A CFD PRIMER: WHAT DO ALL THOSE COLORS REALLY MEAN?

James M. Sorokes
Principal Engineer
Dresser-Rand business
Part of Siemens Power and Gas
Olean, New York, USA

James Hardin
Senior Engineer
The Elliott Group
Jeannette, Pennsylvania, USA

Brad Hutchinson
Global Industry Director
Industrial Equipment & Rotating Machinery
Ansys, Inc.
Waterloo, Ontario, Canada

James M. "Jim" Sorokes is a Principal Engineer with the Dresser-Rand business with more than 40 years of experience in the turbomachinery industry. Jim joined Dresser-Clark (now the Dresser-Rand business) after graduating from St. Bonaventure University in 1976. He spent 28 years in the Aerodynamics Group, became the Supervisor of Aerodynamics in 1984 and was promoted to Manager of Aero/Thermo Design Engineering in 2001. While in the Aerodynamics Group, his primary responsibilities included the development, design, and analysis of all aerodynamic components of centrifugal compressors. In 2004, Jim was named Manager of Development Engineering whereupon he became involved in all aspects of new product development and product upgrades. In 2005, Jim was promoted to principal engineer responsible for various projects related to compressor development and testing. He is also heavily involved in mentoring and training in the field of aerodynamic design, analysis and testing.

Jim is a member of AIAA, ASME and the ASME Turbomachinery Committee. He has authored or co-authored more than fifty technical papers and has instructed seminars and tutorials at Texas A&M and the Dresser-Rand business. He currently holds four U.S. patents and has several others patents pending. He was elected an ASME Fellow in 2008. Jim was also selected as a Dresser-Rand Engineering Fellow in 2015.

James Hardin is a Senior Engineer in the Research and Development department at Elliott Group, in Jeannette, PA, where he performs computational fluid dynamics (CFD) and other aerodynamic analyses for turbines and compressors. Previous experience includes CFD and other analyses on shipboard propulsion and piping systems with Westinghouse Electric Corporation, and turbine design support and testing at Elliott Group. He has 35 years of engineering experience in aerodynamics and fluid systems.

Mr. Hardin received a B.S. degree (Mechanical Engineering, 1981) from Carnegie-Mellon University, and is a registered Professional Engineer in the State of Pennsylvania.

Brad Hutchinson is a member of the Industry Marketing team at ANSYS and is responsible for the turbomachinery portfolio. He has a Ph.D. in Mechanical Engineering and has worked in simulation for more than 30 years, starting as a CFD developer. Much of that time he has spent working with turbomachinery companies worldwide, understanding their issues and processes and advising on both the development and application of ANSYS turbomachinery tools and capabilities.

ABSTRACT

This tutorial provides a general overview of computational fluid dynamics (CFD). It is not intended for CFD experts but rather for those seeking answers to questions such as:

- What is the role of CFD in the design / analysis process?
- What is CFD?
- What benefits can be derived from CFD?
- What are the typical deliverables from CFD?

In addition to answering the above questions, the session will also offer several sample cases illustrating some practical applications of CFD in day-to-day design, analysis and trouble-shooting processes.

INTRODUCTION

Over the past 25 years, three-dimensional or 3-D computational fluid dynamics (CFD) has become increasingly important in the design and analysis or troubleshooting of process centrifugal compressors. Once used exclusively by academia, the aircraft industry and the space program, 3-D CFD now plays a vital role in the day-to-day design process for the full range of turbo products. Aero designers use the codes to assess new component designs and/or combinations of designs. Analysts also apply CFD when trying to determine the cause of test stand or field performance issues, and it has become increasingly common for OEMs to use CFD-generated performance maps or component loss models when attempting to convince end users of the viability or plausibility of a proposed equipment solution. In fact, it is becoming increasingly rare to attend a design review, design audit, or root cause meeting without seeing at least one or two figures containing CFD results.

Regardless of the intended purpose for the CFD study, the governing equations, modeling techniques, solution schemes, turbulence models, and boundary conditions used, as well as the solution deliverables provided (*i.e.*, scalar, vector, contour plots, *etc.*) are essentially the same. Those skilled in the art; that is, those who know the strengths and limitations of CFD; have little difficulty wading through the maze of contour plots, velocity vector distributions, streamline / stream tube diagrams and the like. But what of the person who is not skilled in the science of CFD? What are they to make of the myriad of plots and figures? How much confidence can they have in the results being presented? What is the science behind these complex simulations? And what exactly do all of those colors really mean?

This tutorial will attempt to answer those questions and many others. The main objective is to provide an overview of computational fluid dynamics: what it is… how it's used… and what one might learn from it. This tutorial is <u>not</u> intended to teach attendees how to use CFD but rather to:

- give a very high-level overview of the mathematics, governing equations, and solution schemes;
- introduce the nomenclature or terminology used by "CFDers";
- summarize the types of analyses that can be performed;
- review the factors that can influence the accuracy of CFD results;
- help those unskilled in the art to make sense of the CFD-based information that might be presented to them; and
- provide examples of some of the practice uses for CFD.

THE ROLE OF CFD

Before discussing what CFD is, it might be helpful to explain where it fits in the design / analysis process. Prior to the late 1980s, 3-D CFD was not part of the day-to-day design process for industrial centrifugal compressors. Design personnel relied on one-dimension tools; *i.e.*, so-called 1-D codes; and streamline curvature methods (often called 2-D codes) to develop the meridional (a.k.a. cross-sectional or hub and shroud) geometries and blade / vane shapes. These codes certainly fall under the broad category of computational fluid dynamics but the term CFD is used almost exclusively these days for advanced 3-D solvers. The 1-D and 2-D codes were effective in developing high performance turbomachinery. Many OEM products still offered today were designed using these tools. However, as their names suggest, they were incapable of resolving the three-dimensional nature of turbomachinery flow. It was incumbent on the aerodynamic designer to recognize flow patterns in the 2-D velocity diagrams that might suggest problems. Consequently, designers would occasionally encounter unpleasant surprises simply because: (a) the analytical tools being used could not identify flow abnormalities within components or (b) the analyst failed to recognize potential issues in the 2-D results. Further, most 1-D and 2-D codes were also incapable of analyzing complex, non-axisymmetric components such as inlets, sidestreams and/or volutes. Therefore, before the introduction of CFD in the process turbomachinery industry, the only way to validate or assess the performance of these non-axisymmetric components was to gather data in production compressors or scale model test vehicles.

Though 1-D and 2-D codes and test rigs are still used today, CFD codes give designers and analysts the ability to dig much deeper into the 3-D flow field within components and to identify flow anomalies that might have gone undetected by 1-D and 2-D software systems. Further, running 3-D computer simulations is much less costly than doing rig or scale model testing. That is, constructing a computer-based model and tweaking or optimizing the design numerically is less expensive than building a test vehicle with metal or other materials and having to repeatedly modify the physical model until an acceptable design is obtained (often referred to as the "cut 'n' try" approach). However, it is vitally important to remember that virtually all CFD simulations are done with grids generated using drawings and/or solid model; *i.e.*, with idealized or perfect geometries. Conversely, real compressor flow paths deviate from drawings or solid models because of manufacturing and assembly tolerances as well as geometric deviations due to thermal growth, mechanical deflections or the like. There are also compromises related to gas properties and turbulence models. Therefore, unless rigorous and time-consuming steps are taken to account for all of these real world considerations, CFD simulations remain but a mathematic approximation of the real world. Despite these shortcomings that limit CFD's ability to replace rig testing, such simulations have still become a valuable screen or design filter that helps OEMs (and others) avoid expensive tests of inadequate designs.

The introduction of transient or unsteady simulations further enhanced CFD's role in the design process. These simulations capture the time-dependent aspects of the compressor flow field; *i.e.*, such as the impact of impeller blades rotating past stationary diffuser vanes. Such analyses are far more computationally intensive and take significantly more time to achieve convergence. It might take days or weeks to get a simulation for a single flow condition. This might be acceptable in the gas turbine industry or academia, where design projects can take years to complete. However, a two-to-three week execution time for a single flow condition is impractical in the industrial compressor industry. Still, situations arise when the extra rigor can be justified, such as when centrifugal compressor aero components will experience highly unsteady pressure loads during field operation; *i.e.*, high pressure compressors with vaned diffusers or high power consumption compressors. Given the time required to conduct such simulations, designers will try to minimize the number of runs required.

Today, CFD is a standard step in the design process, often portrayed as a design pyramid (Figure 1).

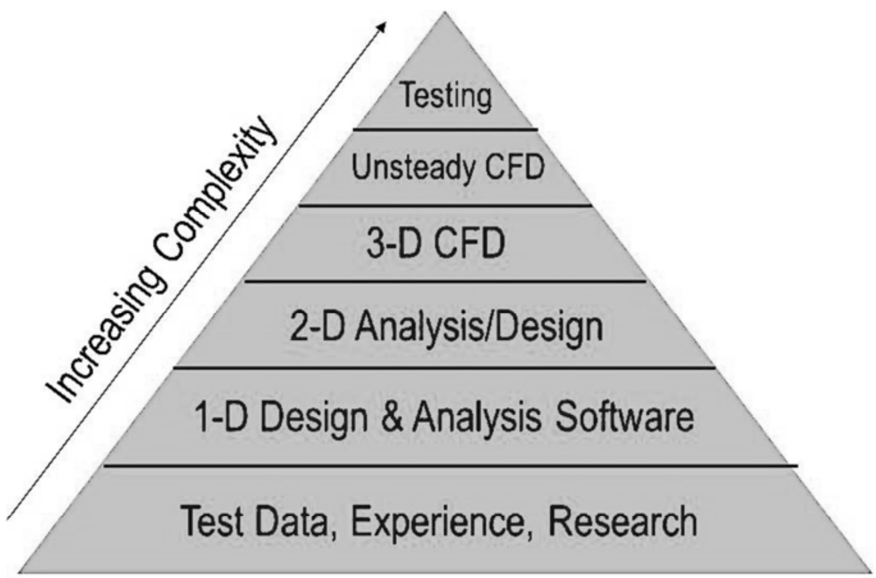

Figure 1. Design / Analysis Pyramid

Briefly describing the pyramid in Figure 1, all OEMs build their design processes on a sound foundation of good test data, extensive experience and focused research. The data and trends derived from this "data" foundation are used to develop the performance models that are built into the 1-D tools used to conduct scoping studies and develop the preliminary flow path geometry. These geometries are then expanded into the 2-D domain via the streamline curvature or blade-to-blade solvers that use approximation techniques to derive pseudo-3-D flow field solutions. Assuming a design meets the requirements of the first three layers of the pyramid, it is subjected to a 3-D CFD, and possibly an unsteady CFD assessment to obtain performance information that most closely simulates the "real world." Of course, if doubts still exist after the CFD work, the new components can be assembled into a test vehicle and subjected to stringent performance tests to obtain "as-tested" characteristics.

The lessons gleaned from each level of the pyramid are used to strengthen other levels of the pyramid to further enhance the assessment criteria used for those layers of the pyramid. For example, loss models developed from CFD studies can form the basis for enhanced 1-D code performance models for impellers, vaned diffusers, return channels, *etc.*

WHAT IS CFD?

The Basics of CFD
The following is a brief introduction to CFD and how it works.

We could be very brief indeed and simply say that computational fluid dynamics is the numerical solution of the Navier-Stokes equations. But such a brief statement omits many details that can be helpful in understanding the nature, process and results of CFD simulation. Some additional insight will be provided by presenting these topics:

1) The full equation set involved, including additional models and empiricisms required
2) The impact of turbulence on the equation set
3) The nature of the resulting equation set
4) The numerical solution process and why it is so
5) Assessment of validity, or conversely, expected sources of error

The topics will be addressed in the following, keeping in mind that the end application is simulation of turbomachinery for the purpose of better understanding its behavior and performance.

The equation set
As a starting point, let us consider the fundamental equations that describe fluid behavior. Fluid mechanics is a specialized branch of continuum mechanics where the focus is narrowed to materials that are fluids only. But there are many mathematical similarities to solid mechanics, which perhaps explains why the numerical methods used in CFD are similar to the finite element methods that are well-known for solid mechanics.

The main ingredients are Newton's Laws as applied to fluids, some basic thermodynamics, some equations that relate the various properties of the working fluid, and some relationships that describe considerations like heat transfer.

Newton's Laws are applied in a manner that asserts conservation of fundamental properties, such as mass, momentum, energy, species, *etc.* Thermodynamic considerations (first and second laws of thermodynamics) provide equations of state, both defining and linking fundamental fluid properties such as pressure, enthalpy, entropy, *etc.* Transport considerations relate stress to strain, heat transfer to temperature gradient (Fourier's law), mass transfer to concentration gradient, with associated properties such as viscosity, thermal and mass diffusivity, *etc.* Further complications develop by including phase change, multiple species and chemical reactions.

The set of equations that arise from the above process are called the Navier-Stokes equations [White, 1991], although they are a little more general than those first derived in the 19th century. Two additional comments are appropriate at this point:

1) The equations described above can be cast into a common form, which is referred to as transport equations. Each equation has four distinct terms, namely: transient (time dependence); convection; diffusion and source.
2) In general there are transport equations for mass, momentum, energy, species and phase. For simplicity we will restrict our attention to a single phase and species. As will be seen, turbomachinery applications present sufficient challenges, with variable physical properties, compressibility, stationary and rotating frames of reference and turbulence, *etc.* without adding more complexity.

The impact of turbulence
Strictly speaking the equations described above apply at any instant in time. The degree of difficulty would be acceptable if flows of interest were laminar in nature. In that case fluid mechanics would be much like structural analysis. Most flows of interest are

predominantly turbulent. Turbulence introduces unsteadiness into the flow field. The nature of turbulence is such that it introduces a wide range of length and time scales, and the higher the Reynolds Number the greater the range. So while the equations could be directly integrated for low Reynolds number flows, at more practical Reynolds Numbers this becomes practically impossible, even with the largest computers over very long times. That approach, known as direct numerical simulation (DNS) is only pursued in a well-equipped research environment for very special cases.

The most common and practical approach to dealing with the turbulence is known as "Reynolds averaging," consequently one often sees the term "RANS equations," which stands for Reynolds-Averaged Navier-Stokes equations [Durbin *et al*. (2003)]. In this process one decomposes the instantaneous quantities (pressure, velocity, *etc.*) into time-average and fluctuating components. The advantage of this approach is that the resulting equations look just like the original equations (or the laminar versions of the equations) but with a few additional, albeit unknown, terms. Additional equations are required to specify these unknown terms, but despite impressive mathematical efforts it is not possible to get away from the fact that there are more unknowns than equations, and at some point, empirical information must be added. This leads to a wide range of "turbulence models". The most common variants are:

1) Zero-order models, which provide an algebraic relationship between the turbulence and the mean flow.
2) Two-equation models, which provide two transport equations to characterize the turbulence; examples include the k-epsilon, k-omega and SST models [Menter, 1994].
3) Reynolds stress models, which provide transport equations for all of the Reynolds stresses (six) as well as equations for the scalar fluxes.

There are many additional variants, but the above provides the main landscape. All contain various levels of approximation and empiricism. In practice, the two-equation models provide the best balance between practicality and accuracy and consequently they and their variants are the workhorse models in use today.

Before leaving turbulence, a cursory mention of the "Large Eddy Simulation" (LES) approach [Guerts, 2004] is required. Again, there are many variants; with some that blend LES and RANS approaches [Menter, 2015]. This approach chooses to simulate the large scales, while modeling the small scales. LES is more computationally intensive (takes greater computer power) and is therefore more time-consuming to use. It is often only beneficial for certain classes of flows, particularly those that are very unsteady and are not dominated by boundary layers. A good example is the combusting flow in a gas turbine combustor. Use of LES methods is likely to grow as the models are refined and validated and computing power grows. Turbulence modeling remains one of the biggest areas of CFD research.

The nature of the resulting equation set

After applying the Reynolds-averaging process, one has a set of transport equations that assert conservation of mass, momentum, energy and turbulence quantities. Equations of state, fluid property relations and several empirical relations supplement this main set. The empirical relations generally deal with wall phenomena, such as shear stress and heat transfer, and are specific to the turbulence model selected.

It should be clear now that the equation set is very complex. The equations are lengthy, contain numerous variables, include partial derivatives, are coupled (linked to one another) and are non-linear (some terms in the equations are products of different variables and partial derivatives of variables), and quantities can vary in both space and time. Given all of these complexities, analytical solutions that account for all of the considerations discussed above have only been found for extremely simplified cases; *i.e.*, two-dimensional laminar flow between parallel plates. Since turbomachinery flow passages are far more complex than two parallel plates, general analytical solutions cannot be achieved. Therefore, one <u>must</u> resort to approximate numerical methods of solution.

The numerical solution process and why it is so

The equations described above apply at every point in space within the domain of interest. Early numerical methods satisfied themselves with solving the equations at only discrete points in the domain. That approach works well for some things, but it fails to enforce an important property known as conservation. Engineers are more satisfied seeing solutions where the amount of flow that enters a domain also leaves it. A similar principle of conservation applies for other quantities such as momentum and energy. For this reason, the "finite volume" approach is most popular for CFD.

In the finite volume method, the equations are integrated over a number of discrete regions within the domain, so the concept of a computational mesh or "grid" arises [Patankar,1980]. Generally speaking, the finer the grid is made, the better the solution will be. A finer grid is needed where gradients are large, for instance at a wall or in a shear layer, whereas the core flow typically requires less resolution.

In CFD, one can truly say that "the devil is in the details." Simply (numerically) integrating the equations is straightforward enough, but how does one represent the convective and diffusive fluxes that arise, and how does one linearize this inherently non-linear set of equations? And then given the resulting large, (temporarily) linear set of equations, how does one solve that set? There are many choices to be made. Each CFD developer has chosen answers to these questions after many years of development. Some choices are generally available to the user (choice of discretization scheme, for instance), while others are not. Most CFD codes provide sensible defaults for a specific range of flows.

The inherent non-linearity of the equation set means that the solution process is always iterative in nature. The type and "order" of the discretization scheme selected influence the accuracy and stability of the solution. Despite so much effort on the part of developers, there remains an "art" component to CFD modeling. But now simple problems are readily solved even by "CFD rookies." More challenging applications, for instance multi-stage blade row plus leakage path simulation, require experience.

Assessment of validity, or conversely, expected sources of error

CFD is a very useful tool, despite the inherent numerical complexities and empirical approximations identified above. But to get the most from CFD, it is helpful to understand the nature and implications of some of the factors that affect quality and timeliness of solutions. Organizations that rely on CFD usually develop their own "best practices." For turbomachinery, these practices are typically fairly common, although there is some variation across companies and machine types.

The first point is to realize that CFD can only solve what you give it, so if you only solve for a subset of the domain of interest then effects outside of the domain will not impact the solution. One must be sure that a representative domain is selected and that the boundary conditions applied to the domain are adequate, if not perfect. There is always a tradeoff; a small domain yields a smaller model and solves faster, but may neglect key influences. A larger domain with more features and additional geometry is more complex and will take longer to solve. For turbomachinery, there is now plenty of experience to guide the user and it is well known what can be expected from analyzing a single passage of an impeller versus a full stage analysis, for instance.

One of the most important factors impacting solution quality is the mesh. Generally the finer the mesh the better the solution, because as the mesh is refined the numerical discretization error is reduced. A best practice is to solve on three grids: coarse, medium and fine, and then look at the changes of key quantities and flow features as the grid is refined. This may seem like a lot of work but experience has shown that it saves time and money in the long run. It is seldom that the grids used are fine enough to reach a "grid independent" solution (the point where further refinement of the grid does not appreciably change the solution), and one must be content to accept some level of error. This grid refinement process helps us understand the proximity of our solutions to "grid independence."

The other important mesh-related factor is its distribution. The distribution is never uniform because it is important to resolve flow gradients: boundary layers; free shear layers; separation regions; shock waves; *etc.* In turbomachinery, these phenomena can change somewhat as the flow rate changes, so the mesh must be constructed with the expected range of flow rate in mind. Good practice (but often neglected) is to use the powerful post-processing now available. This enables one to look at the computed flow overlaid on the mesh and assess the success in adequately resolving the key flow features.

Several user-controllable settings also affect solution accuracy. The user must be sure that the settings chosen are appropriate to the problem at hand. One key setting relates to the "discretization scheme." CFD solvers typically offer "first order," "second order" or "blended" methods. This concept of "order" relates to the numerical approximations made when converting the partial differential equations into non-linear algebraic equations. "First order" methods provide computational stability (*i.e.*, CFD runs are more likely to achieve convergence), while "second order" methods provide better accuracy, but risk instability. Blended schemes attempt to merge the two on a local basis. Best practice is to use a blended method that is mostly, if not fully, second order. Use of purely first order methods ("upwind differencing," for instance), while very stable, is not recommended because they introduce non-physical diffusion and loss of total pressure, unless the mesh is very fine.

The final point that will be considered relates to the turbulence model. In most cases, the best practice is to use the best available two-equation RANS turbulence model, perhaps with some additions for transition, *etc.* The shear stress transport (SST) model is generally recommended. It is not advisable to adjust the (adjustable) constants, unless one is a turbulence modeling expert and has laboratory-quality data to justify doing so. There are situations where other models are more appropriate (LES-type models for a combustor, full Reynolds Stress models for a cyclone, for instance) and yield better results. However, for most turbomachinery applications the greatest body of experience is with two-equation models, and they represent the best compromise in terms of stability, accuracy and time-to-solution.

Other factors can influence accuracy and time-to-solution. The above is only an introduction to some of the main points. CFD users are recommended to consult experts and develop best-practices for each class of component or machinery of interest. Doing so will enable one to make good use of CFD. On the other hand, failure to do so can result in dissatisfaction, frustration and unreliable results.

WHAT REALISTIC BENEFITS ARE AVAILABLE FROM CFD?

Having established what CFD is, the next question is, “What good is it?” CFD has some unique advantages:

- It shows more details of the flow field than would be practical with a test. Only so many instruments can be placed in a machine. This is especially true when analyzing a client unit instead of a test rig. Seeing these details can show *why* performance is what it is, and so drive design changes or troubleshooting fixes based on data.
- It enables predicting performance on machines that have not yet been tested, or perhaps even built. It is true that CFD is sometimes over-hyped as a “virtual test stand,” and its predictions should be taken with several grains of salt. Nonetheless, properly done CFD results are of useful accuracy, even if not perfect accuracy. This is especially important for unique components for which there may not be any other analysis program available.
- It enables evaluating designs quickly. CFD is slow compared to other analytical methods, but is fast compared to building and testing parts.

CFD, however, also has limitations, and cannot necessarily answer every question a customer might have. Even if an analysis is possible, it might not be practical because of the time, cost and computer resources that would be required. CFD has always been limited by computer capacity, making it necessary to simplify analyses or abandon particularly demanding analyses altogether. This is still true today, even with the advances in computer technology over the years. In most analyses, more mesh or a larger model extent would be desirable, and of course, greater speed is always desired on any project. In cases that require many runs, just waiting for the computers to finish can be the largest part of the overall schedule.

However, as the capability of computer hardware continues to improve, what is “practical” keeps changing. Problems that were considered impossible gradually become routine. In 1991, it was a major milestone at one engineering firm when the node count exceeded 10,000 in a model of a single stationary passage. Now analysts routinely run models of 40,000,000 nodes_ 4,000 times larger. Figure 2 shows the relative speeds of a single core in many of the computers that were used for CFD at one OEM over the years. The speed has increased by a factor of 41 over about 18 years.

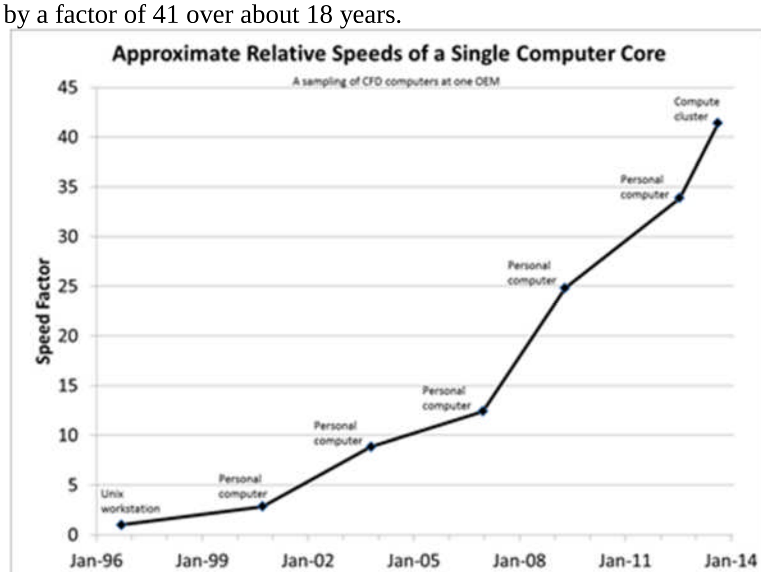

Figure 2. Speed of a single computer core through history

The graph in Figure 2 is for a single compute core on each machine. Not shown is the enormous impact that parallelization has had on speed. Back in 1996, analyses were typically run on a single CPU that had only a single core. Now problems are routinely divided up among many cores running in parallel--128 cores is a fairly typical number, though fewer are often used for small problems and many more can be applied for larger models. For these larger simulations, the speed-up resulting from parallel computing is nearly linear [HP Enterprise, Slagter *et al.* (2014)], so it would be reasonable to estimate that the overall speed of the analysis on those 41-times-faster cores could be more than 100 times faster yet; more than 4,000 times faster than the 1996 workstation.

Increased computer performance is usually used for the following improvements:

- **Increased mesh density**. This improves the accuracy of a solution. It is particularly important near the walls, where using lots of very thin cells enables the use of more sophisticated turbulence models, so that boundary layer separation and reattachment can be better predicted. Very thin wall cells, combined with a reasonable expansion ratio out into the main channel, can drive the node count up significantly, perhaps on the order of a factor of 10 compared to a simpler model.
- **Increased model extent**. All the components in a machine affect each other to some extent. Component performance prediction is often improved considerably by including up- and downstream components in the model, just to make sure the flow field in the component of interest is more correct. Also, if circumferential variation is a concern, increased computer power can make it practical to model the full 360-degree geometry instead of a single representative passage. Adding more components can easily double or triple the size of a model. Changing from a single passage to a full circle increases the model size by a factor of the blade count, which can be a serious problem for axial machines.
- **Transient analyses**. Most industrial CFD are still run as steady-state solutions, which assume a flow field that does not change over time. Yet the flow in a real turbine or compressor is unsteady by nature, at the very least changing as rotor blades pass stator vanes. Steady-state approximations work surprisingly well in many cases, but more difficult flow fields may require transient analyses for reasonable accuracy. Transient analyses have to find a new solution for each small step forward in time, generally needing a few internal iterations on each time step. Changing from steady-state to transient analysis can increase the time required by a factor on the order of 10, though the actual value varies greatly.
- **New types of analyses**. Some analyses that were impractical or impossible in prior years are now available because of the increased computer performance available. Among these are:
    - Extracting time-dependent data from transient analyses to generate Fast Fourier Transforms (FFTs) of any quantity, for vibration analyses of mechanical parts.
    - Fluid/structure interaction, finding time-varying mechanical loads on parts caused by the fluid flow, or even the interaction between the fluid flow and the mechanical response [Ansys documentation 1].
    - Acoustic analyses, finding sound pressure levels caused by the fluid flow [Ansys documentation 2].
- **Improved productivity**. Of course, faster computers can allow analysts to simply get a job done sooner. But most of the increased computer performance is being used to enhance accuracy, by increased node count and model extent, so the actual time required to run an analysis has not decreased as much as one might expect. It was noted above that a typical model might have 4,000 times more nodes than a model from 1996, and that the overall analysis speed might be around 4,000 times faster. The time required for a CFD run is roughly proportional to the number of nodes, so the overall time required to run a model might not have changed all that much. The biggest difference is that the model itself is far superior to that from 18 years ago, so that more useful geometry can be analyzed with greater accuracy.

Even today, with these fast computers running in parallel, computer limitations usually require compromises to be made in CFD analyses. Industrial CFD models seldom have “grid convergence,” where the mesh is so fine that further refinement would not change the answers obtained. Model extents are typically limited to keep the model size reasonable; modeling an entire machine is almost always impractical. Most analyses are still run as steady-state, to keep the required time down. One just has to hope that the curve in Figure 2 continues to climb.

For a small, straightforward, standalone CFD analysis (not integrated into a design process), with a 3D CAD model available, it is likely to take a day or two for geometry manipulation and mesh generation, a day or two of elapsed time for setting up and running the model, and a day or two for postprocessing (calculating results and plotting them in a useable form). So the minimum time for a simple CFD analysis is likely to be about a week. There is no upper limit; multiple transient runs of large models, especially if performed on computers that are not really up to the task, can easily consume months.

The CFD analyst should work with those requesting the analysis before a project begins to determine what compromises are reasonable to keep a CFD analysis practical. These choices depend on why the CFD analysis is being done in the first place; that is, what is to be learned from the analysis and how much time and money can be committed to the effort. In some cases, especially for a field problem that has brought an expensive production line to a halt, thorough CFD analyses of all the relevant cases may simply take too long. Simplified CFD may be required, after making sure everyone understands the limitations or risks associated with simplifying the approach, or CFD might not be the right tool for the job. In all cases, the focus must be on exactly what questions the CFD analysis is expected to answer, and what kinds of analyses would be required as a minimum to answer the questions.

Table 1 lists some of the typical questions that CFD might be needed to answer, with some description of the kinds of analyses required to address the question and the difficulty (and thus cost) of each. Rows are color-coded from green for easy jobs, through yellow, orange, and finally, red, for very difficult jobs. This table is not at all exhaustive. Many different questions about a machine's performance can be addressed by properly applied CFD. The objectives for the study must be discussed with the CFD analyst. The work must not proceed until all involved are convinced that the proposed analysis will meet the stated objectives.

Table 1. Examples of CFD Analyses

Information Needed	**Type of Analysis**	**Difficulty and Cost**
"No surprises" overview of bladed component flow field.	Steady-state single-passage model.	Very easy and inexpensive.
Approximate bladed component or stage performance to compare with a similar case.	Steady-state single-passage model, new and baseline cases.	Easy and inexpensive. Be careful with boundary conditions for single components.
Absolute stage performance, like a virtual test.	Transient, fine mesh, possibly 360-degree model.	Difficult and expensive. Unlikely to produce the desired results anyway. Best left for the OEMs to keep striving for, using their test data.
Approximate stage flow range.	Steady-state single-passage model. Might need transient at low flow. Several runs required.	Moderate difficulty and expense.
Sideload or extraction behavior, spanwise effects only.	Steady-state single-passage model.	Easy to moderate, depending on the detail required.
Sideload or extraction behavior, including circumferential effects.	Steady-state 360-degree model.	Moderate difficulty and expense.
Performance effect of upstream bends or components.	Steady-state 360-degree model.	Moderate difficulty and expense.
Approximate bladed stage performance for multiple stages.	Steady-state single-passage model.	Easy to moderate, depending on the detail required and how many stages.
Performance of entire machine, flange to flange.	Steady-state mix of 360-degree and single-passage models.	Usually not practical. Could be just difficult and expensive for a small, simple machine.
Qualitative idea of flow stability.	Steady-state single-passage model.	Easy to moderate. Look for regions of reverse flow or numerical instability. Might not give a definitive answer.
Specific frequencies and magnitudes of flow instability.	Transient, single-passage or 360-degree model depending on the specific case. Should also run a baseline case with known instability to validate the CFD approach.	Difficult and expensive. Risky; might not produce the desired results.
FFT of flow field from a stationary component.	Steady-state 360-degree model.	Moderate difficulty and expense. FFT is taken from values sweeping around the circle as if it were on a downstream rotating component.
FFT of force on a rotating component.	Transient, multiple-passage or 360-degree model depending on the specific case. There are rules for setting the circumferential extent, depending on the type of analysis.	Difficult and expensive.
Sound level and frequency prediction.	Transient, usually linked with other software for acoustics.	Very difficult and expensive. Can vary widely depending on specifically what is needed.

Information Needed	Type of Analysis	Difficulty and Cost
Fluid/structure interaction, average mechanical loads.	Steady-state, single-passage or 360-degree model depending on the specific case. One-way link to FEA solution, either simultaneous or passing the fluid pressures as boundary conditions after the CFD.	Difficult and expensive. Can vary widely depending on specifically what is needed.
Fluid/structure interaction, dynamic mechanical response.	Transient, single-passage or 360-degree model depending on the specific case. One-way link or full 2-way coupling with simultaneous FEA solution, depending on flexibility of solid parts.	Very difficult and expensive, not practical in many cases. Can vary widely depending on specifically what is needed.

A common desire is to use a single CFD run to predict absolute performance of a machine, like a fast, low-cost alternative to a test. Unless a CFD approach has been carefully validated against multiple tests, CFD generally will not predict absolute performance very well. Efficiency predictions are often off by a few points, usually being too high. But results vary a lot, of course, if the discrepancies were consistent and repeatable, then CFD would be predicting absolute performance very well.

Validating a CFD approach is different from calibrating other kinds of performance analysis programs. Traditional turbomachinery design programs have loss parameters and corrections that can be adjusted based on experience. CFD, by contrast, performs a more pure first-principles analysis. Performance numbers simply come out of the calculation of the overall flow field; there are no virtual “knobs” to adjust for performance. Validation consists of modifying the CFD approach—such things as mesh density and topology, turbulence model, inclusion of geometric details, convergence criteria, and others—until CFD is able to match test data and, more important, successfully predict the outcome of the next test. Even after all this effort, CFD should not be expected to match test data exactly.

There are numerous reasons that CFD will not match test results, many of which were discussed in the section “WHAT IS CFD?”. Some are limitations of CFD itself; others come from the practical limitations discussed above. These reasons include:

- CFD, like any analytical tool, remains an approximation of the real world. Mathematical models of fluid flow are not perfect. This is especially true for turbulence models, but applies to other aspects of the flow as well, such as surface roughness effects, heat transfer convection coefficients, gas properties, transitions between laminar and turbulent flow, and others.
- CFD has to approximate continuous flow with discrete pieces, mesh cells to represent three-dimensional space and finite time steps to represent time in transient runs. Practical limits on computer resources exacerbate this issue, as mesh cell sizes and time step increments are seldom as small as would be optimum.
- Many CFD analyses are run steady-state, even though a turbomachinery flow field is fundamentally transient.
- CFD analyzes CFD models, not actual parts. Models are usually idealized in ways that the analyst hopes will not affect the results significantly. Dimensions are usually exactly nominal, with no tolerances. Models are usually based on cold, stationary geometry, without considering the changes in shape from heat and rotation. When temperature and mechanical strain are included, they must be calculated by other approximate methods, not necessarily matching the actual part in service. When surface roughness is included at all, it is approximated as a uniform sand grain height rather than mimicking the complex roughness of a real part. Welds are usually either neglected or represented by perfect chamfers. Small geometric features are often left out to simplify meshing and reduce the number of nodes in the model. When troubleshooting units in the field, fouling and erosion are usually not considered, and have to be approximated if they are.
- CFD results can be strongly influenced by the boundary conditions applied to the model. The only places that inlet and outlet boundary conditions are really well known are far up- and downstream in the pipes, but this model extent would be much larger than practical in almost all cases. Compressor surge is influenced by the entire downstream piping system, which of course is not included in a practical model. Conditions at the model inlet must be specified, and are often taken as uniform. Inlet profiles can be specified if they are known, but these profiles are also approximations to the real flow. Inlet turbulence parameters must be specified as well, and these are seldom known. Outlet conditions, typically mass flow or average static pressure must be set, and these may not be known exactly. Heat transfer through walls is usually neglected and is difficult to set correctly if considered.

Sensitivity studies can be used to evaluate the effects of these approximations on CFD results. Mesh sensitivity studies are often discussed in technical papers, but the same concept also applies to other parameters. If there is concern about whether a geometric feature has a significant effect, run models with and without the feature. If there is concern that the boundary conditions are affecting the results too much, run cases with variations in the boundary conditions, perhaps including or eliminating a profile, or changing the values of inlet turbulence parameters. The same approach can be used for any feature of a model.

The most common kind of sensitivity study is a study of mesh density, as mentioned in the section "WHAT IS CFD?" The idea is to run models with increasing mesh density until an increase in mesh density does not change the results significantly. The smallest model that gave the same results as a larger one can be used for further analyses. As a practical matter, this "grid convergence" is seldom achieved in industrial turbomachinery models. Models that include all required features yet are small enough to run on available computers usually don't have enough nodes to be free of mesh influence.

It is important to make sure the changes between models are large enough that similar results are not simply caused by running similar models. For example, consider doubling the mesh density. Doubling the overall node count results from increasing mesh density in each coordinate direction by only about 26 percent, a rather small change. Doubling mesh density in all three coordinate directions increases the node count by about a factor of eight, a more reasonable target for a refined mesh.

Other sensitivity studies show how much performance changes as a parameter is varied. Often, input parameters that affect the results are not known precisely. Common examples include inlet turbulence, surface roughness, and inlet flow angle. Models are run with the uncertain parameter varied between reasonable limits. The resulting variation in performance can be considered a tolerance on the results. This kind of study can show that the influence of a parameter is small, increasing confidence in the results. This is often true of parameters in the inlet boundary condition if the model inlet is far upstream of the region of interest. If the variation in performance is large, it may be worth spending additional effort to narrow the tolerance band by trying to determine the most representative value of the parameter.

Sensitivity studies of multiple parameters can add a considerable number of models and runs to a task. Judgment is required to choose the tradeoff between concern about the effect of a parameter and time required to make the additional runs. This judgment improves over time as experience with similar analyses is gained. It is, however, not always necessary to spend so much effort trying to match CFD results to real-world performance. The best way to mitigate the issues with CFD performance calculations is to find relative performance rather than absolute performance. Many CFD analyses involve changes from a baseline, whether a new design iteration, a proposed modification, or a change in hardware from its intended shape. The best approach is to analyze both the baseline and the changed hardware and compare the results. The mesh, geometry simplifications, boundary conditions, and other parameters must be as similar as possible between the two models. The idea is to make the deviations from absolute performance caused by each of these features as similar as possible for the two models, so the difference in performance is caused by the change in geometry.

If the models are similar, then the trend of relative performance between them should be correct, and the magnitude of the performance difference is likely to be reasonably close to the real value. This kind of analysis can show whether the change was an improvement and approximately how much of an improvement. If the real performance of the baseline geometry is known, adding the change in performance from the CFD analyses can allow one to estimate the real performance of the changed geometry.

Despite the extra time and effort required to analyze two geometries instead of just one, it is almost always worthwhile to run a baseline case in addition to the case of interest, so performance can be evaluated relative to a baseline.

CFD cannot replace testing, but it can be used effectively in conjunction with testing. The single most important contribution of CFD to a test program is to eliminate some tests. Using relative performance comparisons, designs can be evaluated, and only those with promise need to be tested.

Modeling the test itself can also be useful for the following reasons:

- CFD can sample the flow field where the instruments are located and compare that with true massflow-averaged values across the passages. This helps determine whether the instrument locations are adequate to calculate performance. This can be particularly important for client tests and field tests, where fewer instruments are available than in development testing. Comparing the values at the instrument locations and across the passage may enable one to estimate the actual full-passage value in the test.
- CFD can help identify problems with the test. If there is an unexpected discrepancy between CFD results and test, both should be investigated. This is particularly true if the match between CFD and test is especially far off at some particular instrument, which might indicate a problem with that instrument.

- CFD can fill in information where instruments are not available. If the match with instruments is reasonable, then it's likely the rest of the CFD flow field is close to what actually exists in the real machine. Values in regions where no measurements were taken can be estimated from the CFD results, and individual component performance, flow angles, losses, and other parameters can be calculated. This can be particularly important for client tests and field tests, where fewer instruments are available than in development testing.

As long as expectations are reasonable, CFD can provide vast amounts of useful data for designing and troubleshooting turbomachinery.

DELIVERABLES FROM CFD

The results of CFD can be presented in many different formats. As noted above, CFD provides pressures, temperatures, velocities, flow angles, gas properties, *etc.* throughout the computational domain so it is possible to calculate performance parameters from CFD in much the same fashion that one can calculate performance results from an actual compressor or from any other analytical tool. Of course, the most frequently used are the ever-popular quasi-3-D images of the velocity vectors or streak lines that begin to resemble abstract works of art. However, it is also possible to generate the simplified tabulations or plots that resemble the typical charts and curves OEMs or third parties typically provide during a sales proposal or the like.

From a very high level, results from CFD simulations tend to fall into the general categories listed below:
1. Scalar plots of key operating or aero assessment parameters including overall performance maps
2. Contour plots of key operating or aero assessment parameters
3. Velocity vector plots showing the flow within a component or components
4. Streak line or stream tube plots showing the flow within a component or components
5. Animations showing the flow through components or how aero parameters vary in time (for unsteady or steady simulations)

Some examples of each will now be offered with the obvious exception of the animations. Those are "difficult" to provide in a static printed paper. Further, the examples provided are but a small fraction of the different plots that can be generated. Please note that further examples of these various styles of plots will be offered in the Sample Cases section at the end of this document; but these are offered now to save time in describing the types of plots later.

Scalar Plots
Scalar plots can be used to display general trends in almost any conceivable aerodynamic parameter in the overall computational domain or within a small region of that domain. For example, the analyst can provide an overall performance map for a single compressor stage (inlet guide, impeller, diffuser and return channel system) or the flange-to-flange map for an entire multi-stage compressor. See Figure 3 for an example of a single-stage performance map generated by CFD. Depending on the purpose for presenting this figure, the presenter might include test data or the performance trends obtained from a competing or previous stage design. [See Drosjack *et al*, 2011]

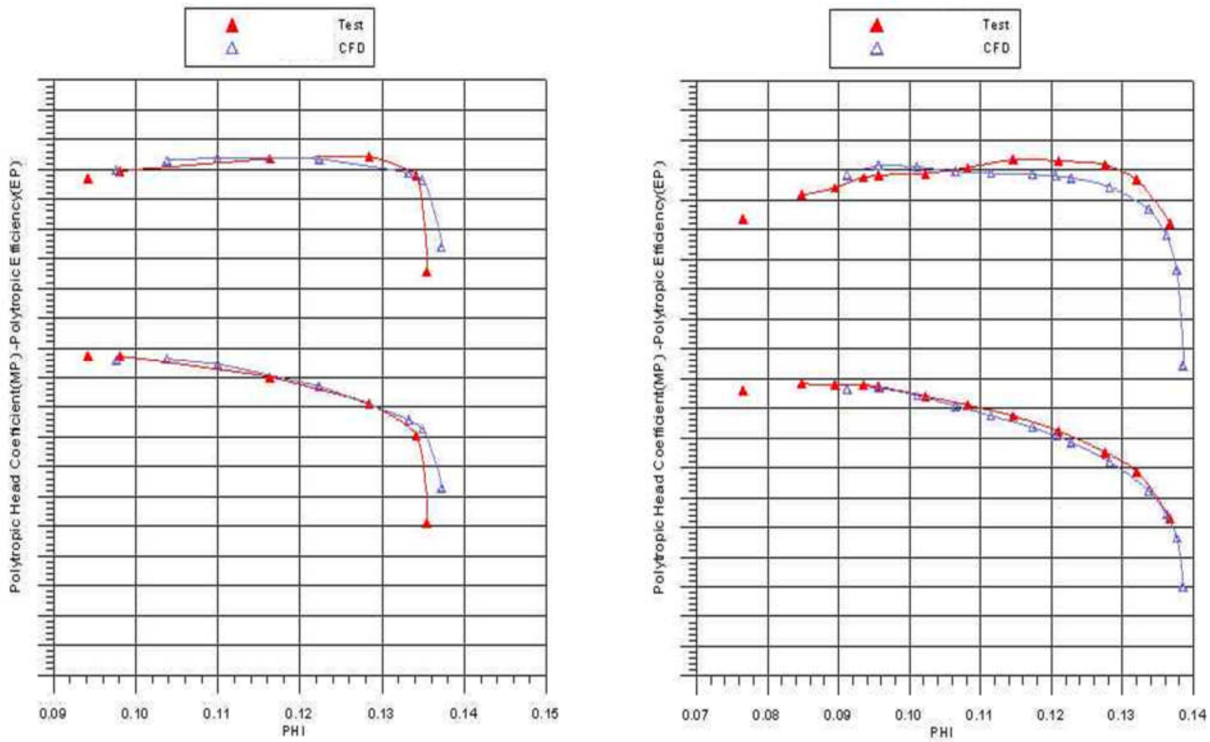

Figure 3. Examples of scalar plots - performance maps generated via CFD compared to test data.

As the discussion digs deeper into the flow field characteristics of the individual components, the analyst might offer plots showing the variation of flow angles, velocities or other operating parameters (pressures, temperatures, enthalpies, entropies, flow distortion terms, *etc.*) exiting or entering components or regions within the modeled flow path. Again, the trends might be plotted against measured data or against results obtained from a previous or competing design. It is important to remember that the CFD results are typically three-dimensional (though some analyses are done in two-dimensional space). Therefore, to obtain scalar properties, the three dimensional or planar values must be averaged. This is typically done via area (averaging a parameter based on the amount of area over which said parameter is of fixed value) or mass averaging (averaging the parameter based on the amount of mass flow that has a specific value for the parameter). It is important for the reviewer to know which averaging method was used. Two examples of scalar plots are given in Figure 4. The caption describes the parameters displayed.

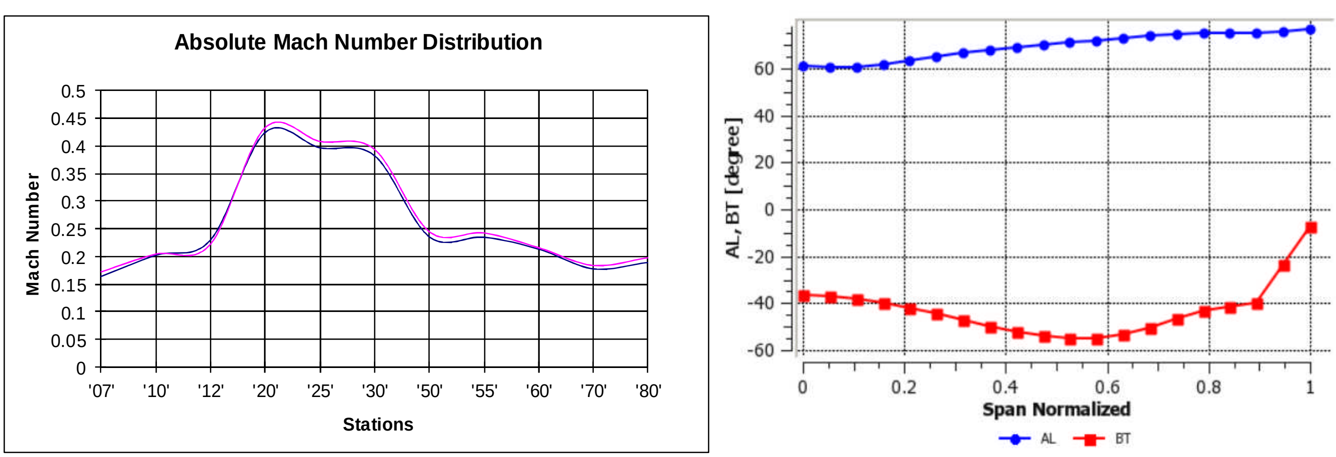

Figure 4. Examples of scalar plots – On left, absolute Mach number distribution through compressor stage; on right, absolute (blue) and relative (red) flow angle at the impeller exit.

Sidebar: It is likely already apparent that it is possible to get overwhelmed, buried or lost in the maze of plots, tables and other "pretty colored pictures" that can be generated from CFD results. In fact, it is possible that this is, in fact, the presenter's strategy… to provide so much detail that the viewer cannot see the forest for the trees. This has led to the alternate definition of CFD as "Clearly For Deception" or "Colors For Dummies." To avoid such pitfalls and to save untold amounts of time sitting in design audits, it is recommended that detailed agreements be hammered out in advance as to what CFD specific results will be discussed and in what format. If there is no clear value in reviewing certain parameters or trends, or if such figures do not contribute to the "story line" being told, leave such figures out of the presentation.

Contour Plots

Contour plots are used to display parameters on planes or surfaces within the flow domain. Again, it is important that the analyst or presenter be very clear about what is being shown because it is very easy to lose one's perspective on these quasi-three-dimensional surfaces. Some of the more common surfaces that are shown include:

- Near blade / vane pressure surface
- Near blade / vane suction surface
- Near shroud surface
- Near hub surface
- Mid-span (hub-to-shroud)
- Mid-pitch (blade-to-blade or vane-to-vane)
- Meridional or radial planar cuts through the flow path (from the inlet to the exit of a component or stage).

It is helpful if the physical geometry or "metal" is included in the figure to help visualize the planes or surfaces being discussed; however, there are times when the geometry obstructs the view of the contour plot details. It is important that those in attendance ask questions if unclear about the details of the data being shown.

There are numerous types of contour plots shown in Figure 5. Again, the caption provides details on the data being portrayed. While most are for typical centrifugal compressor stages, item D shows the results of a mixing calculation for an incoming sidestream.

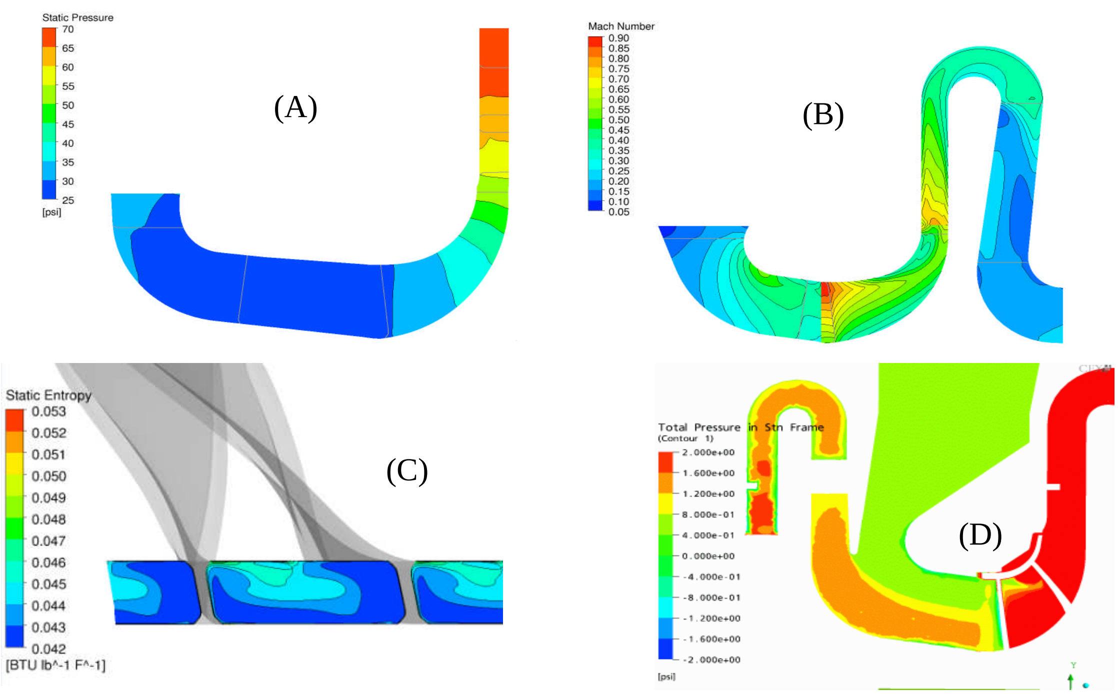

Figure 5. Examples of contour plots: (A) – Static pressure distribution through inlet guide and impeller; (B) Mach number distribution through a centrifugal stage; (C) Entropy distribution at impeller exit plane; (D) Pressure distribution in mixing section of sidestream.

Velocity Vector Plots

Velocity vector plots are essentially the same as contour plots except that they show vectors that represent absolute or relative gas velocities instead of contours of some parameter. The vectors in the plots are typically colored by some quantity such as the velocities themselves, absolute or relative Mach number, flow angle, or almost any other parameter. Vectors colored by velocity or Mach number are certainly the most common.

Vector plots are very popular because most rotating equipment engineers have at least a fundamental grasp on how the flow is supposed to travel through a compressor (or turbine) component or from component to component. Therefore, even someone not knowledgeable in aerodynamics can discern flow patterns in vector plots that suggest problems. That is, a close review of vector plots can quickly uncover regions of flow recirculation or other flow abnormalities that could contribute to more global problems. They can also be useful in understanding the uniformity, or lack thereof, in the flow field exiting the various components within a stage or machine and in determining the root cause of any performance shortfalls (whether in an existing turbomachine or in one being designed).

As with the previous styles, several velocity vector plots are provided in Figure 6 and the caption provides the details of the parameters being plotted. Some cases include the physical part geometry and others do not to allow the reader to see how the presence or absence of the geometry impacts one's perception of the plots.

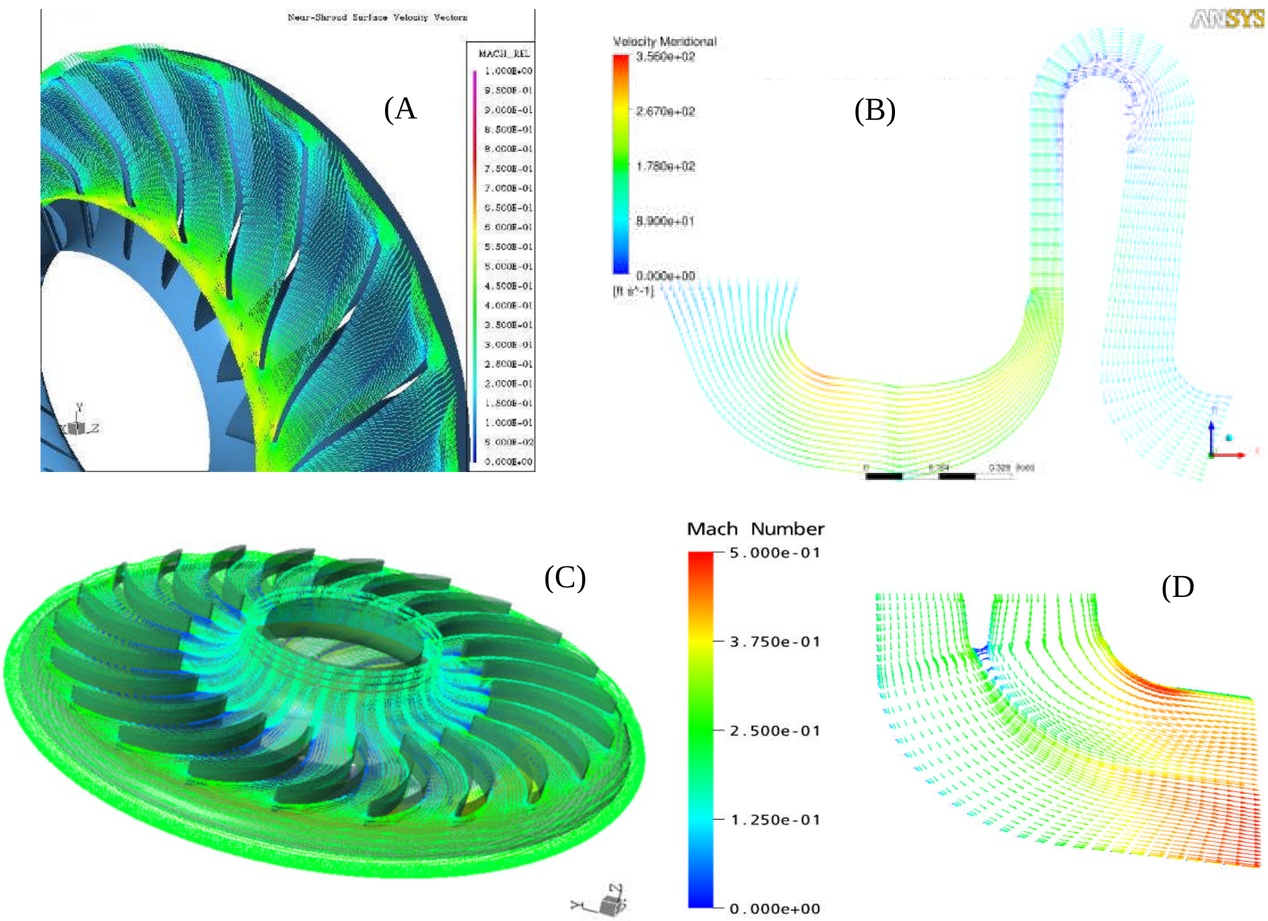

Figure 6. Examples of velocity vector plots: (A) Impeller velocity along shroud surface; (B) Meridionally-averaged velocity through compressor stage; (C) Velocity vectors in return channel passage; (D) Velocity vectors in sidestream mixing section.

Streak Line or Streamtube Plots

Like velocity vector plots, streak line or streamtube plots are popular because they help visualize the flow patterns through a component or combination of components. These plots reflect the path a "mass-less," "volume-less" particle would take if released at certain points in the computational domain. They are much akin to the flow visualization done with smoke on test vehicles with visual access to the flow passages (see Figure 7 for examples). Of course, the effort and expense associated with generating these streak line plots are much less than the cost and time it would take to design, build and do flow visualization testing on a scale model test vehicle; however, the accuracy of the flow traces is only as accurate as the velocities and flow angles calculated by the CFD code. Still, nowhere is the ability to look inside a flow component more obvious than in looking at the flow traces generated via "mass-less" particle tracking.

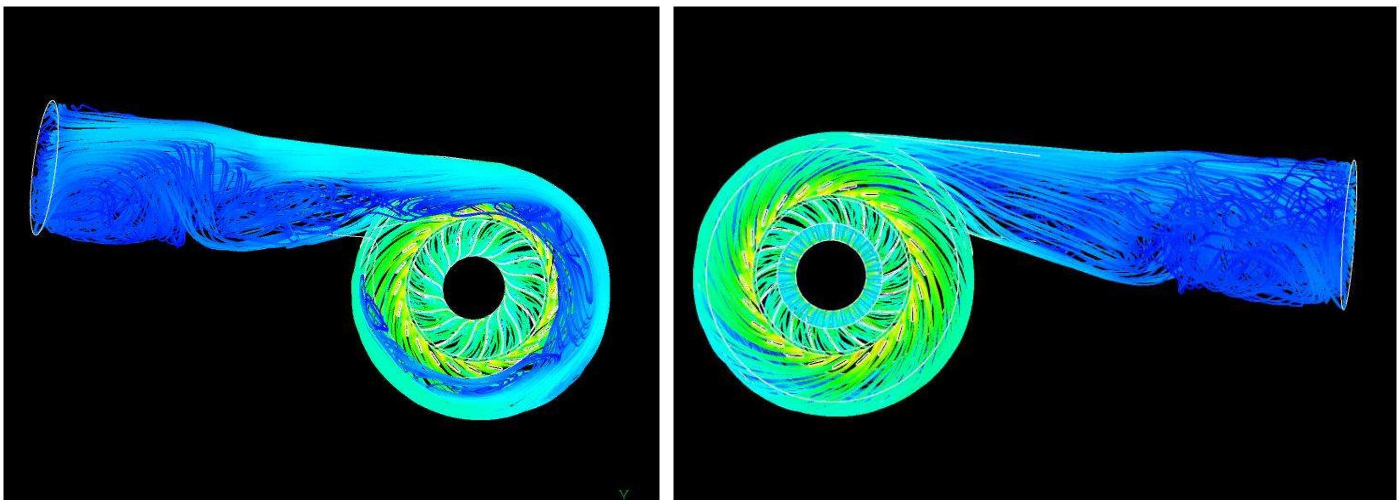

Figure 7. Streak line plots for an impeller with downstream volute. Note distorted flow field in volute and exit nozzle.

Of course, if one is not careful, these plots can become so overwhelmed with streak lines that the individual particle tracks are untraceable. It is incumbent on the analyst to determine the correct particle density and release point to bring out the details in the flow field. For example, releasing the particle in the wrong locations might allow the traces to totally avoid the recirculation zones or vortices that would indicate a problem in the component design or in the compressor operation. However, those with knowledge of these shortcomings in flow traces should note voids in the particle traces and question the lack of flow in those areas.

Animations

Animations are essentially movies or "motion pictures" that are assembled from CFD results. These can be generated using steady state solutions but they are more commonly used to visualize the results of unsteady or transient simulations. Like the other plots described above, the primary objective of such visuals is to help the analyst or the reviewer understand or assess the flow behavior in the components. However, the animations introduce the time component; *i.e.*, provide insight into how the flow field changes in time. This might be observed in the fluctuation of pressure contours with time or the variation in velocity vector magnitude and direction with time, *etc.* Essentially, the "movies" are created by taking a series of snapshots of the computational domain and then assembling these "snapshots." This requires large amounts of data storage during the execution of the unsteady CFD simulation. Gigabytes or terabytes of storage are required, so these unsteady simulations and, today, animations are typically not done as part of the day-to-day design or analysis process.

Obviously, it is not possible to include animations in this paper; however, a series of screen captures is provided in Figure 8 to give the reader a feel for the type of results that appear in the animation.

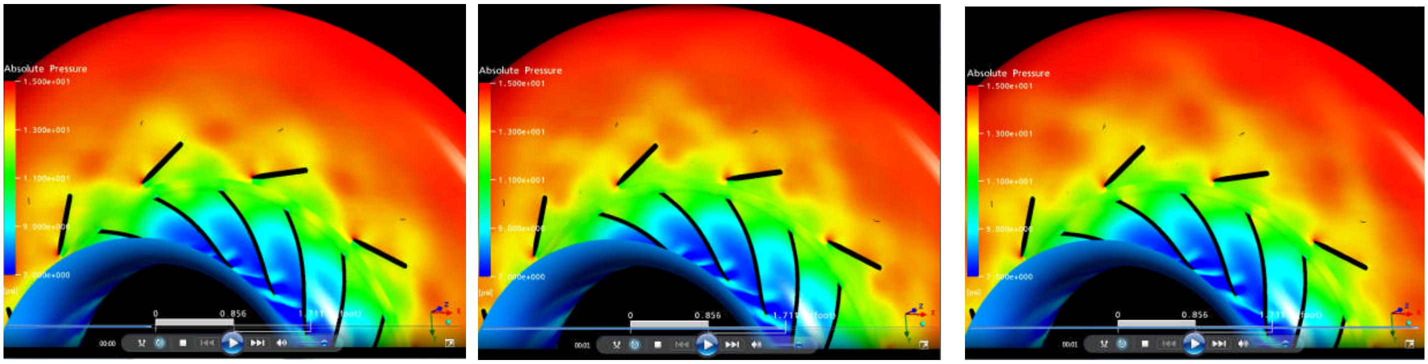

Figure 8. Screen captures of animation showing static pressure profiles as impeller blades rotate past low solidity diffuser vanes.

SAMPLE CASES

Inlet Loss Coefficient

Part of a radial compressor inlet CFD model is shown in Figure 9. Flow enters vertically from a straight pipe, is distributed among the curved vanes, and flows out in an annulus toward the reader, to enter the first impeller. The small white balls near the trailing edges of two vanes show proposed locations for instrumentation.

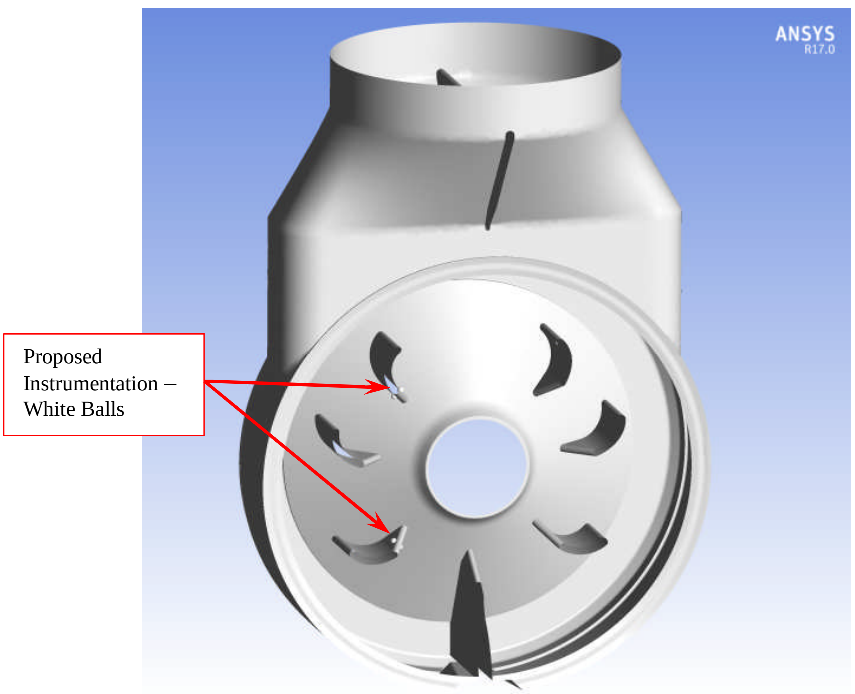

Figure 9. Compressor Inlet CFD Model.

Those looking at CFD figures must get used to some of the things show in Figure 9, such as:

- Some surfaces are removed (to help a reader see inside the part) resulting in a view of the assembly that could never exist in real life. The shaft is not present, and the shroud surface formed by the first diaphragm has been removed.
- Only surfaces that contact the flowing gas are shown. For most CFD models, where mechanical and heat transfer effects are not included, these are the only surfaces that exist in the model. The CFD analysis does not know or care about the actual solid pieces that these surfaces are part of. Consequently, one can see right through the vanes, and the overall geometry seems misshapen because the outer surfaces of the casing and nozzle are missing.

- Fictitious geometry for illustration is blended freely with actual part geometry, such as the small white balls for instrumentation locations. Often, surfaces in the flow path where average values were calculated will be shown, as seen in Figures 10 and 11.

Suppose this inlet was going to be tested and the loss coefficient was to be measured. Though seemingly simple, it is extremely difficult to do. CFD can help explain why and avoid measurement problems.

Total pressure is often used to evaluate losses in stationary components. Total pressure is the pressure that would result from bringing a fluid to a stop isentropically, that is, with no additional losses. It's the sum of the static pressure (the actual pressure of the fluid) and the "dynamic pressure," the effective pressure caused by the movement of the fluid. For an incompressible fluid, the dynamic pressure is half of the density times the velocity squared ($\rho V^2/2$), and that's not a bad approximation for gases moving at low velocity. CFD postprocessing, however, does not use this approximation; it calculates the correct value of total pressure. In a stationary component with no losses, the total pressure would remain constant, no matter how the static pressure changes. In a real stationary component, the total pressure always falls. Naturally, a design goal is to keep the total pressure from falling very much.

In a test, total pressure can be measured with probes that protrude into the flow and capture some fluid, bringing it to a stop with reasonably low loss. The probes must be aligned with the flow direction within some tolerance. Another common approach is to measure the static pressure with static taps, calculate the density from measured pressure and temperature, calculate velocity from passage area, density, and measured mass flow, and finally calculate the total pressure from all of these. The process can be summarized as follows:

- Static pressure P from measurement
- Static temperature T from measurement
- Density ρ from gas properties ($\rho = P/(RT)$ for a perfect gas, where R is the gas constant for the particular gas)
- Velocity $V = w/(\rho A)$, where w is mass flow and A is flow area
- Total pressure $P_t = P + (\rho V^2/2)$

Instrumenting the pipe near the nozzle inlet flange is not difficult, and because of the uniform flow coming from a straight inlet pipe, the measurements should give good approximations of the mass-averaged values across the entire pipe. CFD confirms this. Using static pressure and density values along the outer edge where pressure taps would be, and finding total pressure from these and the mass flow and the area yields a total pressure value of 49.993 psi, whereas, a mass-average of total pressure across the entire pipe gives 49.994 psi… a difference of only 0.001 psi. Therefore, one can assume the inlet conditions are well-known.

Finding the total pressure leaving the inlet, however, is more difficult. Pressure taps that measure the static pressure are often placed on the vanes near the trailing edges. The white balls in Figure 9 represent four possible locations for pressure taps. If temperatures are taken nearby, density can be calculated at these locations. These values are often used to calculate total pressure, since installing total pressure probes that capture both the static and velocity components is often impractical in a shop floor test.

Figure 10 shows the total pressure across the entire flow passage near the trailing edges, something only CFD can reveal. Total pressure is very non-uniform, with lower values opposite the nozzle. Consequently, measurements taken on the suction side of the vane nearer the bottom would show lower total pressure, and thus higher loss coefficient through the inlet.

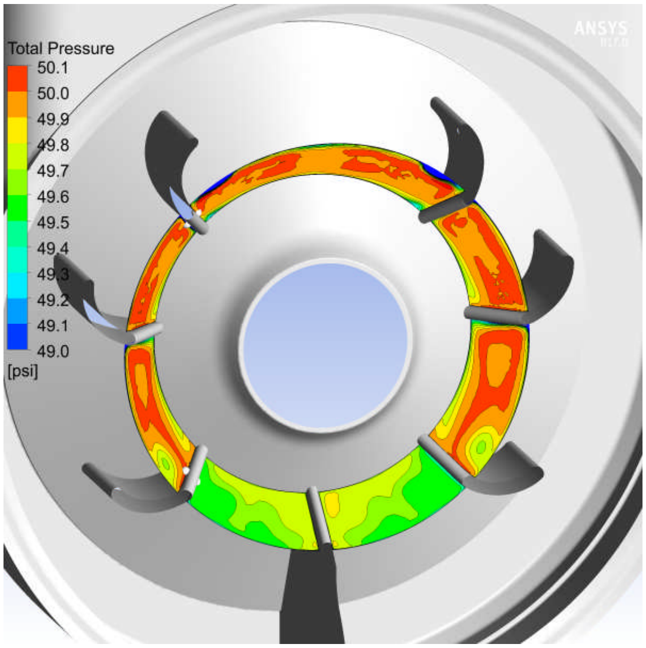

Figure 10. Total Pressure near compressor inlet vane trailing edges.

What is the average total pressure near the trailing edges? That depends on how it's calculated, and the difference can have a big impact on the loss coefficient that's calculated. The mass-averaged total pressure across the passage, a value that can be found only with CFD, is 49.852 psi. Calculating total pressure from test data is usually done by using the static pressure and density along with the mass flow and flow area. This approach yields a total pressure at the vane trailing edges of 49.152 psi, much lower than the mass-averaged CFD value. This difference in total pressure drop is enough to increase the calculated loss coefficient by almost a factor of five.

Why the big difference in calculated total pressures? There are three main reasons:

1. The surface where the averages are taken is oriented parallel to the trailing edges, not necessarily being a true quasi-orthogonal for the flow passage. So the area might be too large, resulting in too small a calculated velocity and thus too low a total pressure. Determining the area to use is a challenge for real tests, not just CFD, when total pressure is calculated from static pressure and mass flow.
2. CFD calculates total pressure using the entire velocity at every node, not just the component perpendicular to the flow surface. Calculations using mass flow and area, by nature, use only the component of velocity perpendicular to the calculation surface, neglecting secondary velocity and thus finding a lower total pressure. Which approach is correct depends on what happens downstream and can lead to long debates among engineers.
3. Mass-averaging total pressure is not exactly correct for finding an equivalent total pressure for a nonuniform flow field. There is no simple procedure to get a truly equivalent average total pressure.

Since one approach uses static pressure to calculate total pressure, the variation of the static pressure field shown in Figure 11 must be investigated.

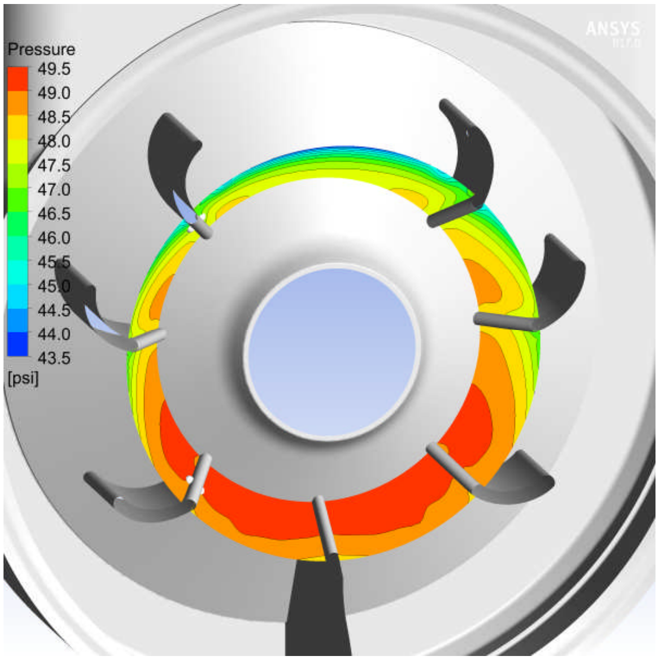

Figure 11. Static pressure near compressor inlet vane trailing edges.

The trend is opposite that of the total pressure. The high velocity near the nozzle results in low static pressure. When calculating total pressure from static pressure and mass flow, however, the same average velocity is assumed for every point, so the calculated total pressure tracks with the static pressure. Consequently, using the pressure tap locations on the lower vane results in higher calculated total pressures than using those on the upper vane, despite the total pressure field shown in Figure 11. In fact, the calculated total pressure is higher than the inlet total pressure, leading to a negative loss coefficient.

Considering the difference in loss coefficients found even for full-passage averages above, it is hard to know what the "right" answer is to compare with. Table 2 shows loss coefficients calculated from each of the possible pressure tap locations using the loss coefficient from mass-average values of total pressure as a reference value.

Table 2. Fractions of Inlet Loss Coefficient

Pressure Tap Location	**Loss Coefficient Relative to Mass-Averaged Total Pressure Value**
Upper vane, pressure surface	9.6
Upper vane, suction surface	10.6
Lower vane, pressure surface	-3.1
Lower vane, suction surface	-2.8

The calculated loss coefficient ranges from about -3 times the reference value (three times the magnitude and negative besides) to more than 10 times the reference value.

The further work required is to decide what the best reference value of CFD-based loss coefficient is, find instrumentation locations that give representative results, and then check those with CFD runs at different conditions to make sure they still give representative answers.

From Figure 11, one can see that even if total pressure probes could be used, locating them correctly would be challenging. CFD values of total pressure across the entire calculation surface shown vary from 45.287 to 50.179 psi. Depending on where a probe was placed, the loss coefficient could range from about -1 times the reference value (same magnitude but negative) to more than 33 times the reference value.

CFD does not provide a simple answer to this measurement challenge, but it illuminates the problem and gives some insight into how to approach it. In general, it is a good idea to check proposed instrument locations with CFD to see whether they represent the overall flow conditions well.

Loss coefficients for inlets and side loads are usually determined by better controlled tests and analyses rather than shop tests because of the uncertainty demonstrated in this example. Even using CFD to locate the instruments can't guarantee accurate calculation of loss coefficient. In this particular case, using the average of the values of static pressure from the vane pressure surface instrument locations results in a calculated total pressure of 49.518 psi, not terribly far off from the CFD mass-averaged value of 49.852 psi. Still, this small difference is enough to change the loss coefficient by a factor of 3.3 from the reference value.

Fortunately, calculations of section performance are far less sensitive to small variations in total pressure. Total pressure values calculated at the vane trailing edges of a side load can be used to find the performance of a downstream section with little error. Loss coefficient is particularly sensitive to small errors in total pressure because the total pressure drop through an inlet or sideload is so small. CFD can be used to locate instruments well enough to provide accurate section performance calculations.

Inlet Pipe Bends

A piping arrangement with an elbow closer to the compressor nozzle inlet than is recommended [Hackel & King] is shown in Figure 12 (white pipe) along with a comparison case with a straight inlet pipe (green pipe). CFD was used to evaluate the effect of the piping configuration on the inlet loss and the first impeller performance. Since CFD results should always be compared with a baseline case, the model with the straight green pipe was built and run as well.

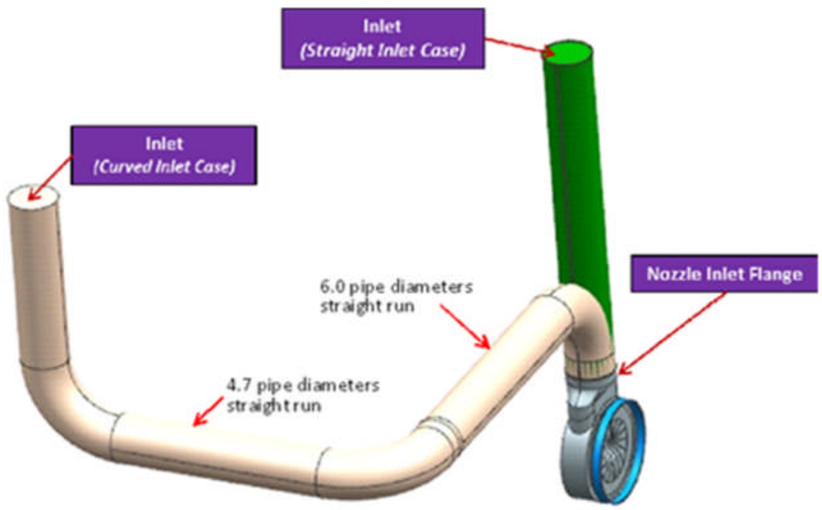

Figure 12. Inlet piping configurations.

The non-uniformity in total pressure caused by the inlet bends are shown in Figures 13 and 14. The contour plot scales are the same for the plots. Information this detailed would be impossible to get in a real machine.

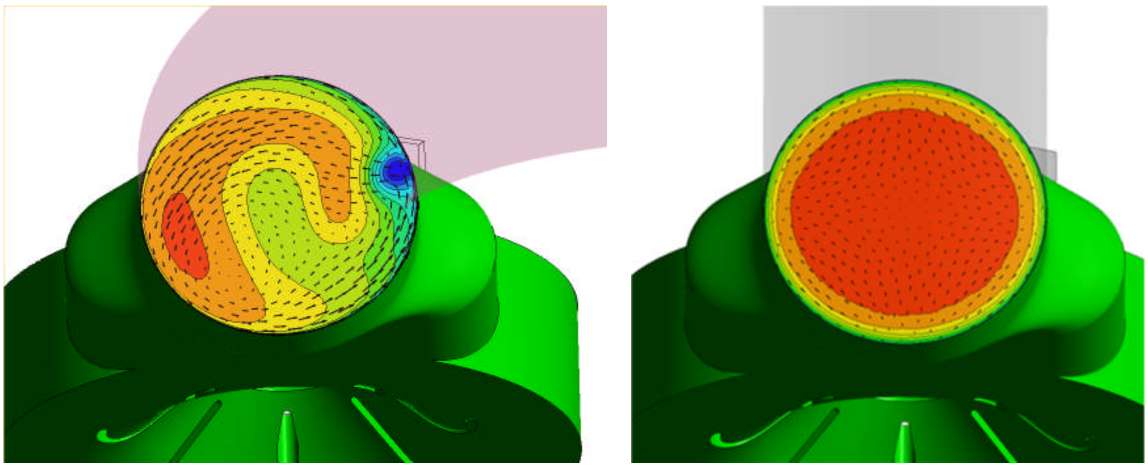

Figure 13. Total pressure at nozzle inlet with upstream bends (left) and without bends (right).

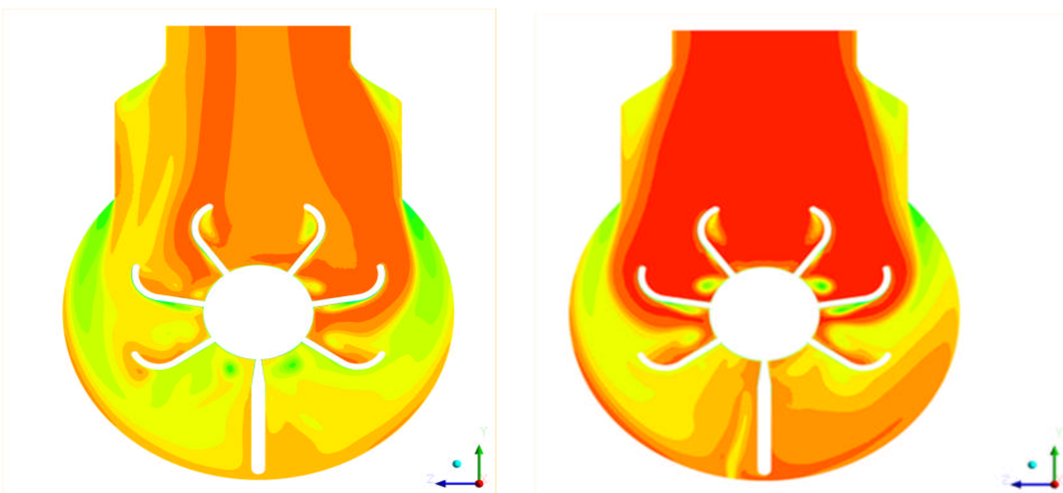

Figure 14. Total pressure on inlet cross-section with upstream bends (left) and without bends (right)

Losses and performance were calculated from the CFD results for the curved and straight inlet pipe cases and compared. With this particular geometry at these particular operating conditions, the penalty from the inlet elbows was small. However, this is not a general rule. Piping configurations that violate guidance such as recommended number of diameters from an elbow to the inlet should be avoided. If they are necessary, the performance impact should be evaluated with CFD and taken into account when setting expectations for the process.

Compressor Sideload

A multi-stage centrifugal compressor with two side loads is shown in Figure 15, where additional flow is added between stages. These are often called sidestreams. Questions arise about losses and mixing in side loads. CFD can give considerable insight into the flow field where the added flow meets the flow coming from the upstream stage.

Figure 15. Sideloads.

The velocity magnitudes in a side load are shown in Figure 16 [Elliott brochure]. The colors show that the velocities of the streams match, which minimizes mixing losses.

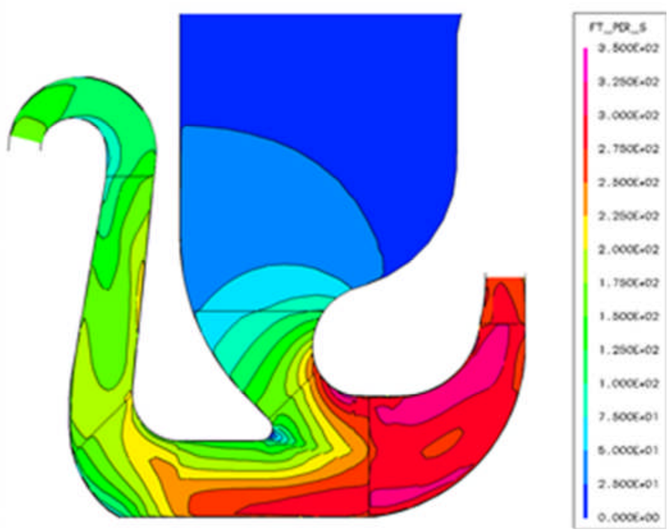

Figure 16. Side load velocity matching.

Of course, because there is little mixing, the fluid from the sideload and from the upstream return channel remain separate entering the impeller. Figure 17 shows the temperature where the two streams meet for a case where the sideload flow is cooler than the main stream flow. The non-uniform temperature entering the impeller can affect impeller performance. Making the flow more uniform would require mixing, with the associated mixing loss, upstream of the impeller. It is generally preferable to avoid taking an additional loss just to make the flow more uniform [Hohlweg *et al.* (2005)]. The trade-off between the harm due to non-uniform temperature and the harm due to mixing loss must be considered in a sideload design. The temperature non-uniformity generally causes little harm to the impeller performance, although the effect varies for different cases.

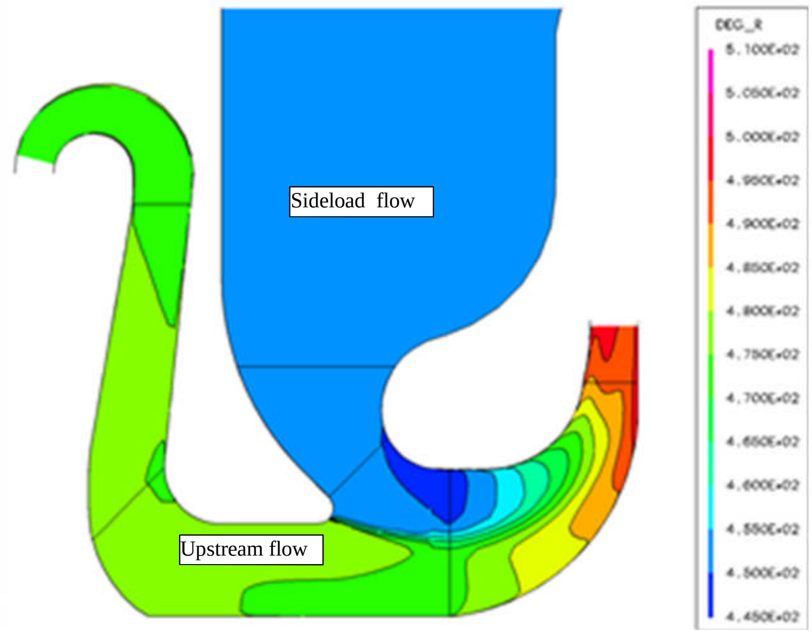

Figure 17. Sideload temperature

CFD can be used to compare the performance of the impeller with a baseline case to assess the impact of the non-uniform temperature. The baseline case could be the same geometry with matched temperatures in the two streams, or a simple axial inlet, or a hypothetical return channel exit bend sized for the entire flow.

Turbine Governor Valve

Turbine governor valves operate over large flow and pressure ratio ranges. In some areas of its performance map, flow through a valve can become unstable [Araki *et al.* (1981), Hardin *et al.* (2003), Zhang *et al*. (2004)]. Designing a new governor valve requires detailed knowledge of the flow field in the valve. As with most CFD analyses and design projects, it's best to start with an existing baseline design.

Figures 18 and 19 show pressure taps installed in an old hemispherical head governor valve design. Even with all this instrumentation, all that is found is pressure at several points on the surfaces of the valve.

Figure 18. Instrumented governor valve plug.

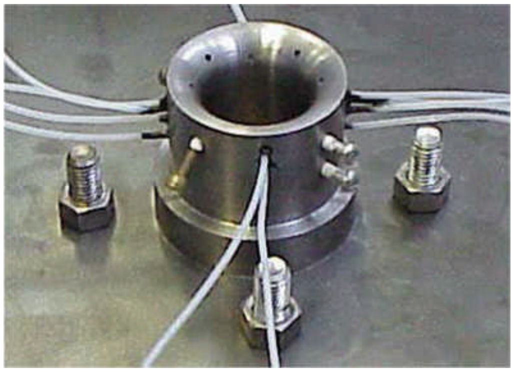

Figure 19. Instrumented governor valve seat.

Compare this with the detail available from CFD analyses of the same valve. Figure 20 shows Mach numbers throughout the valve at two different cross-sections, with different flowing gases and different up- and downstream configurations. These runs enable comparing results from the air test rig with results using steam in different possible steam chests. Among other things, they show the potential error created when evaluating a valve isolated from the actual system in which it would be installed.

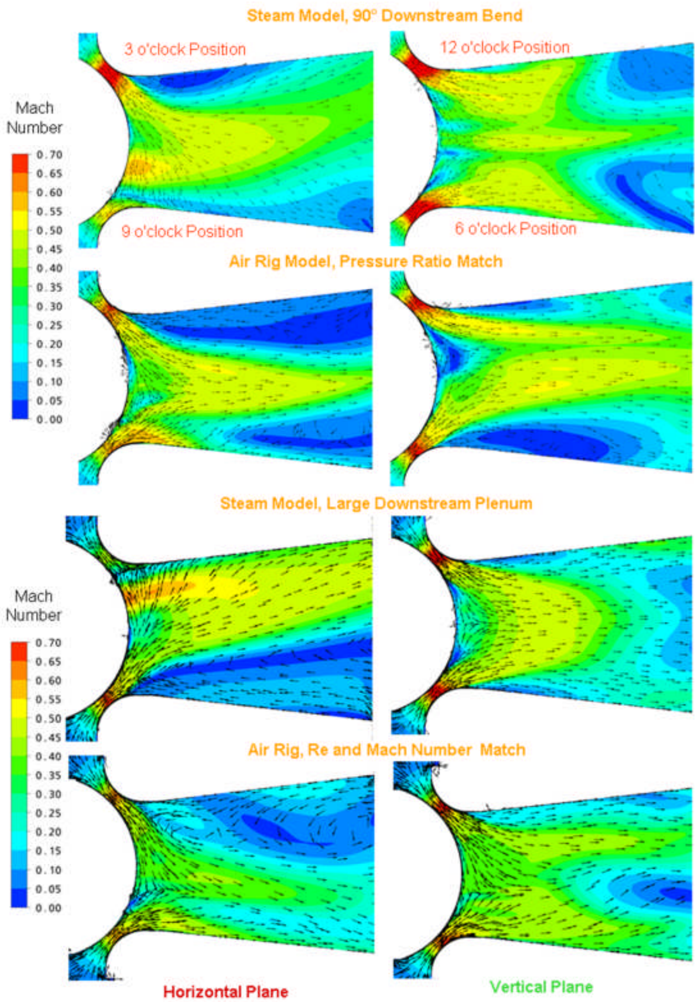

Figure 20. Mach number in governor valve.

Figure 21 shows total pressure through the seat of a proposed governor valve design at various valve openings. Again, this level of detail would not be available from any practical set of instrumentation.

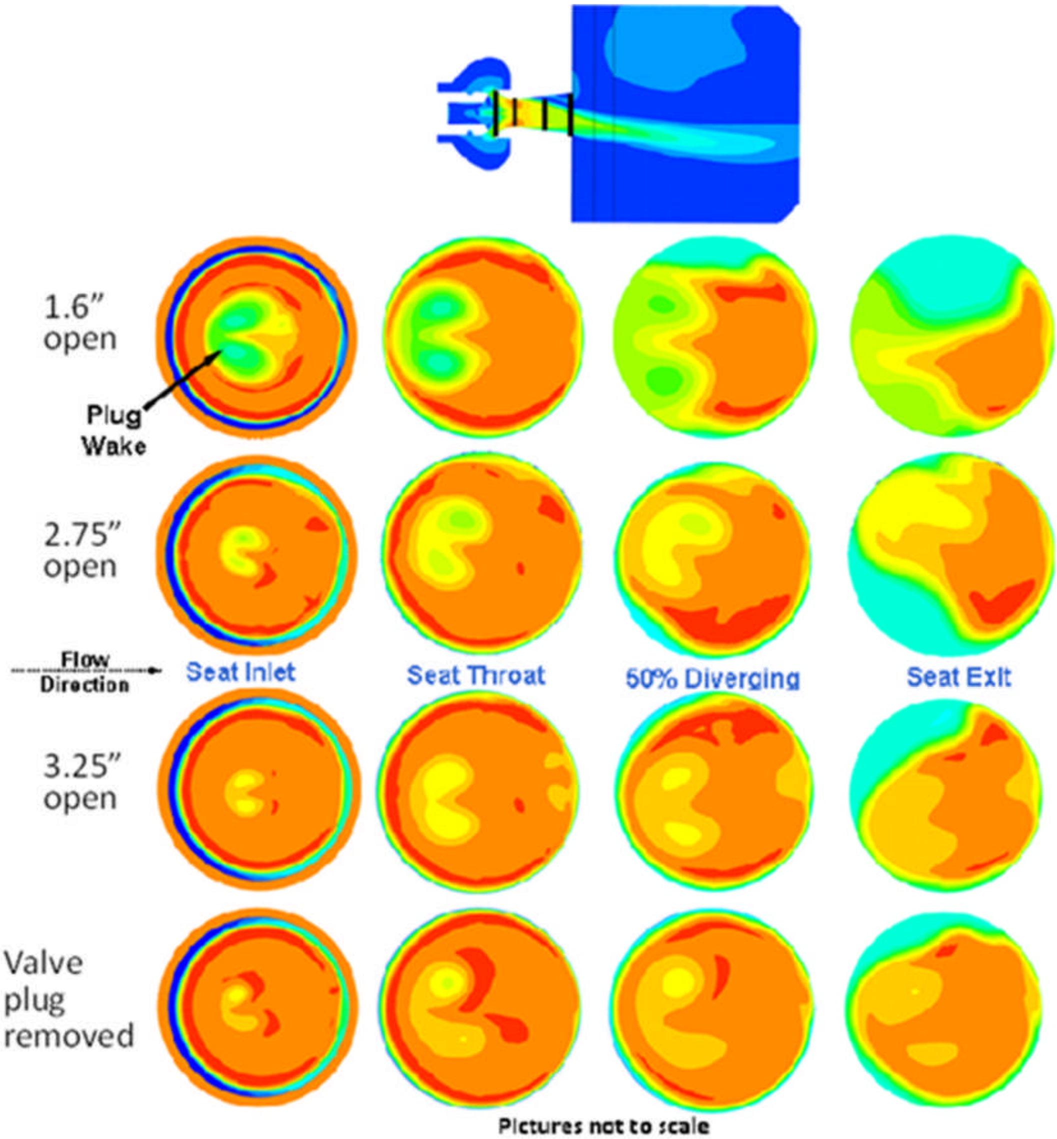

Figure 21. Total pressure in governor valve.

Centrifugal compressor impeller grid refinement study
As stated earlier, the computational mesh has a significant impact on the solution, and it is always wise to understand the magnitude and nature of that effect. A rather simple case has been chosen to demonstrate that point: a centrifugal compressor impeller of moderate pressure ratio. Three high-quality hexahedral meshes were created on a domain that included the impeller plus short vaneless sections upstream and downstream. The coarse mesh had about 140,000 nodes (Figure 22), the medium about 350,000 nodes and the fine about one million nodes. The "character" (nodal arrangement) of all three meshes was the same, and care was taken to ensure that expected boundary layers were well represented and refined as the overall mesh size increased. Note that all of these mesh sizes are for a single blade passage, and not for the full impeller.

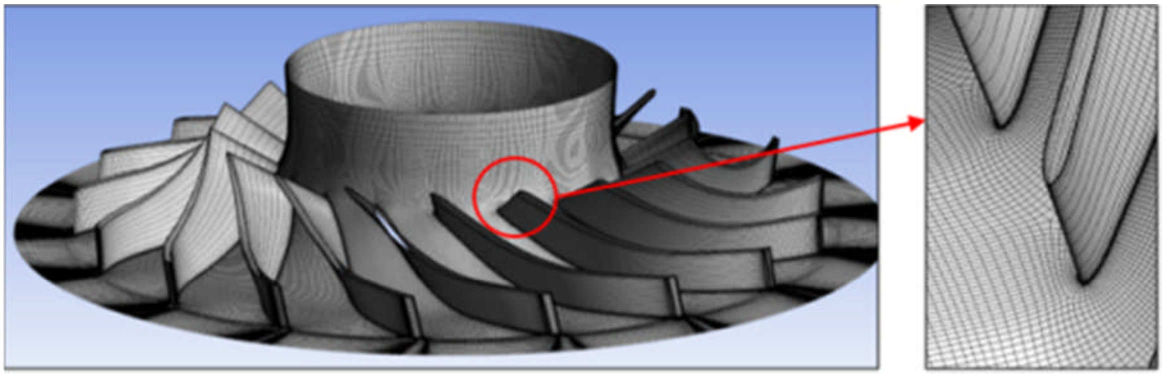

Figure 22. Full coarse mesh displayed on the hub and blades of a centrifugal impeller. The magnified view at the right shows the clustering of the mesh near the blades.

Figure 23 shows the relative Mach number, meridionally-averaged, in the meridional plane. These medium grid results were obtained for one impeller passage using the ANSYS CFX 17.0 solver with standard settings and the SST turbulence model. Near design mass flow (right image), an acceleration of the flow in the inducer section of the impeller at the shroud can be seen. Near stall (left image), the flow rate is lower and the relative Mach number is, therefore, lower. That acceleration region is greatly reduced along the shroud and instead a wake (low Mach number) region (the cyan color) develops as can be seen in the left image.

Figure 24 provides the quantitative results for the flow map, ranging from near-stall at the left to overload at the right. These predictions of total pressure ratio and efficiency (relative to the coarse grid design point efficiency) indicate only small changes in value as the mesh size is refined. But there is both a change and a pattern. Both total pressure and efficiency increase as the mesh is refined. The change is small for the total pressure, but larger for the efficiency, indicating that increasing the mesh size refines the predicted loss and, therefore, work input declines. The finer mesh also serves to reduce numerical errors, which manifest themselves as (non-physical) losses. Since differences can be seen between the medium and the fine mesh predictions, it is clear that grid independence is not achieved. Simulation on a still finer mesh is required to determine at what point the solution stops changing.

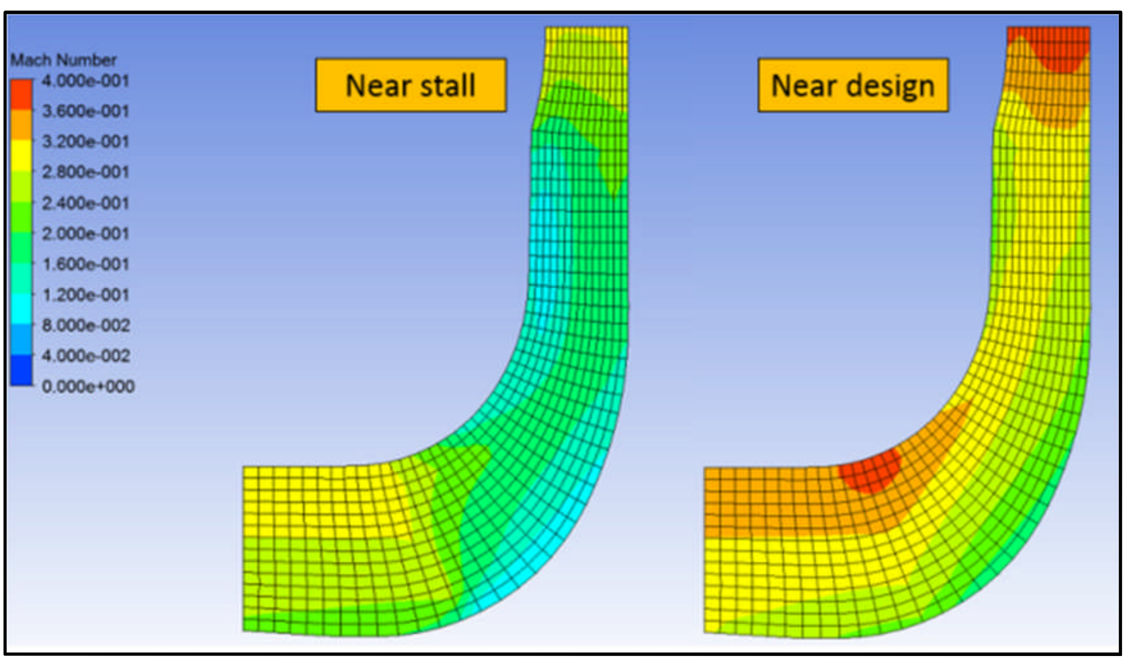

Figure 23. Circumferentially-averaged relative Mach number, near stall (left) and near design (right).

Analysts often face this situation in turbomachinery CFD. That is, due to limitations of time and computing resources one must be satisfied with using meshes that do not produce mesh-independent results. However, it is important to understand the implications of that shortcoming. In this case it means that our solutions are producing results that are a little lower in terms of pressure ratio and efficiency than mesh independent results. One can also note that other features have probably been neglected, such as fillets or wall roughness (this was a smooth wall simulation) that would serve to push the results in the opposite direction. In the end, the analyst must relate the simulations to test and understand the many factors that would explain the differences. In any case, CFD is invaluable in the design process in that it can generally help the analyst decide which is the better design, even if the magnitude of the prediction is not exact.

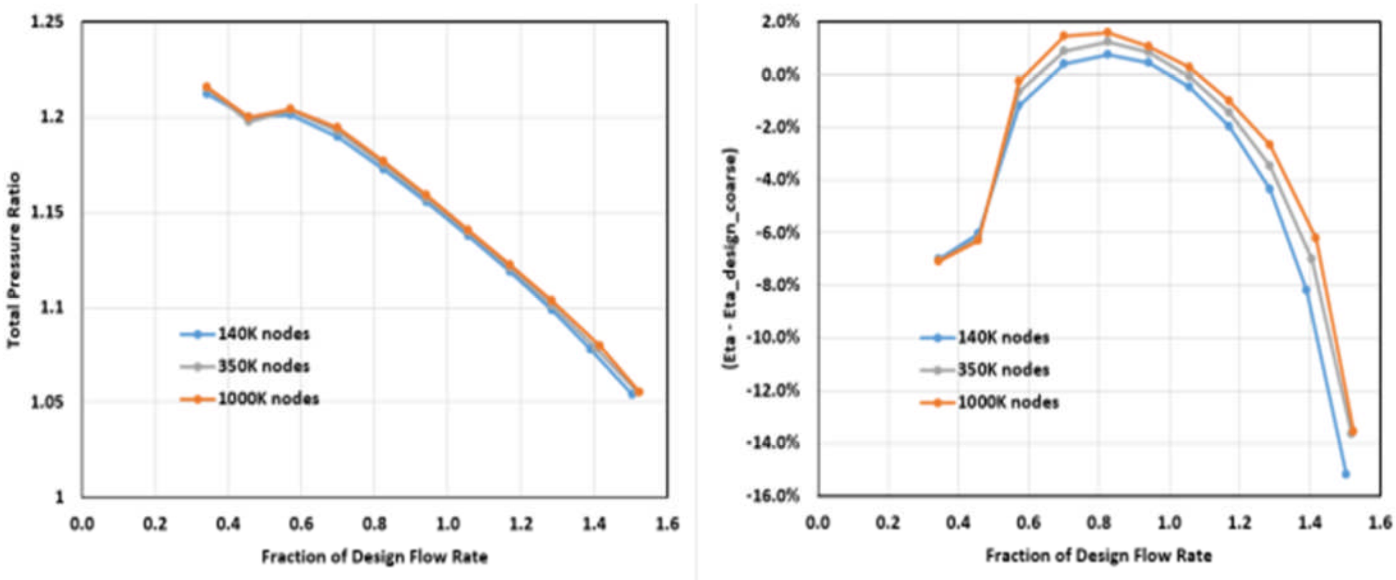

Figure 24. Predicted centrifugal impeller total pressure ratio (left) and isentropic efficiency relative to the design flow, coarse grid result (right).

Design Iterations
As noted previously, one of the most powerful ways to use CFD is as a comparative tool. This can be particularly useful when developing new centrifugal components or when seeking to improve the performance of an existing stage. It is well-established that the relative changes indicated by CFD will be reflected in the "real world." That is, if CFD results indicate that a revised geometry provides better operating range or higher efficiency than an earlier design or design iteration, the designer can be confident that the revised geometry will provide improved performance in the real machine.

As an example, consider the CFD results shown in Figure 25. The contour plot to the right shows the relative Mach number distribution for an existing impeller design. Based on testing, the impeller had been short in surge / stall margin. The distribution on the right shows that a large wake or low momentum region (indicated by the light green and blue colors) is forming along the shroud near the blade suction surface. It was concluded that this low momentum zone was responsible for the limited stall margin. In response, a revised blade geometry was developed with the goal of reducing the size of the low momentum zone. The CFD-generated flow field in the left in Figure 25 shows that the low momentum region (again, indicated by the blue region) is much smaller in the updated design. As expected, the revised impeller provided much better surge / stall margin than its predecessor.

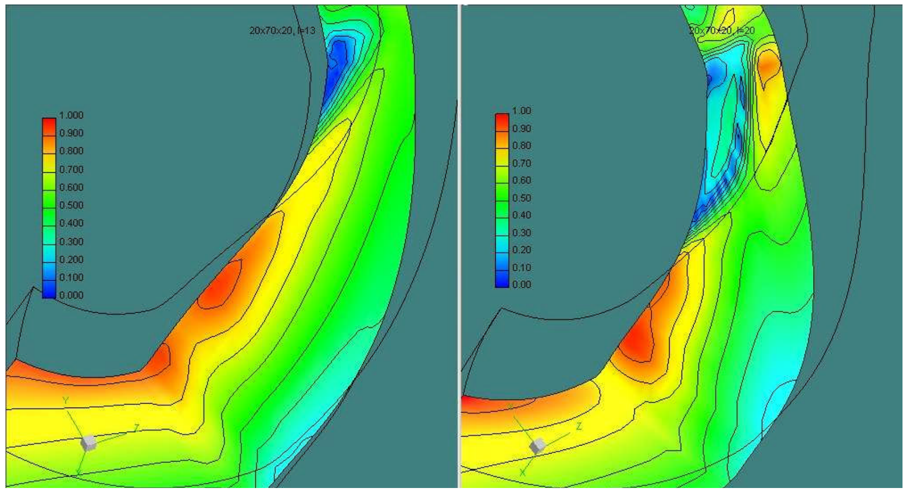

Figure 25. Relative Mach number contour plots for the new (left) and original (right) impellers.

CONCLUDING REMARKS

In closing, this paper provided a general overview of computational fluid dynamics (CFD) and hopefully provided the non-expert with a better understanding of what CFD is, some of the modeling and approximation techniques used to simplify the analyses, where it fits in the design/analysis process, what results can be obtained and how those results might be presented; *i.e.*, the style of plots and/or schematics typically used to review the results of these advanced simulations. Several sample cases were offered to show the practical ways that CFD can be applied.

Those interested in digging deeper into CFD are encouraged to review any of the references listed in the bibliography or to attend a seminar or course offered by the various CFD code providers.

As computer hardware and computational speed continues to advance, developers will likewise grow the capability of CFD solvers. More advanced turbulence models, adaptive gridding schemes and other modeling methods will bring these mathematical simulations ever closer to the "real world," further enhancing CFD's critical role in the development of new products and applications.

REFERENCES

Ansys documentation 1, "Coupling CFX to an External Solver: ANSYS Multi-field Simulations," ANSYS documentation, version 17.0, CFX Modeling Guide section 13.2, ANSYS Inc., Canonsburg, PA, 2016.

Ansys documentation 2, "Aerodynamic Noise Analysis," ANSYS documentation, version 17.0, CFX Modeling Guide section 14, ANSYS Inc., Canonsburg, PA, 2016.

Araki, T; Okamoto, Y; Ootomo, F, "Fluid Induced Vibration of Steam Control Valves",Toshiba Review, Vol. 36, issue 7, ISSN 0372-0462, Tokyo Shibaura ElectricCo., Kawasaki, Japan, pp.648-656, 1981.

Drosjack, M., Sorokes, J.M., Miller, H.F., "Buying/Selling Serial #1," Turbomachinery Symposium Proceedings, Texas A&M, 2011

Durbin, P., and Pettersson Reif, B.A., Statistical Theory of Turbulent Flows, John Wiley and Sons, Ltd, UK, 2003.

Elliott brochure, "Engineered Solutions, Superior Technology from a Global Supplier," Brochure ENG2 1501FL, Elliott Group, Jeannette, PA.

Guerts, J. B., Elements of Direct and Large-Eddy Simulation, Edwards Inc., Philadelphia, USA, 2004.

Hackel, Ross A. and King, Raymond F. Jr., "Centrifugal Compressor Inlet Piping - A Practical Guide," Compressed Air and Gas

Institute, Vol. 4, No. 2, Cleveland, OH.

Hardin, James R. and Boal, Charles F., "Using CFD To Improve Stall Margin," Proceedings of the American Society of Mechanical Engineers Fluids Engineering Division - 1999, FED-Vol. 250, pp 115-121, 1999.

Hardin, James R., "A New Approach To Predicting Centrifugal Compressor Sideload Pressure," Proceedings of IMECE2002, Technical Paper IMECE2002-39592, American Society of Mechanical Engineers, 2002.

Hardin, J., Kushner, F., Koester, S., "Elimination of Flow-Induced Instability From Steam Turbine Control Valves," Proceedings of the Thirty-Second Turbomachinery Symposium, Turbomachinery Laboratory, Texas A&M University, College Station, Texas, pp. 99-108, Sept. 2003.

Hohlweg, W., Blahovec, J., Wright, J., "Design Considerations for Centrifugal Compressors in Refrigeration Duty," Paper 42b, 2005 AIChE Spring National Meeting, American Institute of Chemical Engineers, April 2005.

HP Enterprise, "Scalability of ANSYS 16 applications and Hardware selection," white paper, available from www.ansys.com Resource Library, Hewlett Packard Enterprise, Palo Alto, CA.

Hutchinson, B.R., Shi, F., Sorokes, J.M., Koch, J.M., "Investigation of Advanced CFD Methods and their Application to Centrifugal Compressors," IMECE2002-39588, 2002

Koch, J.M., Sorokes, J.M., Belhassan, M., "Modeling and Prediction of Sidestream Inlet Pressure for Multi-stage Centrifugal Compressors," Turbomachinery Symposium Proceedings, Texas A&M, 2011

Menter, F.R., "Two-Equation Eddy Viscosity Models for Engineering Applications," AIAA Journal, 32(8), 1598-1605, 1994.

Menter, F.R., "Best Practice: Scale-Resolving Simulations in ANSYS CFD", V2.0, Ansys, Inc., 2015

Patankar, S.V., Numerical Heat Transfer and Fluid Flow, CRC Press, 1980.

Slagter, Wim *et al*, "Industry Perspectives on Extreme Scalability for High-Fidelity CFD Simulations," presentation, available from www.ansys.com Resource Library, ANSYS Inc., Canonsburg, PA, 2014.

Sorokes, J.M., "The Practical Application of CFD in the Design of Industrial Centrifugal Impellers," Turbomachinery Symposium Proceedings, Texas A&M, pp 113 - 124 (1993).

Sorokes, J.M. and Hutchinson, B.R., "The Practical Application of CFD in the Design of Industrial Centrifugal Compressors," Challenges and Goals in Industrial and Pipeline Compressors, PID-Vol. 5, pp 47 – 54, 2000

Sorokes, J., Pacheco, J, Malnedi, K, "Pushing the Envelope -- CFD simulation contributes to increasing the operating envelope of a centrifugal compressor stage," Ansys Advantage Magazine, Fall 2013

White, F.M., Viscous Fluid Flow, 2nd Edition, McGraw-Hill, Inc., 1991.

Zhang, D.; Engeda A.; Hardin, J. R.; Aungier, R. H.; "Experimental Study of Steam Turbine Control Valves," Proceedings of the I MECH E Part C Journal of Mechanical Engineering Science, Vol. 218, No. 5, 1 May 2004, pp. 493-507.

Zigh, Ghani and Solis, Jorge, "Computational Fluid Dynamics Best Practice Guidelines for Dry Cask Applications," NUREG-2152, United States Nuclear Regulatory Commission, Washington, DC, 2013.

"Standard for Verification and Validation in Computational Fluid Dynamics and Heat Transfer," ASME V&V 20-2009, American Society of Mechanical Engineers, New York, NY, 2009.

"Performance Test Code on Compressors and Exhausters," ASME PTC 10-1997, American Society of Mechanical Engineers, New York, NY, 1998.

ACKNOWLEDGEMENTS

Mr. Sorokes thanks the various members of the Aero/Thermo Design team (present and past) for their assistance in assembling the graphics used in the paper and thanks Dresser-Rand, A Siemens Business for allowing him to participate on this tutorial. Mr. Hardin thanks Elliott Group for allowing him to participate in this tutorial, and personnel at QuEST Global Services NA and Michigan State University for their parts in some of the projects presented here. Dr. Hutchinson thanks Ansys, Inc. for allowing this paper to be submitted and their participation in the tutorial, and Bill Holmes of ANSYS for running the centrifugal impeller grid refinement test case.

