

Bakker A., Marshall E.M. (2004) *Computational Fluid Dynamics*. Encyclopedia of Chemical Processing. Editor Prof. Sunggyu Lee. Marcel Dekker Inc. New York.

Computational Fluid Dynamics

André Bakker and Elizabeth Marden Marshall
Fluent Inc.

Keywords

Fluid dynamics, turbulence, computational fluid dynamics, CFD, fluidized bed, computing

Introduction

Computational fluid dynamics (CFD) is an engineering field in which flow fields and fields of related scalar variables, such as temperature and chemical species concentrations, are calculated in great detail for the domain of interest. The flow field calculations are performed by solving discretized forms of the equations for the conservation of mass, momentum, energy, and other relevant variables. The final results are detailed maps of the flow field showing fluid velocities, temperatures and species concentrations throughout the system. These results can be used to gain a better understanding of the system of interest. CFD results can illustrate how a piece of equipment operates, how to troubleshoot problems, how to optimize performance, and how to design new equipment.

During the 1950s and 1960s, CFD was used mainly in the aerospace and defense industries. This was followed by the power generation industry (especially the nuclear industry) and the automotive industry in the 1970s. The process industries did not start to use CFD seriously until the 1980s and 1990s. One of the important reasons for the relatively late adoption was that chemical reactors are characterized by so many complex phenomena. These include heat transfer, multiphase flow, and homogeneous and heterogeneous reactions. The development of methods to model these complex physical phenomena has occurred in recent years. Much progress has been made, and CFD is now used to model accurately many pieces of plant equipment, such as fluidized beds, packed bed reactors, cyclone separators, spray dryers, bubble columns, stirred vessels, static mixers, and more.

One way that CFD results differ from traditional analysis methods is in the amount of detail that they provide. Whereas traditional methods, such as design based on correlations of experimental data or overall thermodynamic balances, provide general design guidelines, they do not usually provide insight into the internal workings of the equipment. For example, in a chemical reactor where several competing reactions take place, the final product composition may depend on non-uniformities in the

concentration, flow, or temperature fields inside the reactor. Traditional methods that make use of average values inside the reactor are not able to take into account the effect of such non-uniformities. CFD methods, which typically calculate flow field variables at hundreds of thousands of points inside the reactor to come up with overall reaction rates, are far better suited for the analysis of such systems.

Another difference between CFD and traditional design methods is the minimal reliance of CFD on experimental data and extrapolation of that data to different scales, a process known as scale-up. CFD relies on solving the fundamental equations of motion and conservation. These equations are scale-independent, and can be solved directly for the full-scale equipment.

While CFD provides a number of advantages over other analysis methods, high fidelity CFD can come at a cost. It typically requires the use of specialized CFD software that offers extensive modeling capabilities, the availability of a skilled CFD analyst to set up the model and interpret the results, and high-end computing resources. In recent years, however, the cost of high-end computing has dropped substantially as the speed and power of processors have increased at a similar rate. Clusters of inexpensive computers can now do jobs that previously required supercomputers. In addition, new CFD software products have been introduced with limited capabilities and easy-to-use front ends. These new tools are designed to make CFD available to non-specialists who want quick decision support for product design or trouble-shooting.

In this article, the conservation equations that describe fluid motion are reviewed, and the general methodology behind CFD is summarized. The steps that are typically followed during an analysis are described and illustrated by an example of a reacting fluidized bed.

Conservation equations

CFD is based on the principle of solving conservation equations for all relevant variables. The conservation equations include the transport of the variable throughout the domain, as well as either its creation or its destruction. Conserved variables include:

- Mass,
- Momentum,
- Enthalpy,
- Turbulent kinetic energy,
- Turbulent energy dissipation rate,
- Chemical species concentrations,
- Local reaction rates, and
- Local volume fractions for multiphase problems.

The equation for conservation of mass, also termed the continuity equation, has the form:

$$\frac{\partial \rho}{\partial t} + \frac{\partial}{\partial x_i} (\rho U_i) = 0 \quad (1)$$

Here ρ is the fluid density, and U_i is the fluid velocity in the x_i direction. When an index, such as i , is repeated in the same term, it means that the term is a summation over all possible values of the index. The first term on the left hand side describes the change in fluid density over time, and the second term describes the transport of the fluid. For incompressible fluids, which have a constant density, the continuity equation reduces to the following, simpler form:

$$\frac{\partial U_i}{\partial x_i} = 0 \quad (2)$$

The statement for the conservation of momentum, also known as the Navier-Stokes equation, is as follows:

$$\frac{\partial (\rho U_i)}{\partial t} + \frac{\partial}{\partial x_j} (\rho U_i U_j) = -\frac{\partial p}{\partial x_i} + \frac{\partial}{\partial x_j} \left[\mu \left(\frac{\partial U_i}{\partial x_j} + \frac{\partial U_j}{\partial x_i} - \frac{2}{3} \frac{\partial U_k}{\partial x_k} \delta_{ij} \right) \right] + \rho g_i + F_i \quad (3)$$

The first term on the left hand side describes the variation of the fluid momentum in time; the second term describes the transport of the momentum in the flow (convective transport). The first term on the right hand side describes the effect of gradients in the pressure p ; the second term, the transport of momentum due to the molecular viscosity μ (diffusive transport); the third term, the effect of gravity g ; and in the last term, F_i lumps together all other forces acting on the fluid. Techniques for solving the set of four equations (one continuity and three momentum equations) are discussed in a later section of this article. When the flow is compressible, it is usually necessary to close the system of equations listed above using a thermodynamic equation of state (such as the ideal gas law) that calculates the density as a function of temperature and pressure.

For steady state flows, the first term on the left hand side is zero. For turbulent flows, which are unsteady, it saves time and computer resources to calculate a single, time averaged flow field only, instead of solving for the full time dependent flow field. In order to do this, the velocity is written as the sum of a constant and fluctuating component ($U_i + u_i'$), and the momentum equation is averaged over time, resulting in the following modified conservation equation for momentum:

$$\frac{\partial}{\partial x_j} (\rho U_i U_j) = -\frac{\partial p}{\partial x_i} + \frac{\partial}{\partial x_j} \left[\mu \left(\frac{\partial U_i}{\partial x_j} + \frac{\partial U_j}{\partial x_i} - \frac{2}{3} \frac{\partial U_k}{\partial x_k} \delta_{ij} \right) \right] + \frac{\partial}{\partial x_j} (-\rho \overline{u_i' u_j'}) + \rho g_i + F_i \quad (4)$$

A new term is introduced, the so-called Reynolds stresses $\overline{u_i' u_j'}$. The overbar denotes a time average. This term is the correlation between the turbulent velocity fluctuations u_i' and u_j' , and it describes the transport of momentum in the mean flow due to turbulence. This term is difficult to model, and over the years a variety of turbulence models have been developed. Turbulence models are necessary for calculating time-averaged flow fields directly, without first having to calculate a fully time-dependent flow field and then doing time averaging. The use of turbulence models is therefore much more computationally efficient. A detailed discussion is beyond the scope of this article, but it is important to note that not all turbulence models are equally suited for all types of flow. Table 1 summarizes the most common turbulence models and their properties.

For the other conserved quantities, conservation equations of a form similar to the momentum conservation equation are solved. These usually include a time-dependent term, a convective transport term (describing the transport of the variable due to the mean flow), a diffusive transport term, and a generalized source term describing all other physical effects.

Table 1: Summary Of Turbulence Models

Model	Description, Advantages and Disadvantages
Standard k-ε	The most widely used model, it is robust, economical, and time-tested. The Reynolds stresses are not calculated directly, but are modeled in a simplified way by adding a so-called turbulent viscosity to the molecular viscosity. Its main advantages are a rapid, stable calculation, and reasonable results for many flows, especially those with a high Reynolds number. It is not recommended for highly swirling flows, round jets, or for flows with strong flow separation.
RNG k-ε	A modified version of the k-ε model, this model yields improved results for swirling flows and flow separation. It is not well suited for round jets, and is not as stable as the standard k-ε model.
Realizable k-ε	Another modified version of the k-ε model, the realizable k-ε model correctly predicts the flow in round jets, and is also well suited for swirling flows and flows involving separation.
Spalart-Allmaras	The Spalart-Allmaras model was developed for external flows in aerospace applications. It provides good answers for attached flows and flows with mild flow separation. It is not commonly used for process industry applications.
RSM	The full Reynolds stress model provides good predictions for all types of flows, including swirl, separation, and round and planar jets. Because it solves transport equations for the Reynolds stresses directly, longer calculation times are required than for the k-ε models.
LES	Large eddy simulation is a transient formulation that provides excellent results for all flow systems. It solves the Navier-Stokes equations for large-scale turbulent fluctuations and models only the small-scale fluctuations (smaller than a computational cell). Because it is a transient formulation, the required computational resources are considerably larger than those required

	for the RSM and k- ϵ family of models. In addition, a finer grid is needed to gain the maximum benefit from the model and to accurately capture the turbulence in the smallest, sub-grid scale eddies. Analysis of LES data usually requires some degree of planning in advance of building the model.
--	---

CFD Methodology

Overview

The general methodology followed during a CFD analysis is comprised of the following steps:

1. Geometry creation
2. Grid generation
3. Defining the model by specifying physical models, boundary conditions, and initial conditions
4. Specification of the numerical methods to be used for solving the model equations
5. Performing the calculation
6. Analysis and interpretation of the results

Preprocessing

The first steps, geometry and grid creation, are also referred to as preprocessing. In the first step, geometry creation, a computer model is created of all flow passages and those solid components that need to be included in the heat transfer analysis. The geometry is created in a way that is similar to three-dimensional solid model creation in computer aided design (CAD) software, and in software used for structural analysis. In fact, it is becoming more and more common to import CAD geometries into the CFD preprocessor instead of creating the whole geometry from scratch using the preprocessor. An important difference between the computer models constructed for structural analysis applications and those constructed for CFD is that the former require details of the solid parts, whereas the latter require details of the flow passages, or the space enclosed by the solid parts. Hence, when CAD models are imported into CFD preprocessing software, new volumetric entities for the flow passages usually have to be created from the imported solid surfaces. Furthermore, certain geometric details that can be present in the CAD model (e.g. boltheads, weldlines, and flange details) but that are not relevant to the CFD analysis can (and should) be removed.

Although it is common to use three-dimensional geometries, some problems can be solved suitably in two- dimensions. For example, the flow in a straight circular pipe in the absence of gravity or asymmetric heating is independent of the angular position. In such cases, it is sufficient to create a two-dimensional geometry, taking into account the axial and radial directions, but ignoring the circumferential (“tangential”) coordinate.

Once the geometry has been created, a computational grid, or mesh, has to be generated. During grid generation, the flow domain defined in the geometric model is subdivided into a large number of computational cells. The CFD solver will later solve conservation equations inside each computational cell, keeping track of what enters and leaves through each of the cell faces. Generating a suitable grid is therefore of utmost importance for the success of the model calculations. Different cell types with different shapes can be used to this effect. For 2D boundary layer flow, for example, it is better to use up to ten layers of quadrilateral cells than triangular cells in the near-wall region. To accurately capture a 2D jet entering a plenum, at least ten cells should be used for the jet cross section to avoid excessive spreading. These guidelines can be extended to 3D cases, and depend on the overall dimensions and requirements of the case at hand. In most software, the initial grid can be refined or coarsened during the course of the solution, to better resolve gradients or capture other effects, if needed. Figure 1 shows some of the most commonly used types. The corner points of the cells are referred to as nodes. For two-dimensional models it is common to use triangular or quadrilateral cells. For three-dimensional models it is common to use hexahedral, tetrahedral, prismatic, or pyramid cells. Of course, these are not all of the possible cell types. For example, clipping corners of hexahedral cells can create other cell shapes. To accommodate such shapes, some CFD software allows the use of arbitrary polygonal cells, with any number of faces.

There are different ways in which the cells can be distributed throughout the domain. When the mesh consists solely of quadrilateral or hexahedral cells distributed throughout the domain in a regular, stacked pattern such as that shown in Figure 2, the mesh is referred to as a structured mesh. The use of such structured meshes, in which all cells are referenced by an I, J, K index in the computational domain, was common in the past because of the advantages it offered in the form of simplified programming and straightforward data interpolation routines. The main disadvantage of using structured meshes lies in the difficulties they pose in representing complex shapes. Now it is more common to use unstructured meshes, such as that shown in Figure 3. In unstructured meshes, cells can be arranged in an arbitrary fashion. The user can choose the cell arrangement that is most suitable for reproducing the details of the geometry. Unstructured meshes can use only one cell type (e.g. be all hexahedral or all tetrahedral), or mix different cell types as shown in the mesh in Figure 3. Meshes that combine different cell types are usually referred to as hybrid meshes.

For very complex geometries, meshing can be simplified by subdividing the flow domain into different volumetric zones that can be meshed independently, choosing the best mesh type for that particular zone. The meshes for the different zones then can be combined into one mesh, usually called a multi-block mesh. When combining the meshes for the different zones, the cells and the nodes at the shared faces do not necessarily have to be lined up exactly. Multi-block meshes where the nodes at the shared faces do line up are called conformal meshes, and those where they do not line up are called non-conformal meshes.

Model definition

Once the mesh has been created, the scope of the problem to be solved has to be specified. This step includes specifying the material properties, defining the physical models, setting the boundary conditions, and specifying the initial conditions.

A wide range of physical models is available in most commercial CFD software. At a minimum the flow field will be calculated by solving the conservation equations for mass and momentum. In addition to flow, many of the problems encountered in the process industry involve heat transfer as well. For such applications, the temperature field can also be calculated, which is commonly done by solving a conservation equation for enthalpy. For problems involving chemical reaction, the transport equations for the chemical species involved in the reaction(s) will be solved. The creation and destruction of the species due to the reaction are modeled by means of source terms in these equations. The reaction rates determining these source terms are calculated locally, based on values of species concentrations and temperature at each variable storage site (a point at the center of each computational cell). The quality of the CFD predictions for reacting systems depends largely on the availability of accurate reaction rate constants. Many real life systems involve a certain number of main reactions and another number of side reactions that progress at much slower rates. When the total number of reactions is very large, or when the reaction constants for the less important side reactions are not known, the reaction model is often simplified to include only the most important reactions.

Multiphase flows are common in many types of process equipment. In such equipment there is usually one main continuous phase (a primary phase), and one or more dispersed (or secondary) phases. Multiphase flows can be modeled in a number of different ways. The most straightforward method is to calculate a single-phase flow field and then to calculate the trajectories of a finite number of particle streams for the dispersed phase that include the effects of drag, buoyancy, and turbulence. (Here the term particle is used in the most general sense; it includes drops and bubbles as well.) This so-called Lagrangian particle tracking calculation shows how the particles are dispersed throughout the domain. Physical phenomena such as particle evaporation, boiling, or surface reaction can be included in the tracking algorithm as well. The effect that the particles have on the flow field can be modeled by including particle trajectory-dependent source terms in the momentum (or other) equations for the fluid. Particle tracking models are best suited for systems with low volume fractions ($< 10\%$) of the dispersed phase. Examples of systems that are commonly modeled using this method are spray dryers and liquid fuel or coal combustors.

Eulerian multiphase flow models work differently. Rather than compute individual particle trajectories, the primary and secondary phases are modeled as interpenetrating continua using separate sets of conservation equations. In addition to a set of momentum and continuity equations for each phase, a volume fraction equation for each secondary phase is also solved. Momentum exchange terms are included that depend on the buoyancy, drag, and lift of the secondary phase particles. CFD predictions are for the flow fields for both phases, which are usually different, depending on the system characteristics. Eulerian multiphase flow models can be formulated for any number of

phases. Break-up and coalescence can be included for droplets or bubbles, along with mass transfer between the phases. When one of the secondary phases consists of a granular material, a maximum packing density can be prescribed for that phase, and a special treatment is used for the granular viscosity and granular pressure. Eulerian multiphase flow models are commonly used for fluidized beds, stirred vessels, and bubble columns.

Other types of multiphase flows are free surface flows, where there is a well-defined interface between two continuous phases. Examples of such flows can be found in liquid separators, unbaffled mixing vessels where surface deformation occurs when a central vortex forms, mold filling applications, and blow molding applications. With free surface flows, mass and momentum conservation equations are solved, along with an equation that describes the movement of the interface.

The material properties that have to be specified depend on the physical models chosen. For flow field calculations only, it is usually sufficient to specify the density and the viscosity. The fluid can be incompressible (constant density) or compressible (either a compressible liquid or an ideal or non-ideal gas). The viscosity can be either Newtonian or non-Newtonian (independent or dependent on shear rate, respectively). Any type of non-Newtonian fluid can be modeled, including shear-thinning fluids, yield stress fluids, and visco-elastic fluids, assuming that the viscosity behavior is known. Unfortunately, this is not always the case. For fluids with complex rheological behavior such as paper pulp, dough, and other visco-elastic fluids, it is often difficult to determine the rheology. In such cases, simplified rheologies often have to be used, and it is not uncommon to model the fluid as Newtonian using an effective average viscosity that is known to give good results for such quantities as pressure drop or torque predictions in rotating equipment.

When heat transfer is involved, the density and viscosity can be functions of temperature, and the thermal properties of the fluid have to be prescribed as well. These properties include heat capacity, thermal conductivity, molecular weight (for gases), and the thermal expansion coefficient. For problems involving chemical species, all physical properties have to be specified for all of the species, along with a method to calculate the average property for mixtures in each cell using the properties of the component species.

To complete the model specification, boundary conditions have to be specified. These describe the flow conditions at the domain boundaries. At flow inlets one can usually specify the fluid velocity, a mass flow rate, or an inlet pressure. Depending on the problem definition, the inlet temperature, species concentrations, turbulence properties, and volume fractions of any secondary phases must also be supplied. At flow exits, one usually specifies an outlet pressure, and if entrainment through a flow exit is anticipated, the exit should either be moved downstream or have boundary conditions carefully specified for the entrained flow. For problems with only a single flow exit – far from recirculation zones - no such boundary conditions are required. One can also set guidelines regarding how much of the mass entering the domain from the inlet(s) exits through each of the outlets. At the domain walls, the wall can be prescribed to move at

exactly the same velocity as the adjacent fluid with no imposed frictional drag (a so-called slip condition), or the wall can be assigned the more common no-slip condition, where a constant wall velocity (tangent to the fluid flow), usually zero, is set. The shear stress at the wall can also be prescribed. When the problem involves heat transfer, boundary condition options include a specified heat flux (which can be zero to model well insulated, or adiabatic, walls), a specified temperature, or an external heat transfer boundary condition that may include heat losses due to radiation.

For problems involving repeating geometry and flow conditions, it often is sufficient to model the flow in only one section of the domain that is identical to that in other sections of the same size. At the appropriate boundaries of that section, either periodic or symmetry boundary conditions can be prescribed.

For problems where the fluid behavior is steady and the final solution is independent of the initial conditions, calculations can usually be carried out in a steady state fashion. For problems that describe transient behavior or that require a time-marching solution technique, initial conditions must be supplied in addition to boundary conditions. These include the state of the flow field, concentration fields, and temperature fields at time zero at all locations in the domain.

As part of the specification of the physical models and boundary conditions, the user has to specify the body forces on the fluid. Body forces are those forces that act on the entire volume of fluid throughout the domain. These forces include gravity, electromagnetic forces (if relevant), and the Coriolis force for rotating domains.

Numerical methods

There are a variety of numerical methods available to solve the conservation equations. The most commonly used method in commercially available CFD software today is the finite volume method. Excellent descriptions of this method can be found in Patankar (1980) and in Versteeg and Malalasekera (1995). With the finite volume method, an integral form of the conservation equations is solved by performing a mass and momentum balance over all faces of each computational cell. There are, however, many other methods available, such as the finite element method (where the equations are solved in differential form, instead of in integral form), vorticity based methods (where conservation equations for vorticity are solved instead of those for momentum), and spectral element methods (where the equations are transformed into Fourier space before being solved). The main advantage of the finite volume method over other methods is that it conserves mass on coarse as well as fine meshes, including those with large numbers (tens of millions) of computational cells, and performs well for turbulent flows.

One disadvantage of the finite volume method, however, is that because velocities and other variables have to be interpolated from the cell centers (the storage sites) to the cell faces so that face fluxes can be computed, a certain amount of numerical error is introduced. These errors tend to have the effect of decreasing the sharpness of gradients in the flow when compared to real life, similar to how diffusion smooths out gradients in

all scalar fields. This error is therefore termed numerical diffusion, and it is present in all CFD simulations performed using the finite volume method. It can be minimized by using meshes that are finer in regions where steep gradients in the flow occur, and by choosing accurate interpolation methods (also called numerical schemes) that result in a small error (usually at the expense of increased computation time). However, for flows where gradients tend to be relatively small, such as in time averaged turbulent flows, the numerical diffusion may be insignificant. The finite element method is not as prone to numerical diffusion, and therefore tends to give solutions that are more accurate for a given number of cells for flows that have sharp gradients. Examples of such flows are laminar flows at very low Reynolds numbers. The finite element method is not as computationally efficient as the finite volume method for turbulent flows or for meshes with large numbers of cells. Most applications in the process industries involve turbulent flows in complex geometries that require fine meshes. The finite volume method is used, therefore, for most applications. The finite element method tends to be the better choice for applications involving very viscous fluids, such as those encountered in blow molding and extrusion processes.

Solution

Once the solution method has been chosen, the actual calculations are performed using an iterative procedure. The solution procedure is as follows. Starting with an initial solution provided by the user (which may be a value of zero for all variables throughout the domain), the solver determines the error in the mass and momentum balances for each computational cell. This error is also known as the residual. Based on that error, the solver determines what adjustments need to be made to the various variables to improve the balances for each cell. These adjustments are applied to all cells in the domain, and the procedure is repeated until the error in the solution is smaller than a preset tolerance specified by the user. At that time, the solution is considered converged and the calculation is halted. There are many different ways in which this concept can be implemented, and many mathematical techniques are available to enhance the rate of convergence and to improve the solver's ability to converge to a stable and accurate solution. A discussion of these techniques is beyond the scope of this article.

Postprocessing

The final step in a CFD simulation is to analyze and interpret the results, a step also referred to as postprocessing. Typical CFD results include values for the pressure, the three velocity components, turbulence properties, temperature, and species concentrations at all grid cells in the domain (which is typically on the order of several hundred thousand in today's applications). The total amount of available data is thus extremely large and detailed. In order to be useful to an engineer, the data has to be presented in an understandable way.

The objective of the simulation might be to obtain quantitative estimates for certain global variables, such as an overall reaction rate, the temperature in a certain location, or the pressure drop across the domain. These can be calculated from the CFD solution in a

relatively straightforward manner. The results also can be presented graphically, and there are many options available for doing so.

Iso-surfaces can be constructed by connecting all points in the flow domain that have the exact same value of a given variable. These surfaces can be colored by contours of any other variable. Iso-surfaces are useful to visualize certain flow field features, such as the central vortex in the cyclone separator shown in Figure 4. Contour plots can be used to show the local values of scalar variables. A cross-section of the flow domain is created and colored by the local value of the variable of interest, e.g. temperature. A color scale is used indicating how each color corresponds to a certain value of the variable. Such plots can be used to quickly see the variation of important variables throughout the domain. An example of a contour plot is found in Figure 5, which is described in the example in the following section.

A commonly used graphical visualization of the flow field is through plots of velocity vectors. An example, also taken from the simulation described in the following section, is shown in Figure 6. The velocity vectors point in the direction of the fluid flow where the vector originates. The length of the vector can be made proportional to the local velocity magnitude, such that longer vectors indicate higher velocities. The vectors can also be colored by the velocity magnitude, or by any other variable, as desired.

Other methods to visualize the flow field make use of flow lines, of which different types are available. Streamlines are useful in 2D simulations. They are curves that are tangent to the flow field and whose spacing is such that a fixed mass of fluid passes between any pair of lines. For 3D geometries, pathlines can be constructed by calculating the trajectories of massless particles that follow the mean flow. Streaklines are formed by a series of particles that are continuously injected into the flow from a number of fixed locations. For steady 2D flows, streamlines, pathlines, and streaklines are the same. However, for unsteady (time-varying) flows, they can be very different.

Other types of flow lines also can be used. Oilflow lines are pathlines that are constrained to a given boundary surface. When calculating the pathline, velocity components that are tangent to the given boundary surface are included and normal velocity components are ignored. This is useful for visualization of the flow field over an object. Timelines are formed by placing a series of particles along a line in the flow field at time zero. The particles are tracked for a given amount of time, and the curve connecting the particles is redrawn. Timelines also are called material lines, and they are commonly used to study the amount of stretching and deformation that occurs in mixing applications.

EXAMPLE: REACTING FLUIDIZED BED

One example that illustrates many of the concepts described in this article is the conversion of ozone to oxygen in a fluidized bed. The particles in the bed serve as a catalyst where surface reactions take place. To model this complex process, the commercial software FLUENT is used. The fluidized bed is simulated using the Eulerian granular multiphase model, and a reaction is specified in the bed to account for the decomposition of ozone and subsequent production of oxygen.

A 2D axisymmetric model of a column 0.23 meters in diameter and 0.25 meters high is used, following Fryer and Potter (1976). At rest, the bed is 0.12 meters high, and contains sand particles impregnated with iron oxide, 117 microns in diameter with a density of 2.65 g/cm^3 . Two gas phase species (ozone and oxygen) are specified. Ozone enters from inlets at the base of the bed at speeds that range from 4 to 14 cm/sec. The gas lifts and separates the particles in the bed, allowing for better interaction between the ozone and iron oxide at the particle surfaces. The decomposition rate is expressed as

$$K = 1.57 \alpha [\text{O}_3] \quad (5)$$

where α is the volume fraction of the catalyst and $[\text{O}_3]$ is the concentration of ozone. After the reaction has taken place, a mixture of ozone and oxygen exits through the top of the column. Of particular interest are the conversion characteristics of the bed, where conversion is defined as

$$1 - \frac{C_{out}}{C_{in}} \quad (6)$$

where C_{out} and C_{in} are the outgoing and incoming concentrations of ozone, respectively. If C_{out}/C_{in} is small, conversion is high and the bed is operating efficiently.

Figure 5, discussed earlier, shows contours of the gas volume fraction 0.6 seconds after the flow is initiated. Light gray (or red) indicates regions of pure gas, and dark gray (or blue) indicates regions where the sand volume fraction is maximum. The flow field is the same whether the reaction in the gas phase takes place or not. Bubbles are formed near the bottom of the bed and migrate upwards. The simulated bubble shape and size are grid dependent. If a coarse mesh is used, the bubbles are fewer in number and rounder. On finer meshes, the bubbles are denser and more irregular. The number and size of the bubbles has a significant impact on the conversion. The more bubbles in the domain, the greater the interaction between the ozone and sand particles, so the higher the conversion.

Figure 6 illustrates the velocity fields for the gas (left) and granular (right) phases after 0.6 seconds of operation. The vectors are plotted on every fifth grid cell to make the display easier to read. The velocity fields are different in several ways. First, there is no

granular motion on the top of the bed, where only the oxygen and ozone gases are escaping. Second, while some of the larger flow features are the same for both phases, some of the smaller ones are not. For example, a rather large recirculation pattern dominates the gas flow on either side of the bed. In the sand phase, perhaps owing to the increased inertia of the material, the recirculation patterns are not as strong, and have slightly different shapes.

In Figure 7, curves for C_{out}/C_{in} (pink) and gas holdup (red) are plotted as functions of the gas velocity. The gas holdup is defined as the ratio of the gas volume in the bed to the total volume of the bed. The figure shows that the gas holdup increases with gas velocity up to a point, after which saturation occurs. This is because at the low end, the increasing gas velocity forces the bed to lift more. When saturation occurs, the bed can no longer rise and hold additional gas.

The curve of C_{out}/C_{in} is in good agreement with data from Fryer and Potter (1976). The results follow the trends in the gas holdup. At low gas speeds, the residence time is long, so there is more time for the ozone to come into contact with the particles. C_{out}/C_{in} is small, so conversion is high. At high speeds, the residence time reduces, and the rate of conversion tends toward a constant value.

CONCLUSION

Before embarking on any CFD simulation, it is important to realize that the care that goes into the setup of the model is directly related to the quality of results that are obtained. The grid should be fine enough and of optimal quality to capture the flow details of interest. The choice of physical models, for turbulence or multiphase, for example, should be based on the flow regime and expected flow features. Boundary conditions and physical properties should be chosen carefully so that they represent the full range of anticipated behavior. Approximations – what to keep in the model and what to neglect – can also play a significant role in the type of effects that the simulation can capture. Finally, once the setup is complete and the simulation is running, it is important to make sure that proper convergence has been obtained before using the results for any type of careful analysis.

If these issues are heeded, the potential benefits that can be extracted from a simulation are numerous. With the computational resources available today, both in terms of speed and power, engineers are more able to use high density grids and multiple, detailed physical models than at any time in the past. In the years to come, computing power only will continue to increase, so CFD results will become more and more accurate as turnaround times shorten. Interoperability between CFD and other computer-aided engineering (CAE) packages is on the rise. In many industries, this has led to a recent expansion of CFD from its use in the analysis and design departments to being a major contributor to automated plant operation and process optimization.

Acknowledgements

The authors wish to acknowledge their colleagues who have helped to make this work possible: Dr. Bin Yang, Dr. Wei Zhou, Dr. Ahmad Haidari, Dr. Ravi Prasad, Dr. Mike Slack, and Mr. John Murray.

References

Fryer, C., and Potter, O.E., *Experimental Investigation of Models for Fluidized Bed Catalytic Reactors*, AIChE J., **22**, p. 38-47 1976.

Patankar, S. V., *Numerical Heat Transfer and Fluid Flow*, Hemisphere, Washington, DC, 1980.

Versteeg, H. K., and W. Malalasekera, *An Introduction to Computational Fluid Dynamics: The Finite Volume Method*, Longman Scientific & Technical, Essex, UK, 1995.

Notation

α	Volume fraction of gas in fluidized bed (-)
C_{in}	Inlet concentration (kg/m ³)
C_{out}	Outlet concentration (kg/m ³)
δ_{ij}	Kronecker delta function (-)
ε	Turbulent kinetic energy dissipation rate (m ² /s ³)
F_i	Net force in the i direction (N)
g	Gravitational acceleration (m/s ²)
k	Turbulent kinetic energy (m ² /s ²)
K	Rate of ozone decomposition (kg-mol/m ³ -sec)
μ	Molecular viscosity (kg/m-s)
p	Pressure (Pa)
ρ	Density (kg/m ³)
t	Time (s)
u_i'	Fluctuating component of the velocity in the direction i (m/s)
U_i	Velocity in the direction i (m/s)
x_i	Spatial coordinate in the direction i (m)

Figure Captions

1. Examples of three-dimensional element shapes that can be used in CFD simulations
2. An example of a two-dimensional structured grid
3. An example of a three-dimensional unstructured grid
4. An iso-surface of axial vorticity in a cyclone, colored by velocity magnitude, is used to show the central vortex; pathlines are used to show the swirling flow
5. Gas volume fraction in a reacting fluidized bed 0.6 seconds after operation begins
6. Velocity vectors of a) the gas phase and b) the granular phase (sand) in a fluidized bed after 0.6 seconds of operation; the sand motion is different from that of the gas, and is restricted to the bed region
7. Ozone conversion (pink) and gas holdup (red) as functions of gas velocity for the fluidized bed

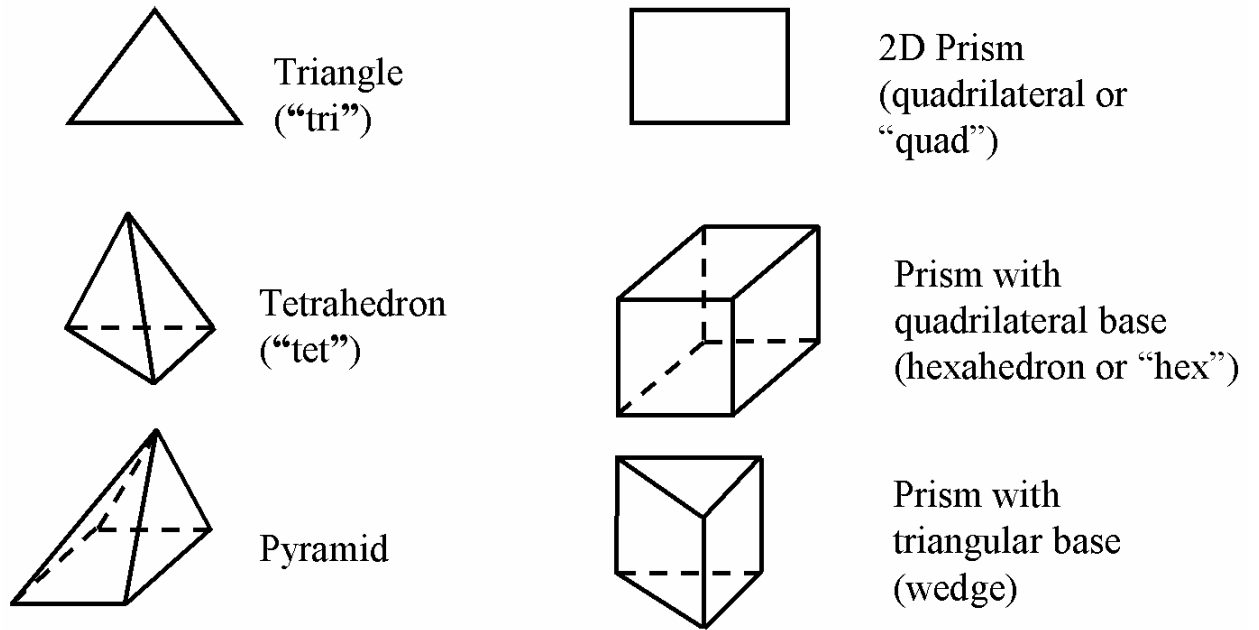


Figure 1. Examples of three-dimensional element shapes that can be used in CFD simulations.




Figure 2. An example of a two-dimensional structured grid.




Figure 3. An example of a three-dimensional unstructured grid.




Figure 4. An iso-surface of axial vorticity in a cyclone, colored by velocity magnitude, is used to show the central vortex; pathlines are used to show the swirling flow




Figure 5. as volume fraction in a reacting fluidized bed 0.6 seconds after operation begins.

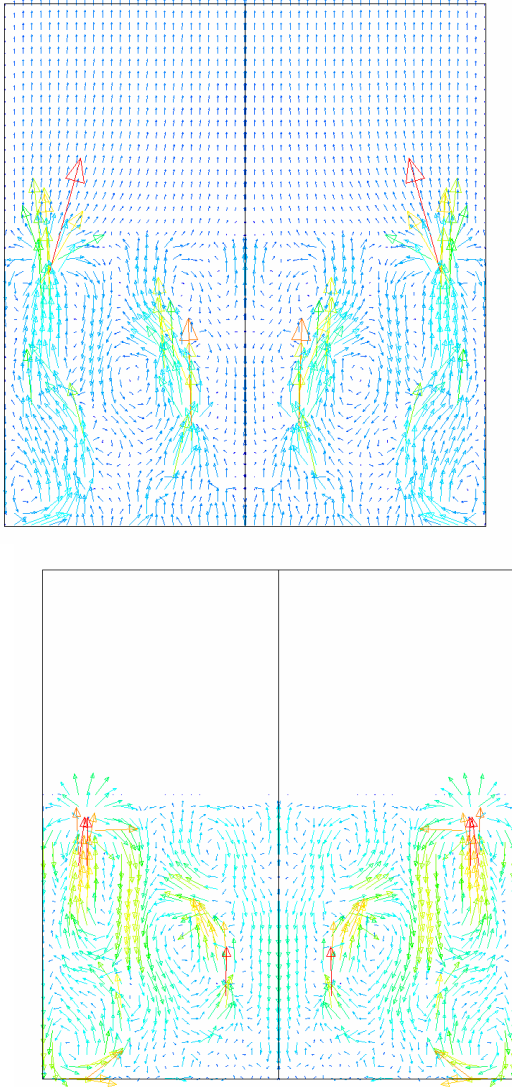


Figure 6. Velocity vectors of a) the gas phase and b) the granular phase (sand) in a fluidized bed after 0.6 seconds of operation; the sand motion is different from that of the gas, and is restricted to the bed region

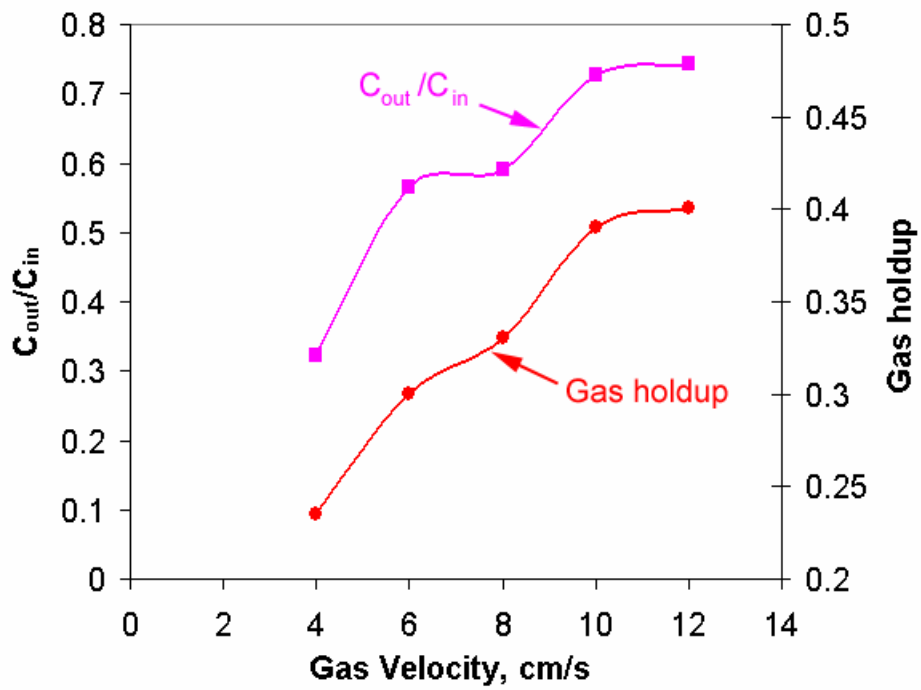


Figure 7. Ozone conversion (pink) and gas holdup (red) as functions of gas velocity for the fluidized bed