

Lab 0: COMSOL Multiphysics Minicourse [7]
Length: 1 week
This is the standard COMSOL Multiphysics tutorial used with permission by COMSOL. The geometry is an aluminum film deposited on a silicon substrate. The simulated phenomena are steady state electrical conduction currents with heat generation and heat conduction in both substances. The model investigates the temperature in the device after the current is applied. Students learn to create the geometry in COMSOL, to apply the proper PDEs and boundary conditions for the physics involved, and then to mesh, solve, and visualize the results.

Lab 1: Channel Entrance Length
Length: 2 weeks
Students formulate and solve for the entrance flow of an incompressible fluid in a two-dimensional channel. The accuracy of the solution is examined and the fully developed flow profile is compared with the exact result [8]. A parametric study is done in which the entrance length is determined as a function of the various flow parameters, such as channel thickness, fluid properties, and fluid flow rate. The student then compares his or her results to an accepted experimental/numerical correlation [9].

Lab 2: Boundary Layer Development on a Flat Plate
Length: 1 week
This lab is an examination of the flow development along a sharp-edged flat plate. Students examine the flow far from the leading edge of the plate and compare their results with the Blasius boundary-layer solution [10]. The pressure singularity that forms near the sharp leading edge is also explored. One of its consequences is a small overshoot in the horizontal velocity in the boundary layer near the leading edge of the plate. Students explore whether the velocity overshoot is real or an artifact of the modeling process and whether or not it could be seen in a laboratory experiment.

Lab 3: Flow in Microchannels
Length: 1.5 weeks
Students build and solve a three-dimensional flow model for two distinct sections of a long serpentine channel in a biosensor application: a straight section and a section with a 180° bend. The flow fields are examined at two different flow rates. At the higher flow rate, spiraling secondary flows occur. Students are challenged to explain their observations of this flow using fundamental physical mechanisms from fluid mechanics.

Lab 4: Two-Dimensional Flow Past a Circular Cylinder
Length: 1 week
In this lab, students consider two-dimensional, steady and unsteady flows over a circular cylinder. The development of bound vortices in the wake and the onset of the von Kármán vortex street [11, 12] are examined. In the last part of the lab, the student spins the cylinder about its centerline and sees what happens.

Lab 5: Potential Flow
Length: 2 weeks
This lab is an exploration of various inviscid potential flows. Simple potential flow solutions are visualized and superposition is used to produce more complicated flows, including the potential flow past a spinning cylinder. The velocity and pressure fields are visualized and examined. The last model considered is the flow in a two-dimensional venturi, which is solved numerically (not using superposition). The results are compared with a one-dimensional theory [13] and the student is challenged to discover the source of the discrepancy between the two.

Lab 6: Thermal Convection
Length: 1.5 weeks
This is the first multiphysics fluid flow problem of the course. Students combine the Navier–Stokes equations, continuity, and the conduction–convection heat transfer equation with temperature-dependent material properties to examine the flow in a two-dimensional cavity heated from the side and the flow in a liquid layer heated from below. The physics behind thermal convection is explored in the side-heating geometry. Multiple solutions due to thermal instabilities are explored in the liquid-layer geometry heated from below [14, 15].

Lab 7: Turbulent Flow
Length: 2 weeks
Students examine turbulent pipe flow and the use of a high-speed impinging jet for cooling a heated surface. These problems

use the k–ϵ turbulence equations to solve for the axisymmetric, time-averaged velocity, pressure, and temperature fields. This lab shows students the difficulties involved in solving these equations and also their limitations in predicting real turbulent flows.

A laboratory guide was prepared for each of the seven CFD labs. Each guide introduced its particular fluid flow problem, helped the student build and solve the numerical model, directed the exploration of the flow field, and asked questions about the flow that were answered in the student's lab report. The lab guides followed the same format and style as the COMSOL Multiphysics Minicourse [7] that the students used in the first week of the course. Generally, there were three sections in a lab guide. First, there was a brief introduction that discussed the basic problem and hinted at what kind of flow physics would be learned.

The second section was a discussion of the flow problem and the numerical model that would be used to explore and understand the flow. First, the background and history of the flow problem was discussed. This gave the student a better idea of how the problem fitted in within the field of engineering and why it was important to understand the physics that was involved. Then the geometry for the flow model was presented. Any assumptions used to simplify the geometry were clearly identified. The fluid domains and the appropriate boundary conditions were presented. Anything that would need to be checked, like the extent of the far field domain for external flows, was highlighted. The last part of this section was a presentation of the flow equations that were to be solved. For the first few labs, the equations were presented and solved in dimensional form. After that, non-dimensional equations were used. In these later cases, the guide showed the student how the equations were scaled and normalized and derived the appropriate dimensionless groups.

The last section of a lab guide used a tutorial style that led the student step by step through the process of model building, solution, and exploration. In the first part of this section, the guide instructed the student how to choose the appropriate COMSOL application mode for the physics of the problem, create the geometry for the flow model, set the material properties for flow domains and solid domains (if present), and set the appropriate boundary conditions on the geometry. This is a critical section for two reasons. On a pragmatic level, if the student made a mistake somewhere in this part, subsequent steps or instructions may not work or even make sense. Even if the model-building process was completed, a mistake in the model may cause failure of convergence to a solution or certain features in the model referenced in post-processing may not be available. This section had to be carefully written and checked beforehand by the instructor. It was very helpful to include images of the geometry as it was being created and images of the dialog boxes from the COMSOL interface showing the proper settings that were used. On a more conceptual level, this part of the lab is where the correspondence of the model problem to the real-world system is specified. Simplifications and assumptions about the geometry, the physics, coupling between different physical processes, and appropriate boundary conditions are all presented here. Even if the numerical model is solved to machine accuracy everywhere in the domain with every flow feature fully resolved, the results may be useless if the model built in this section is flawed.

Once the model is built, the guide led the student through the process of meshing the domains, selecting solver parameters, if needed, and then solving the numerical problem. The proper mesh is important for obtaining an accurate solution of the model. The student was usually instructed to build a mesh with enough spatial density to get an initial converged solution. The steps needed to ensure an accurate solution were done in the last part of this section. The solver parameters used in these labs were almost always the default parameters. These worked well for most of the labs. However, in a few situations a different solver needed to be used because of computer memory limits or the automatic scaling of the equations needed to be modified in order to get a converged solution. The reasons for these problems when they arose were carefully explained. This section is where most of the abstraction of the mathematics was done. One can envisage a time when solution algorithms have evolved to where accurate and fully resolved solutions of the model problem are always found, if present. When this time occurs, this entire section could be replaced by a single button marked *Solve*.

The last part of the third section of the lab is the heart of the course. Students are asked to visualize and post-process the numerical solution. The first thing that was done in every lab was to verify the quality of the solution. Mesh studies were done to ensure mesh independence of the result, the domain was extended (for external flows) to ensure that this does not alter the solution, and sometimes different boundary conditions were used to see how these influenced the flow. The students were asked in several labs to verify that simple control volume balances of mass, momentum, and energy were satisfied by their model or to compare their results to published data, especially to experimental results if available.

Once the student has performed all of the checks needed to ensure the accuracy and quality of their numerical solution, the guide directed the student to explore the flow. Sometimes the directions were explicit. For example, in the first lab on channel flow the student was asked to plot the longitudinal velocity profile at several locations along the channel and to compare the results to show the evolution of the flow to fully developed conditions.

At other times the directions were deliberately vague. In the microfluidics lab, the student is only asked to examine the flow in the 180° bend in the microchannel at two different flow rates and to describe any qualitative differences. A bit later, they are asked to describe the physical mechanisms for any flow characteristics that are found, *especially if they are unusual*. Such vague instructions encourage the student to explore the flow using their own insights and skills with the software tools. When they are truly interested in finding some way to visualize a particular flow feature, they are often more motivated to learn new tools for this purpose or to find more creative ways to use the tools they do know. By the end of the course, the students had experience using almost all of the available flow-visualization tools. This includes the generation of time-based or parameter-based movies using the transient or parametric solver.

When a more organized path of exploration is needed, the lab guide described where and what to look for so that the student could see the flow features of interest and examine them to the desired level of detail. For example, in the lab for the flow past a circular cylinder the free stream velocity was set to specific values so that the student could examine a low Reynolds number flow, the appearance of bound vortices in the wake of the cylinder, the growth of the bound vortices, and finally the instability of this symmetric flow to time-dependent vortex shedding.

As the student progressed through the course, the level of detail in the lab guides was reduced. Whenever a new feature or tool in COMSOL was used to perform some task, the guide described how to perform the task with complete step-by-step instructions. The next time the same feature or tool was used to perform a similar task, most of the detailed instructions about the tool or feature were omitted and only the specifics of the task were described. This made the process of writing the guide less tedious and also made following the guide easier for the student. Easier, this is, if they remembered how to use that particular software feature or tool.

It is important that someone who is very familiar with the COMSOL software, either the instructor or a competent teaching assistant, is present in the computer lab during the times when the students work on the lab. Students have many questions related to the use of the software during this time. The errors that occur are tremendously varied, but most can be traced to not following the lab guide well enough. However, uncovering the source of such errors is very difficult, as any computer-support person can tell you. If the student gets stuck and frustrated during the lab, particularly if the error is traced back to an error in the lab guide, they are not happy and they waste too much time with the software rather than exploring the flow.

To conclude the lab, each student was required to write a report. The format was much like a research paper for publication. The report included an introduction of the flow problem and a short discussion of the basic flow equations and boundary conditions. Then the data collected to verify the quality of their solution were presented. The bulk of the report included the answers to the questions posed in the lab guide. At the end, each student was asked to include any suggestions for improvement and to state how much time they spent on the lab. This information was very valuable in structuring subsequent labs. These reports were time consuming for the student to write and the instructor to grade, but it gave the students important practice in communicating the observed fluid physics in meaningful engineering terms.

The class lectures that accompanied each CFD lab were straightforward. At the beginning of the lab, the flow problem was discussed and modeling issues that would become important for that particular lab were identified. When this discussion was completed, students were released to the computer lab to start work. Near the end of the time allocated for each CFD lab there was a lecture and class discussion on the fluid mechanics that was being investigated. If an exact solution for the problem was available, as in a channel flow [8] or the Blasius boundary layer [10], it was presented and discussed. Any experimental or numerical results from the literature that were used for comparison were discussed, such as the channel entry length correlation [9], the von Kármán vortex street [11], or the instability of a liquid layer heated from below [14, 15]. And, of course, the flow physics being examined was discussed. Students asked a lot of questions during this lecture and discussion.

The final component of the course was about two weeks of lectures on the fundamentals of the finite-element method. There were two main reasons for these lectures. First, proper mesh resolution is still an important aspect of finding high quality solutions of the governing partial differential equations. Designing a good mesh means knowing how the mesh interacts with the underlying finite-element approximations of the solution. Secondly, the weak form of the finite-element solution is a concept that is used extensively in the COMSOL software. Some knowledge of the weak solution is needed to make use of several advanced features in COMSOL and to enable the coupling of different physics into a model using weak-form application modes.

The finite-element lectures followed a standard introductory approach to the method based on Reddy [16]. A general second-order ordinary differential equation was used as a model equation. The method of finding an approximate solution of the differential equation with a finite series and the method of weighted residuals was discussed in both the strong and weak forms. The advantages and disadvantages of the weak form were part of this discussion. The extension of the method of weighted residuals to a finite-element mesh was then presented. The weak form of the differential equation on each element was derived and the properties of linear and quadratic approximation functions on the element were presented. From here, the assembly of the element equations to form the set of linear equations for the global problem and the basic solution technique for the global problem were described. The depth of these lectures gave the student a basic understanding of the finite-element method and an understanding of what the underlying finite-element mesh was supposed to accomplish. In addition, it familiarized the student with elementary terminology used in discussing the finite-element method and which appears frequently in the COMSOL documentation.

Since fluid exploration was the main purpose of this course, each CFD lab had specific flow structures that were studied. The next six sections of this paper highlight and discuss some of the more interesting flow structures from five of the labs. A few examples of the month-long student projects are also presented.

TWO-DIMENSIONAL CHANNEL FLOW

The first fluid dynamics lab was steady flow in a two-dimensional channel. The flow geometry was a rectangular domain with no-slip boundary conditions on the top and bottom, a uniform inlet velocity profile on the left side, and a zero pressure outlet boundary condition on the right side. A typical surface plot of the velocity magnitude is shown in Fig. 4.

This was an excellent first flow lab for the students because they had all studied this flow to some degree in their introductory fluid mechanics class, which was a prerequisite for this course. Students examined the fully developed velocity profile and compared their numerical results with the exact parabolic solution [8]. They plotted the developing boundary layer flow along the solid walls and saw how the boundary layers merge and the flow evolve to become fully developed.

The students were also asked to determine the entrance length L_E of the flow for different flow rates and to compare their results with an accepted correlation [9]

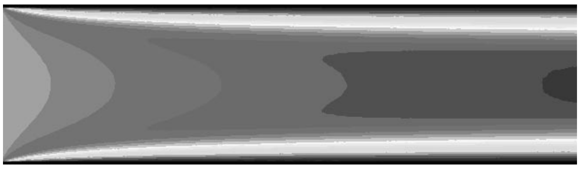

Fig. 4. Grayscale surface plot of the velocity magnitude in a two-dimensional channel flow. Velocity magnitude is zero for white to a maximum for dark gray.

$$\frac{L_E}{h} = \left[(0.631)^{1.6} + (0.0442\ Re)^{1.6}\right]^{1/1.6},$$
$$Re = \frac{\rho U_0 h}{\mu}, \tag{1}$$

where Re is the Reynolds number based on the channel height h and the average velocity U_0. This exercise was challenging to the student because they had to decide for themselves when the flow became fully developed. Should this point be based on when the centerline velocity became constant, when the pressure gradient became constant, or when the inviscid core flow [17] disappeared? Once these questions were explored, the most careful students got results that agreed quite well with the entrance-length correlation in Equation (1). The length of the inviscid core flow was determined by plotting the constant from Bernoulli's equation along the center streamline from the entrance to the end of the channel, as shown in Fig. 5. The result should be constant in the inviscid region and then decrease when viscous effects become important. In Fig. 5, the plot is nearly constant until about $x = 0.8$ m and then it decreases significantly. For this particular value of the inlet velocity, the entrance length was near $x = 2$ m. This result was surprising to the students because many undergraduate textbooks show the merging of the boundary layers in the channel and the end of the inviscid core as being coincident with the entrance length of the channel, e.g., Munson, Young, Okiishi, and Huebsch [17] and Çengel and Cimbala [18].

Another interesting feature of this lab was the presence of a pressure singularity on the top and bottom rigid walls at the channel entrance. The students explored this flow feature and noted that the singularity could not be resolved by any mesh refinement. They were led to realize that the effect was caused by the way the constant velocity profile

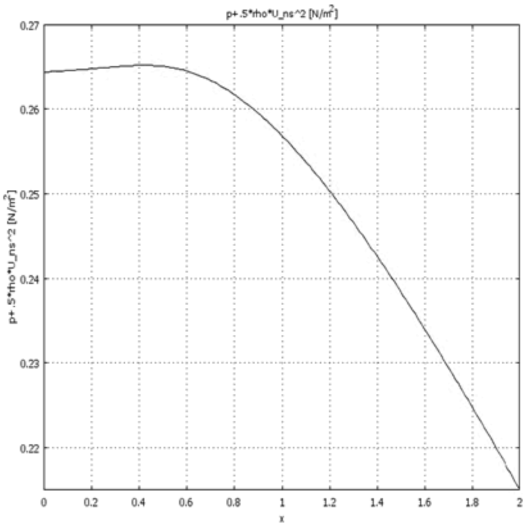

Fig. 5. A plot of Bernoulli's equation along the center streamline of the channel.

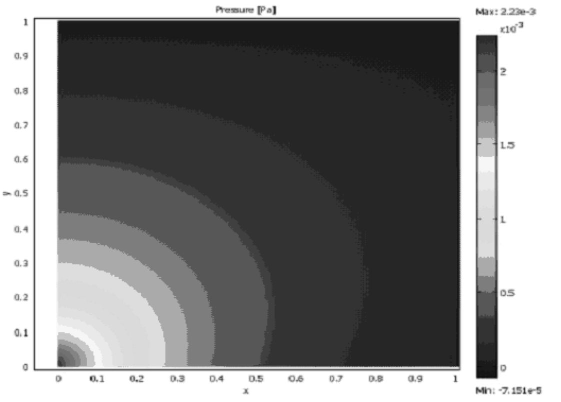

Fig. 6. The pressure singularity near the leading edge of the flow over a flat plate.

was implemented on the inlet surface where it met the no-slip surfaces at the top and bottom walls. This is an obvious numerical artifact and so it was not explored further in this lab.

The pressure singularity was seen again in the next lab on boundary layer flow over a sharp-edged flat plate. A surface plot of the singularity at the leading edge of the flat plate is shown in Fig. 6. In this lab, the students explored how the singularity could be removed by better modeling of the inlet velocity or resolved by a more accurate model of the leading edge of the flat plate.

FLOW IN A MICROCHANNEL

Microfluidic devices have the potential to improve or accelerate all sorts of different chemical and biological processes used in technological devices and systems. The students of today are well aware of this and are very excited to study the relevant flows. Such flows are ideally suited to CFD because they are almost always laminar flows due to the small dimensions of the flow geometry. The microfluidics flow lab gave the students an opportunity to simulate such a flow and to see some unexpected phenomena.

The flow lab considered two different three-dimensional geometries: a length of straight channel and a channel with a 180° bend. These two geometries were the only three-dimensional geometries considered in the course. Students saw why this was the case immediately because they were expected to create the geometries themselves. Thus, they learned many of the software tools needed to create, manipulate, and view the 3-D geometry. They experienced the process of meshing these geometry using both structured and unstructured meshes. Solving these systems showed them that even a simple flow can take a large amount of computational resources (memory and time). And finally, they learned that even if you have a good solution, examining the flow, determining how to use the tools needed to view specific flow structures, and then understanding the flows is not an easy task.

The 180° bend geometry had a channel cross-section of 300 μm by 150 μm with two straight sections that were 1000 μm long. The fluid was water and flows with two different maximum velocities were simulated: 0.007 m/s and 0.7 m/s corresponding to a Reynolds number of 0.94 and 94 respectively. The students were challenged to describe the differences between the two flows. Typical results are shown in Fig. 7 where the inlet is on the left leg of the geometry.

The major problem with this lab was determining how to examine the flows. The lab guide showed how to make slice and boundary plots, but other visualization tools were needed to fully appreciate the flow structures, especially in the higher Reynolds number case. For a small Reynolds number (Fig. 7(a)), the flow is dominated by viscous stresses and the velocity profile as the flow travels along the bend is symmetric about a center plane that passes between the two legs of the geometry. This is seen in Fig. 7(a) because the velocity maximum occurs near the inside of the bend just before and after the bend.

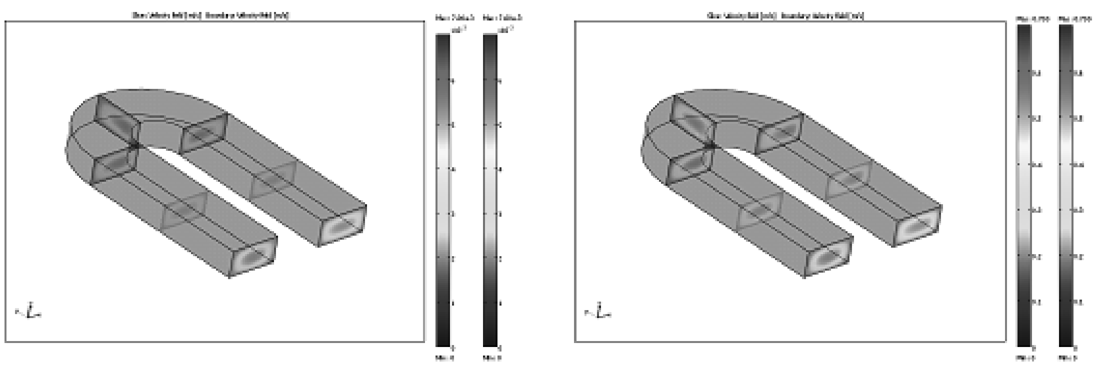

Fig. 7. Cross-section plots of the velocity magnitude in a microchannel with a maximum velocity at the inlet of (a) 0.007 m/s, and (b) 0.7 m/s.

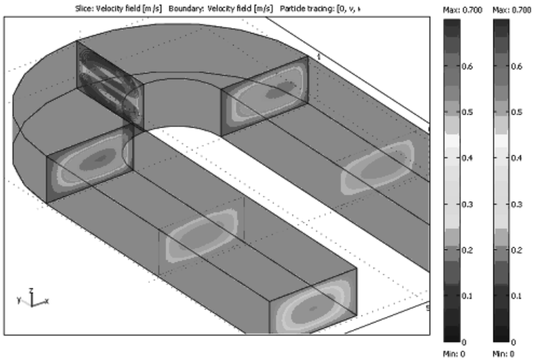

Fig. 8. Projected path lines on a central plane from the particle tracing tool show the secondary flow in the bend for a Reynolds number of 94.

This symmetry is in contrast to the larger Reynolds number case shown in Figs 7(b) and 8. Here, the velocity maximum is located near the inside of the bend on the left leg and near the outside of the bend on the right leg. The students learned that the reason for this is inertial effects that cause the pressure to increase towards the outside of the bend, which in turn produces a secondary vortex flow in the bend. They were asked to explore the various visualization tools in the software and find a way to visualize this flow. The best result used the particle-tracing tool based on the fluid velocity projected on a plane at the middle of the bend as shown in Fig. 8. The resulting closed flow lines with the velocity in the middle of the cross-section directed to the outside of the bend show the secondary flow in the bend very clearly.

In the lecture part of this lab, the classic tea-cup experiment [19] was performed in which a few loose tea leaves were placed in a cup of water, which was then stirred. The tea leaves were heavy enough to sink to the bottom, but large enough to be moved by the flow. The result was a demonstration of the same type of secondary flow seen in the microfluidics lab, i.e., the tea leaves resting on the bottom of the cup migrated toward the center. The mechanics of the flow was completely described in qualitative terms and discussed with the students. They were then asked to use this same mechanism to explain the secondary flow seen in the flow lab and to write the explanation in their own words as part of their lab report. Their report provided a powerful re-enforcement and a demonstration of the depth of their understanding of this flow phenomenon.

FLOW PAST A CYLINDER

Not many flows are more accessible to the student than the flow past a cylinder. It is only examined in a simple way in an introductory fluid mechanics course because of time constraints and because the flow solution has to be done numerically. This flow can be both dramatic and expensive, as shown by the satellite image of a similar flow past Jan Mayen Island in the North Atlantic Ocean [20] (Fig. 9) and the collapse of the Tacoma Narrows Bridge [21] by a twisting structural self-excitation with a time-dependent vortex shedding flow. With these motivations, the students were very excited to examine the flow for themselves. The geometry creation was very easy and students studied the flow for increasing velocities measured by a Reynolds number based on the diameter of the cylinder from 1.37 to 137.

Students learned one thing very quickly in this lab; adequate mesh size and spatial resolution is very important to accurately resolve the flow structures seen for small and large Reynolds numbers. When the upstream flow speed is small so that the Reynolds number is 1.37, the flow is dominated by viscous effects and the flow field around the cylinder extends many cylinder diameters in all directions. Figure 10(a) shows a surface plot of the velocity magnitude for such a flow. The small white circle is the cylinder and the large extent of the flow domain is easily seen. When the students explored this flow, they were surprised to find that there was no recirculation region or wake behind the cylinder. This kind of a wake is an inertial effect that appears at a specific Reynolds number. Figure 10(b) shows an example of this kind of a wake. The two recirculating regions are called bound vortices and they were visualized using the particle-tracing tool. During the lab, students ran the simulation for a range of different upstream flow speeds to find the critical

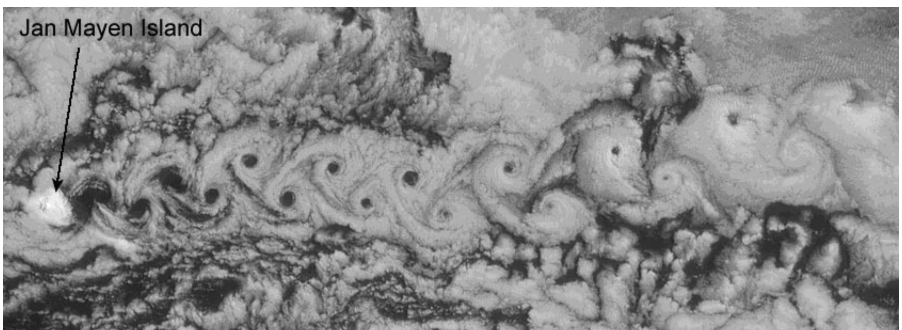

Fig. 9. Satellite image of a von Kármán vortex street behind Jan Mayen Island in the North Atlantic Ocean on June 6 2001 [20].

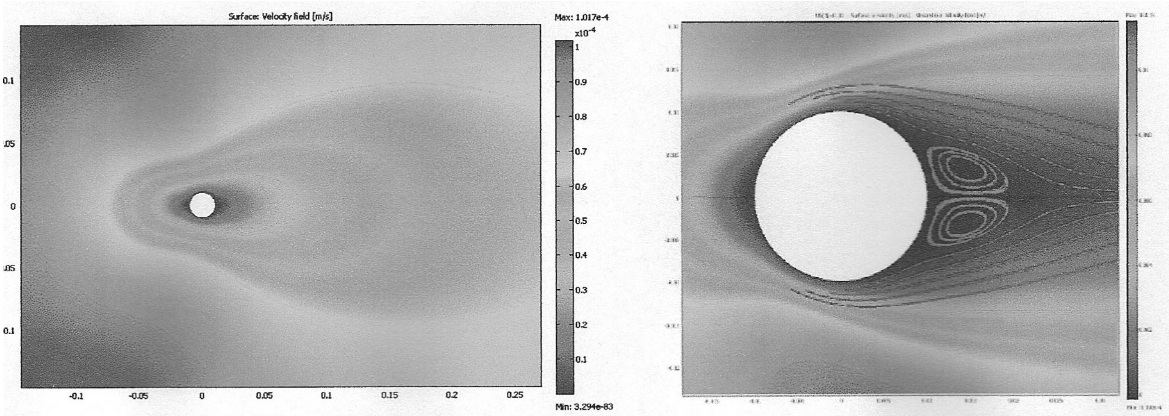

(a) (b)

Fig. 10. Steady two-dimensional flow past a cylinder. (a) $Re = 1.37$ and (b) $Re = 13.7$.

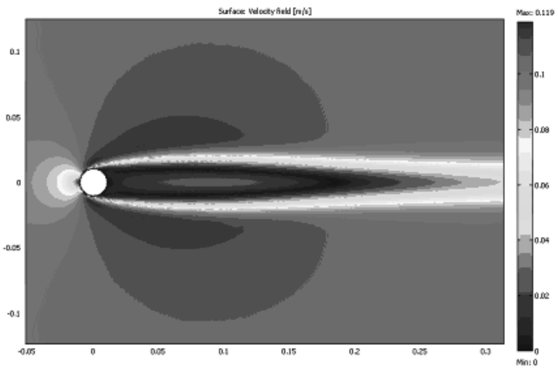

Fig. 11. Steady two-dimensional flow past a cylinder for Reynolds number $Re = 137$.

Reynolds number at which the bound vortices first appear in the flow. Then, as they increased the Reynolds number, they determined the length of the vortices and compared their results with analytical results from the literature [22].

For the largest upstream flow speed, the Reynolds number is 137. The simulation of the steady flow at this Reynolds number was challenging for the students because it usually did not converge on their first try due to inadequate mesh resolution. After careful examination of the flows at smaller Reynolds numbers and repeated attempts at refining the mesh most students were successful at computing the flow as shown by the surface plot of velocity magnitude in Fig. 11. This flow is exactly what the student expected: symmetry about the midplane, a high-pressure region at the front of the cylinder, and an extended wake behind. The next portion of the lab taught a lesson that every student of fluid mechanics needs to know. The lab guide instructed the student to switch to the time-dependent solver and to solve for the flow using an initial guess that was a uniform horizontal velocity. The resulting data were displayed as a movie and showed how the flow field settled down to the steady flow computed previously. The initial guess for the flow was then changed to a uniform velocity

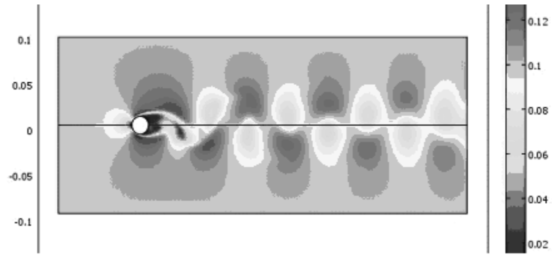

Fig. 12. The time dependent von Kármán vortex street behind a cylinder.

directed 20° to the horizontal. The resulting movie of the time-dependent flow showed the development of vortices that were shed off alternating sides of the cylinder to produce the well known von Kármán vortex street [11, 12] shown in Fig. 12.

Students certainly expected to see such a flow by now in the lab because of the discussion at the beginning of the lab guide. However, they were usually surprised at how easy it was to simulate this flow and a bit puzzled to see that they have now computed two reasonable solutions to the flow field for the same Reynolds number, one steady and one unsteady. The lecture part of the lab that took place after the students had done this simulation was a discussion of flow instabilities and multiple (non-unique solutions) to the underlying partial differential equations.

POTENTIAL FLOW

Potential flows and superposition are not covered in the semester-long introductory fluid mechanics course at Georgia Tech. Therefore, the pre-lab lecture for this lab covered the basics of potential flow theory. Students saw that a velocity potential satisfying Laplace's equation governs the flow field under the assumptions of an inviscid, irrotational, incompressible flow. They also saw that Bernoulli's equation produced the corresponding pressure field for the flow. After this

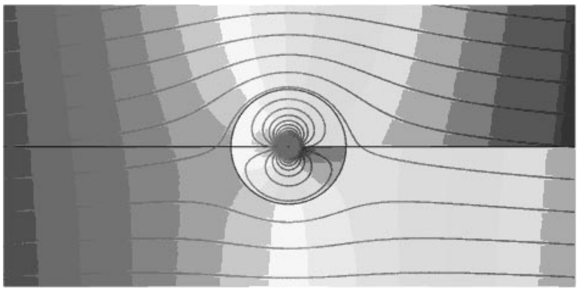

Fig. 13. The predefined potential flow model for a uniform flow past a spinning cylinder. The figure shows the streamlines for the velocity field and a surface plot of the potential field.

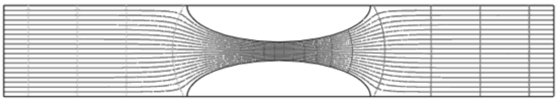

Fig. 14. Two-dimensional potential flow in a venturi. The mostly horizontal lines are streamlines and the mostly vertical lines are constant-potential lines. The flow is from left to right.

derivation, the ideas of a linear flow field and superposition were easily presented. Simple fundamental potential flows such as uniform flow, source/sink flow, doublet flow, and free-vortex flow were presented as general solutions to Laplace's equation that could be verified easily by substitution.

The potential flow lab included the only predefined model used in the course. The reason for this is that the source/sink, doublet, and free-vortex flows had singularities in the potential field that needed to be carefully treated in order to produce accurate solutions. For the first part of the lab, students copied a COMSOL model file based on the potential flow equations and were directed to turn on and off various fundamental potential flows by setting various constants to either a nonzero or zero value. When the resulting streamlines for the flow were plotted, the students saw all of the fundamental flows mentioned above and, using superposition, the flow past a Rankine half body, flow past a cylinder, and flow past a spinning cylinder. This last flow is shown in Fig. 13.

After this predefined model, the students were asked to numerically compute the potential flow for simple geometries such as the flow past a circle and a square by solving Laplace's equation with no superposition. The most interesting portion of this section of the lab was the investigation of two-dimensional flow in a venturi. The streamlines and constant-potential lines for this flow field with the flow going from left to right are shown in Fig. 14.

Students were also asked to derive an equation for the flow rate in the venturi using only Bernoulli's equation, conservation of mass, and the assumption of a uniform flow at the throat of the venturi and upstream of the throat. The result for the flow rate Q in terms of the pressure difference between these two locations is the following:

$$Q = \left(\frac{2\Delta p/\rho}{h_u^2 - h_t^2} \right)^{1/2} h_u h_t, \tag{2}$$

where Δp is the pressure difference between a point well upstream of the throat and at the throat of the venturi, h_u and h_t are the heights of the channel at these two locations, and ρ is the fluid density. This result is similar to the well-known result [13] for the flow rate in a pipe used by pipe flow meters. When the student extracted the required pressure readings from their numerical potential flow solution and compared the numerical flow rate Q_n with the predicted flow rate she found that the relative error $(Q–Q_n)/Q_n$ was 1.4%. Furthermore, she found that this error did not go to zero with increasing mesh resolution, but actually got more precise. The challenge for the student was to correctly explain this discrepancy as a result of the non-uniform velocity at the throat of the venturi. Surprisingly, this proved to be difficult for many students. Explanations in terms of numerical errors, measurement error, and the effects of friction on the flow were very common and incorrect. Most likely, this difficulty was because this was the first time some students had to think in terms of the difference between two model flows rather than the difference between a numerical flow solution and a real flow. Developing this kind of reasoning skill is very important to engineers who use CFD results to assist in the design process.

THERMAL CONVECTION

The thermal convection lab was the first one to do multiphysics, which in this case is a combination of fluid flow with heat transfer. Two geometries were considered, heating from the side and heating from below.

Heating from the side

The side heating model was very simple: a fluid contained in a rigid square box, a fixed temperature on the left wall that is higher than the fixed temperature on the right wall, and no heat flux through the top and bottom walls. The lab guide led the student though the entire modeling process for this geometry. Step by step the model was posed, solved, and visualized to explain the physics. At first, constant fluid properties were assumed. The solution to this flow is simple one-dimensional conduction from left to right and a stagnant fluid. The student was asked to think about how the fluid could move due to this heating? The idea of temperature-dependent fluid properties was presented and how the change in the density would lead to a body force that could drive fluid motion. The Boussinesq approximation [23] was then presented, in which all fluid properties are assumed to be constant except for the fluid density, which varies only in the body-force term

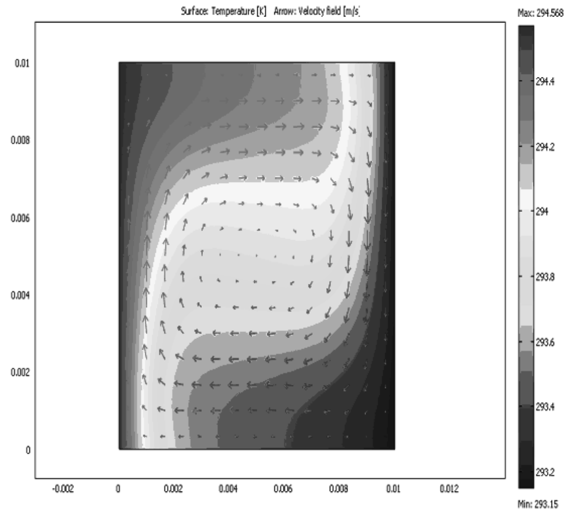

Fig. 15. Thermal convection in a side-heated enclosure. The grayscale surface plot is the temperature field. The left side is hot and the right side is cold. The arrows are the velocity field.

of the equations. The resulting solution is shown in Fig. 15.

In this lab, the governing equations for the model were posed in dimensionless form. The dimensionless groups that appear are the Rayleigh number, which is the magnitude of the driving body force, and the Prandtl number, which is the ratio of the momentum to the thermal diffusivity. Students were asked to explore the model as these two parameters were varied and to explain what they saw in terms of physical mechanisms. They saw the development of thermal and velocity boundary layers as the Rayleigh number increased. As the Prandtl number decreased, thermal diffusivity becomes more important and the overall velocity and pressure magnitudes decreased because the body force was less able to drive the flow. The resulting flow field became more circular and a pressure minimum appeared at the center of the cavity. This is due to an inertial effect in the flow related to the same physical mechanism that produced the secondary flow in a bend seen in the earlier microfluidics lab. Students were asked to explain this effect and were reminded of this earlier result.

One quantity of particular interest in this problem is the net heat flux through the fluid, which in dimensionless terms is the Nusselt number. Students could easily obtain this quantity from their simulation using the integration tools included in the COMSOL software. Using the parametric solver they computed the Nusselt number as a function of the Rayleigh number from 2000 to 200 000 and compared their results with a standard heat transfer correlation found in undergraduate heat transfer texts [24]. The numerical results were consistently about 10% below the correlation.

A final exercise in this part of the lab that the students found very powerful was the inclusion of the full temperature dependence of all fluid properties in the model. COMSOL includes a standard materials library with this data in terms of callable functions. It was a simple matter to place these functions in the appropriate input boxes of the software and then resolve the model to see the net effect on the flow. The changes were minor, but definitely observable and in some instances more significant than expected. Surprisingly, the Nusselt number results for the simulation, which were consistently less than the standard correlation using the Boussinesq approximation, crossed over the standard correlation for a Rayleigh number of approximately 120 000. The two results now appeared to be in better agreement.

Heating from below

The geometry of an infinite liquid layer confined between rigid walls and heated from below is a classic instability problem in fluid mechanics [14]. For a Rayleigh number Ra less than 1708, the velocity field is stagnant and the temperature field is given by one-dimensional conduction through the fluid, i.e., it is linear with distance above the lower wall. For $Ra > 1708$, the layer is unstable because warmer-lighter fluid is beneath cooler-heavier fluid. The result is convection rolls or cells in many different forms depending on the system [25].

In this part of the lab, the students explored the buoyancy-driven instability by solving their model using the transient solver for Rayleigh numbers of 1000 and 2000. The initial conditions for the velocity and temperature fields were an approximation of the unstable convection solution that the students input directly into the model in terms of simple polynomial and sinusoidal functions. The transient results showed either the decay of the fluid motion back to a stagnant conducting layer for $Ra = 1000$, or the development of the steady convection rolls shown in the bottom image of Fig. 16 for $Ra = 2000$.

The students were then asked to compute the steady flow solution for $Ra = 1709$, a point just above the onset of the instability. The result is the top image of Fig. 16. The rest of the time in the lab was spent comparing the two flows for $Ra = 1709$ and 2000 in terms of flow and heat transfer characteristics using whatever software tools the student felt were needed.

For their lab reports, students were instructed to include clear explanations for their observations in terms of relevant physical mechanisms. This type of exercise helps them develop an intuitive understanding of the physics involved in the flow, which can then be extended to other problems. For example, the major qualitative difference in the two flows shown in Fig. 16 is the loss of vertical symmetry in the temperature field for the higher Rayleigh number case. This was correctly explained by most students as an effect of fluid convection. Most students presented their results in terms of separate surface plots for the flow and

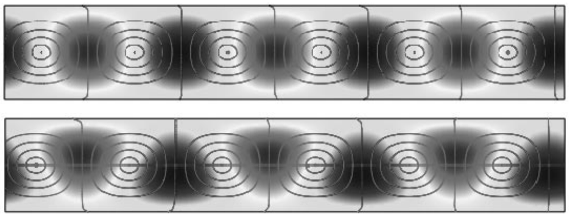

Fig. 16. Simulations of a fluid layer between rigid plates heated from below. The surface plot is the deviation of the temperature field from the conduction solution. Fluid cells with the darkest gray are hot, cells with a mid-gray are cold, and the very light gray between cells has a medium temperature. Solid lines are streamlines. Top, $Ra = 1709$; bottom, $Ra = 2000$.

temperature fields, rather than the combined steam line and surface plots shown in Fig. 16. However, the combined plot shown in the top image in Fig. 16 leads to a very simple question for discussion with the students, i.e., what is the direction of the velocity disturbance in each convection cell, clockwise or counter-clockwise, and why?

STUDENT PROJECTS

The student computing projects were an opportunity for each student to explore a fluid flow problem of their own choosing. They were given about two weeks to choose the project and about a month to build the model, perform the simulation, analyze the data, and write their report. Each project was screened before the model was built to ensure that it was doable in a month and that it had enough physics to make the fluid exploration interesting. Students were strongly encouraged to choose a project that could be done using a two-dimensional simulation. The reasons for this are that a three-dimensional model is typically much harder to build, mesh, and solve than a two-dimensional model. Also, students did not have much experience with the three-dimensional software tools in COMSOL since they had used them only once in the earlier microfluidics lab. And, most importantly, previous experience with student CFD projects showed that most students could not do a three-dimensional simulation in one month.

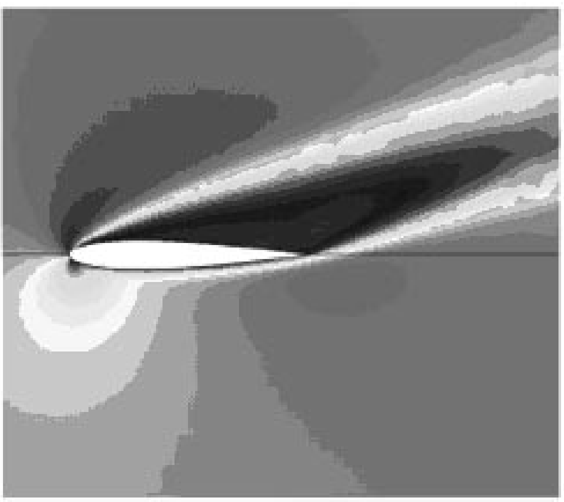

Fig. 17. Onset of stall for flow over a NACA 0012 airfoil.

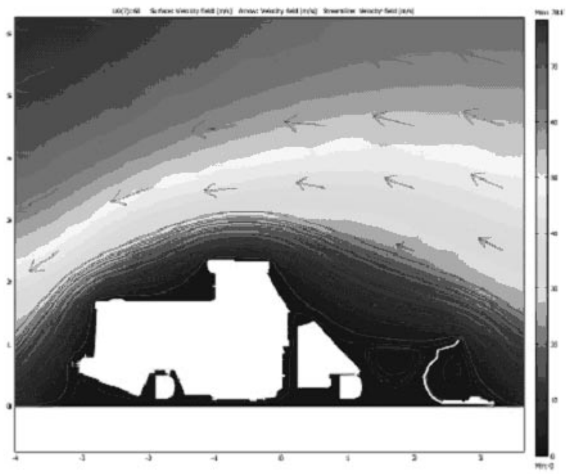

Fig. 18. Aerodynamic drag on a bulldozer.

As expected, there was a wide variety of student projects. The simplest were external flows. Fig. 17 shows a result for one project on the flow over an airfoil. The student computed the lift and drag on an NACA 0012 airfoil section [26] versus the angle of attack of the oncoming flow. The numerical data were compared with wind tunnel data and the critical angle for the onset of stall was identified. The student discussed how lift was provided by the flow and the physical mechanism for the stall phenomenon.

Another student compared the aerodynamic drag on a number of different vehicles, from race cars, to trucks, to the bulldozer shown in Fig. 18. The fluid mechanics was just a simple comparison of the drag force on the vehicles, but the student spent 30 hours on the project. Most of this time was spent trying to import the vehicle shape data from 3D CAD models. COMSOL has a CAD import module that makes this process much easier, but it was not available for this course.

A good multiphysics project was the flow and heat transfer in a cross flow past a cylindrical tube bank. The student exploited the use of symmetry in the tube bank to produce the model shown in Fig. 19. He examined many aspects of the flow, such as flow separation behind the tubes, pressure, temperature, and velocity distributions, particle paths in the flow, and total heat transfer as a function of the inlet velocity.

A final example of a student project was time dependent heat transfer and convection in a solar

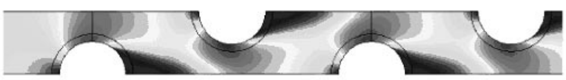

Fig. 19. Surface plot of the temperature field for cross flow in a heated cylindrical tube bank.

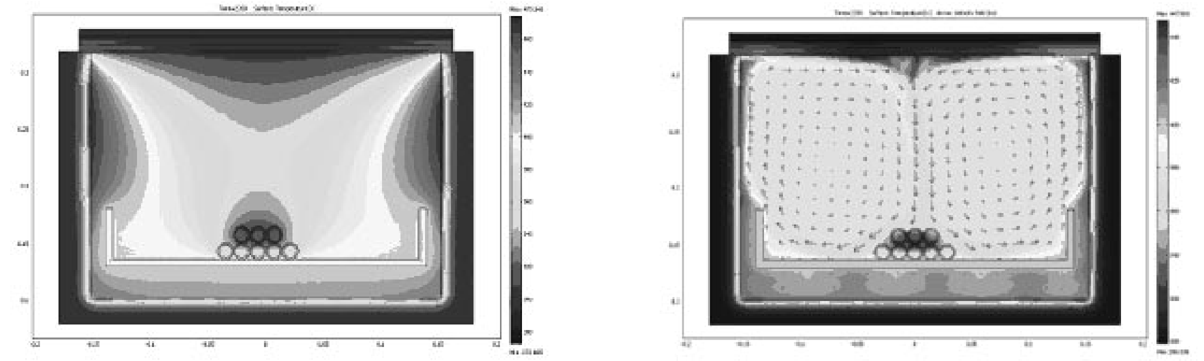

Fig. 20. A surface plot of the temperature in a solar oven after 2200 seconds of operation without convection (left) and with convection (right). The arrows are the velocity field at this instant in time. The flow is not steady.

oven as shown in Fig. 20. This project was more complicated than any of the other projects. However, it was part of this student's senior-design group project and since he was very interested and motivated by the project, he was allowed to do it. He encountered many problems during the work and solved them all, but he spent over 50 hours outside of class to complete the work.

The project was to design a simple solar oven for the sterilization of hypodermic syringes in remote third world environments. The basic physics was the same as the thermal convection lab discussed above, so the student had a good start. The geometry came from his group's design. Since this was a solar oven the student had to include thermal radiation in the governing equations for heat transfer. Radiation is included in the heat transfer module that was used in the course, but it was not covered in any of the labs. The student did his own study of the COMSOL help files and was able to successfully implement radiation into his model. However, one modification that was needed was that the upper glass surface of the oven operated like a band-pass filter, allowing thermal radiation to pass through the glass surface at visible frequencies but blocking infrared radiation from leaving the oven. This filtering effect was not part of the standard settings used in the COMSOL heat transfer module. Thus, the student read the documentation on the implementation of the radiation boundary condition for the glass surface and modified it to mimic this effect. The resulting simulation (shown in Fig. 20) proved to be very helpful in the successful completion of his senior-design project.

CONCLUSIONS

The course was a success for the students and the instructor. The major objective of the course, which was to develop a greater physical understanding of complex fluid mechanical phenomena, was achieved. The COMSOL software with its integrated toolset and its graphical user interface was easily learned by everyone. This allowed the students to spend most of their time investigating flow fields and trying to understand flow physics rather than wrestling with the software to build a model or get a solution. The following subsections summarize comments and observations by several students and the instructor about the course. Student comments were expressed to the instructor either personally or in the online course evaluation that is given to all Georgia Tech courses. Suggestions for improvements in the course and the software are also given.

Students' perspective

On the positive side, all of the students liked the course and most would recommend it to other students. They liked using computers and enjoyed being able to examine flows in real-world-type systems rather than in the idealized systems typically studied in an introductory fluid mechanics course. The few students who had used other CFD software in the past appreciated the ease with which they could interact and modify the governing equations and system parameters by simply using MATLAB-type functions and expressions. These students especially became very competent with the software. All students became more proficient at discussing fluid flows in terms of the physical processes involved rather than just saying things like ' . . . the drag increased when the velocity increased . . .' for example. The students particularly liked the ability to examine and visualize the field variables of the solution, such as velocity, pressure, and temperature, even for simple flows. This gave them a better understanding or 'feel' for the flow than they got by just seeing a single plot or analytical function representing the solution.

Some students thought that the workload for a few of the labs was excessive. This was particularly true for those students who did not have much experience with similar software tools, such as computer-aided-design (CAD) software. For a couple of the labs, such as the one on thermal convection or turbulent flow, finding converged solutions was still difficult and frustrating. However, many of the problems encountered were the result of an inappropriate setting for a system parameter or a boundary condition or an incorrect geometry. These kinds of errors were hard to track down and they usually led to many more hours of work than was necessary to finish to lab.

Even when there were no problems with the model or the solution, some labs took too much time because of the large number of explorations and questions that the students were asked to do. This kind of problem was easily corrected once the instructor had a better idea of the typical amount of time a student needed to work through a lab.

Finally, the lectures on the fundamentals of the finite-element method were given during the last couple of weeks of the course. The students felt that this was too late and that the lectures were not deep enough. It would be better to do these lectures at the beginning of the term right after the COMSOL tutorial, but concurrent with the first couple of labs. This would give the students a better idea of some of the mesh concepts and convergence criteria they were using in the labs.

Instructors' perspective

The main component of the teaching materials used in this course was the computer labs. It was a full-time effort to design the labs and to write the lab guides. Each lab problem had to be an interesting flow with interesting physics. It had to be doable by a student working about ten to fifteen hours per week. The lab as written had to work correctly. Students wasted too much time and became very frustrated if it didn't work because of an error or typo in the lab guide. The exploration questions in the manual had to be carefully posed. The students' level of understanding of the physics is not as high as the instructor's so subtle interactions in the physics are often beyond the ability of the student to notice and appreciate. The physical interactions being explored had to be very basic, easy to see, and easy to explain.

After the basic concept of the lab was discussed in the preparatory lecture, the students were usually excited to begin the lab work. The instructor–student interaction when they had a problem with the lab or a question about the fluid mechanics was usually very good. The most difficult questions though were what to do when the lab results were not going according to the lab guide. As any computer-support person knows, trying to figure out where the user (student) made his mistake is time consuming and frustrating. This is why it is essential that the lab works correctly as written.

As noted by many students, it was clear that the finite-element lectures should be given at the beginning of the term rather than the end. If given earlier, it would also be possible to give an assignment to write a MATLAB program to solve an ordinary differential equation using the finite-element method. An example of this kind was part of the finite-element lectures, but having the student write their own code would be a much better way for the student to understand the technique. One of the flow labs would have to be dropped if this assignment was given and the trade-off for this would have to be evaluated by the instructor. It is not clear that all students would benefit equally from such an assignment, particularly those students who are less familiar with programming and those who do not plan to work full time with CFD.

Future improvements

There are several improvements or additions to the COMSOL software that would help improve the course, help students to become better users of the software, and help them to improve their understanding of the fluid mechanics under question.

First, searching and finding converged solutions of complicated flows can be frustrating and time consuming. What are the best techniques available to do this? Can these techniques be directly implemented into the software? At least, it would be helpful to have such techniques collected and described in a separate section of the documentation.

Better tools are needed to inspect the quality of the solution. The predefined residual variable that is available is not very good. It should be normalized in some way so that it is a quantitative measure of the error in the solution. In the microfluidics lab the geometry was built to micron dimensions. As a result, the residual of the solution was very large and it was difficult to judge the error in the solution from the magnitude of the residual alone.

Finally, improvements in the flow visualization tools are needed when these tools are used to inspect smaller portions of the flow field. The current tools should be easier to use and more intuitive. This is especially true for the particle-tracing tool. This tool can be used to produce particle paths, which are streamlines in steady flows. However, it is often very difficult to make such plots look good and to make sure they include the most interesting flow features, such as regions of recirculating flow. One needs to know where such features exist in order to seed the flow with the particles for visualization. This problem can be circumvented in two dimensions by defining the stream function and making a contour plot. Perhaps this could be done automatically in the software when a stream line plot is selected.

In closing, the course was a success from the perspectives of the students and the instructor. It gives the students the kind of engineering experience that they need in order to decide what career they want to do when they graduate and it looks good on their resume when they are interviewed by prospective employers. Since the students typically enjoy the course and are motivated to learn, the instructor has a more fulfilling interaction with the students. It becomes fun to teach and to watch the students learn and appreciate the physics of fluid flows.

REFERENCES

1. MATLAB R2007a, The Mathworks, Inc.
2. MATLAB R2007a, Function Reference, The Mathworks, Inc.
3. Wolfram Mathematica 6, Wolfram Research, Inc.
4. FLUENT, Fluent Inc.: www.fluent.com
5. COMSOL Multiphysics, COMSOL AB.
6. *ME 4342 Computational Fluid Dynamics*, The George W. Woodruff School of Mechanical Engineering, Georgia Institute of Technology, Atlanta, GA 30332-0405.
7. COMSOL Multiphysics Minicourse, version Sept 2006, COMSOL AB.
8. B. R. Munson, D. F. Young, T. H. Okiishi and W. W. Huebsch, *Fundamentals of Fluid Mechanics*, 6th edn, John Wiley & Sons, Inc., (2009) pp. 309–310.
9. F. Durst, S. Ray, B. Unsal and O. A. Bayoumi, The development lengths of laminar pipe and channel flows, J. *Fluids Eng*., **127**, 2005, pp. 1154–1160.
10. B. R. Munson, D. F. Young, T. H. Okiishi and W. W. Huebsch, *Fundamentals of Fluid Mechanics*, 6th edn, John Wiley & Sons, Inc., (2009) pp. 474–478.
11. M. Van Dyke, *An Album of Fluid Motion*, The Parabolic Press, Stanford, CA (1982) pp. 56, 57, 130,
12. T. von Kármán and H. Rubach, *Über den Mechanismus des Flüssigkeits- und Luftwiderstandes, Collected Works of Theodore Von Kármán*, Vol. I, 1902–1913, Chap. 17, Butterworths Scientific Publications, London, (1956) pp. 339–358.
13. B. R. Munson, D. F. Young, T. H. Okiishi and W. W. Huebsch, *Fundamentals of Fluid Mechanics*, 6th edn, John Wiley & Sons, Inc., (2009) Eq. 8.37, p. 441.
14. P. G. Drazin and W. H. Reid, *Hydrodynamic Stability*, Chap. 2, Cambridge University Press (1981) pp. 32–68.
15. M. Van Dyke, *An Album of Fluid Motion*, The Parabolic Press, Stanford, CA, (1982) pp. 82–83.
16. J. N. Reddy, *An Introduction to the Finite-element Method*, 2nd 3edn, McGraw-Hill, Inc. (1993) pp. 67–103.
17. B. R. Munson, D. F. Young, T. H. Okiishi and W. W. Huebsch, *Fundamentals of Fluid Mechanics*, 6th edn, John Wiley & Sons, Inc., (2009) Fig. 8.5, p. 388.
18. Y. A. Cengel and J. M. Cimbala, *Fluid Mechanics Fundamentals and Applications*, 1st edn, McGraw-Hill Higher Education, (2006) Fig. 8-8, p. 325 and Fig. 8-10, p. 326.
19. P. K. Kundu and I. M. Cohen, *Fluid Mechanics*, 4th edn, Academic Press (2008), Fig. 10.31, p. 389.
20. News Release, NASA Eyes Intricate Pattern on Cloud Street, http://www.jpl.nasa.gov/releases/2002/release_2002_103.html, May 01, 2002.
21. K. Y. Billah and R. H. Scanlon, Resonance, Tacoma Narrows bridge failure, and undergraduate physics textbooks, *Am. J. Phys*., **59**(2), 1991, pp. 118–124.
22. F. M. White, *Viscous Fluid Flow*, 3rd edn, McGraw-Hill Higher Education, (2006), Figs 3–46, p. 198.
23. P. G. Drazin and W. H. Reid, *Hydrodynamic Stability*, Cambridge University Press, (1981) p. 35.
24. F. P. Incropera and D. P. DeWitt, *Fundamentals of Heat and Mass Transfer*, 3rd edn, John Wiley and Sons, Inc., (1990) Eq. 9.51, p. 561.
25. P. G. Drazin and W. H. Reid, *Hydrodynamic Stability*, Cambridge University Press (1981) Table 2.1, p. 51.
26. E. N. Jacobs, K. E. Ward and R. M. Pinkerton, The Characteristics of 78 Related Airfoil Sections from Tests in the Variable Density Wind Tunnel, NACA-TR-460, (1933). Summary located at NACA Airfoil Series, www.aerospaceweb.org/question/airfoils/q0041.shtml

Marc K. Smith is a Professor of Mechanical Engineering at the Georgia Institute of Technology. He received his Ph.D. degree in Engineering Sciences and Applied Mathematics from Northwestern University in August 1982. After post-doctoral work at MIT and Cambridge University and a faculty appointment at the Johns Hopkins University, he moved to Georgia Tech in 1991. His research interests include interfacial fluid mechanics, hydrodynamic stability, surface-tension-driven flows, and boiling. His recent work includes acoustic enhancement of boiling heat transfer and vibration induced droplet atomization. His teaching interests include fluid mechanics, thermodynamics, and numerical methods for engineers.