Emerging CFD Capabilities and Outlook – A NASA-Langley Perspective

Robert T. Biedron, S. Paul Pao, and James L. Thomas
NASA Langley Research Center

COMSAC Symposium
September 23, 2003

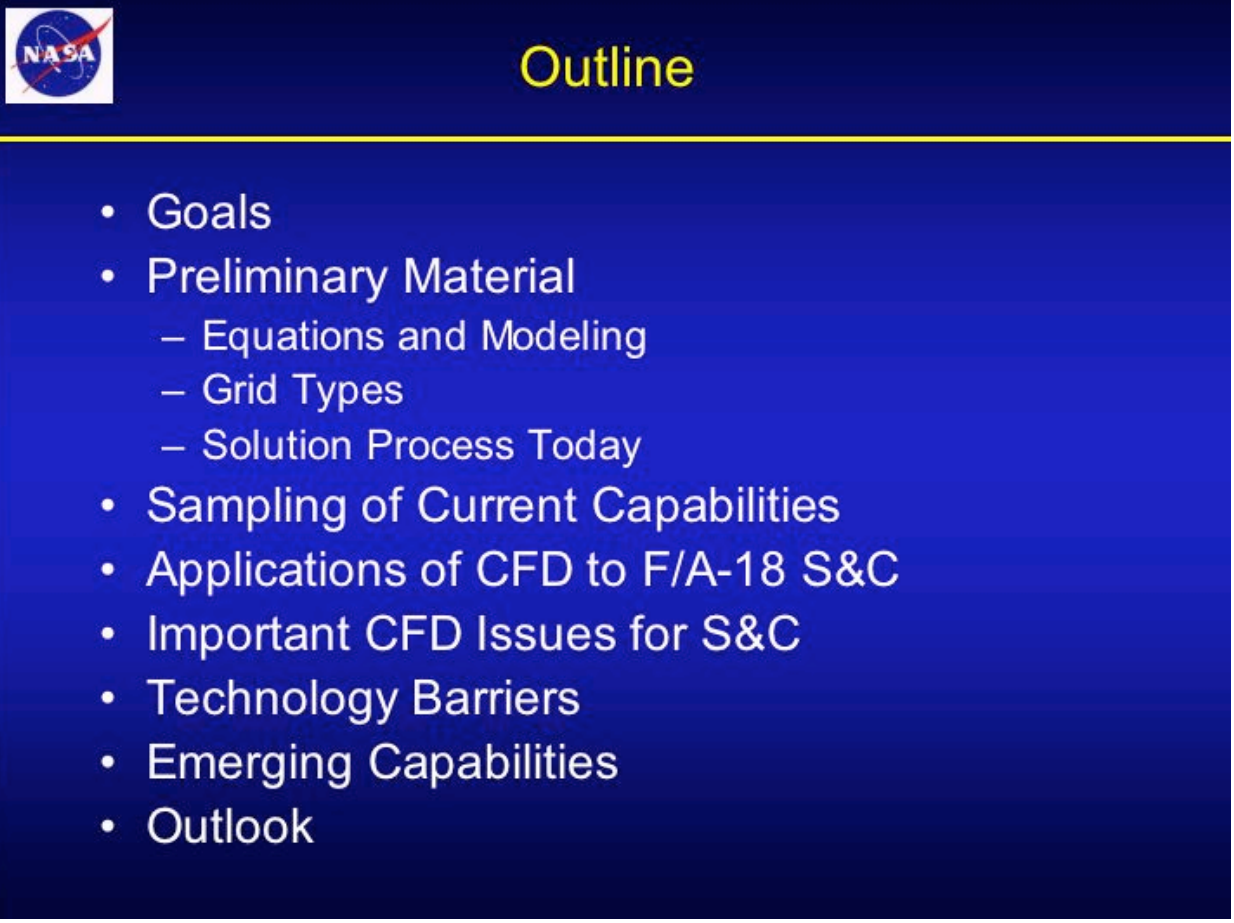
Outline
• Goals
• Preliminary Material
– Equations and Modeling
– Grid Types
– Solution Process Today
• Sampling of Current Capabilities
• Applications of CFD to F/A-18 S&C
• Important CFD Issues for S&C
• Technology Barriers
• Emerging Capabilities
• Outlook

COMSAC goals include increasing the acceptance of CFD as a viable tool for S&C predictions, as well as to focus CFD development and improvement towards the needs of the S&C community. We view this as a symbiotic relationship, with increasing improvement of CFD promoting increasing acceptance by the S&C community, and increasing acceptance spurring further improvements.

In this presentation we want to provide an overview for the non CFD expert of current CFD strengths and weaknesses, as well as to highlight a few emerging capabilities that we feel will lead toward increased usefulness in S&C applications.

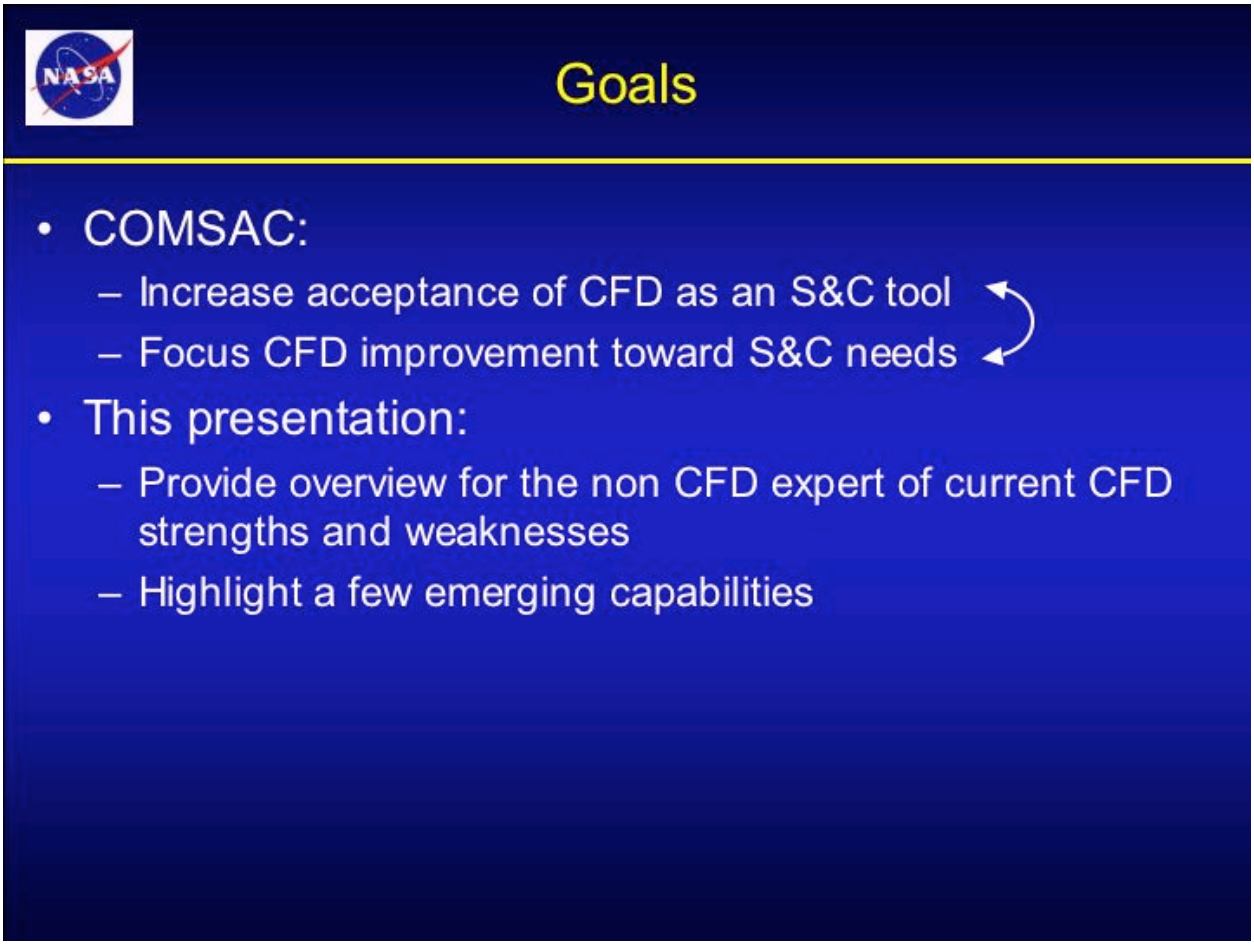

"CFD" can imply different things to different people. To put everyone on the same footing for this presentation we are going to restrict the definition of CFD to imply the numerical solution of the Navier-Stokes equations, with the Euler equations as a subset.

There are of course many levels of approximation to the Navier-Stokes equations, principally distinguished by how turbulence is treated. In principle one can use a fine enough mesh and a small enough time step to resolve and track all the important scales of the turbulence. Such solutions for full aircraft geometries are many decades away. So for practical applications we must resort to some level of turbulence modeling.

At the highest level of turbulence modeling lie the Reynolds Averaged Navier-Stokes (RANS) solvers. These solvers require O(10^7) points for a complete configuration in order to have an adequate resolution of the flow field. In a RANS code, the effect of turbulence is entirely

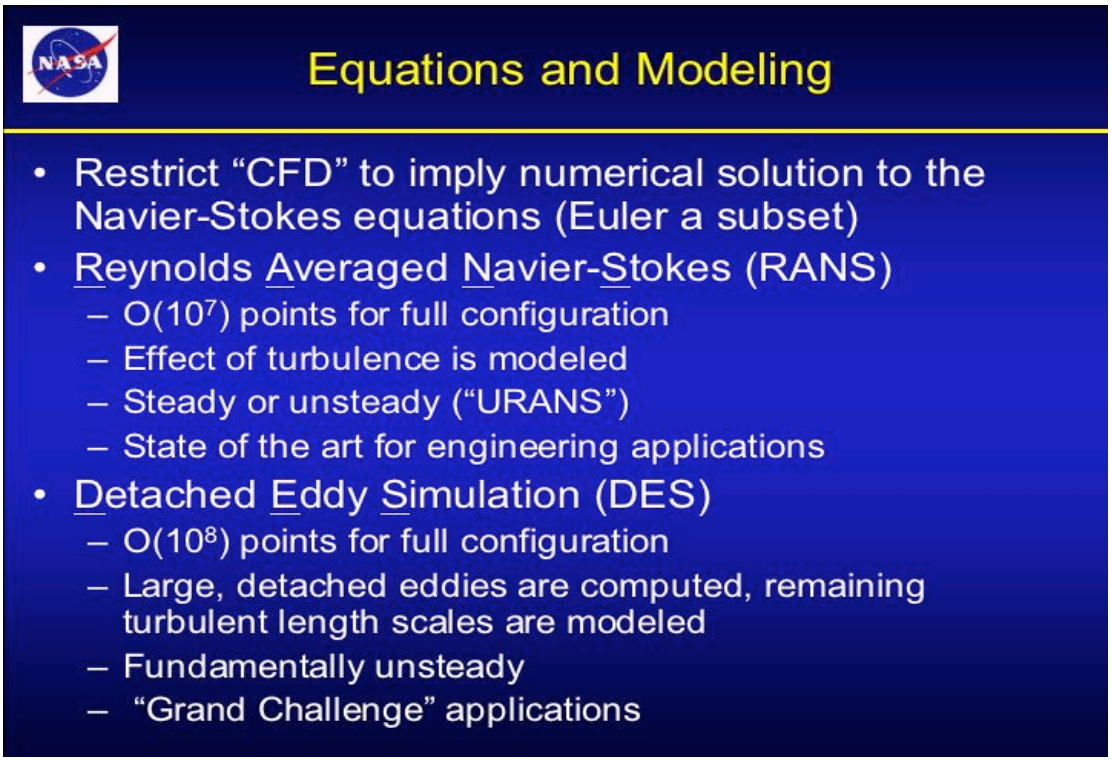

modeled; the grid is not dense enough to realistically track individual turbulent eddies anywhere in the field. Many RANS codes can simulate both steady and unsteady flows; those that simulate unsteady flows are often referred to as URANS. One might label RANS solvers as "state-of-the-art" for engineering applications.

RANS codes tend to do a poor job with massively separated flows. Detached Eddy Simulation (DES) is one of a number of hybrid methods that have been proposed to better deal with these flows. A well-resolved DES calculation may require O(10^8) points. Large, detached eddies in the separated flow region are computed and tracked, while the remaining turbulent length scales (near the body) are modeled with the RANS approach. A DES simulation is fundamentally unsteady. The need for very large numbers of grid points and time-accurate simulation puts DES out of the engineering realm at this time, into the "Grand Challenge" category.

All CFD codes employ some sort of grid or mesh, and may be categorized into two general types by the kind of grid used.

Structured grids are comprised of body-fitted hexahedral cells. In order to fit around complex geometries, structured grids must usually be made up of multiple blocks or zones of points. These zones may either abut against one another, or may overlap. Overlapped grids are often referred to as overset grids. Flow solvers utilizing structured grids make use of the inherent connectivity (or structure) between the cells, resulting in relatively fast flow solvers.

The other major category of flow solvers utilize unstructured grids. Such grids are also body fitted, but are usually comprised of tetrahedral cells, although prisms, pyramids and hexahedra

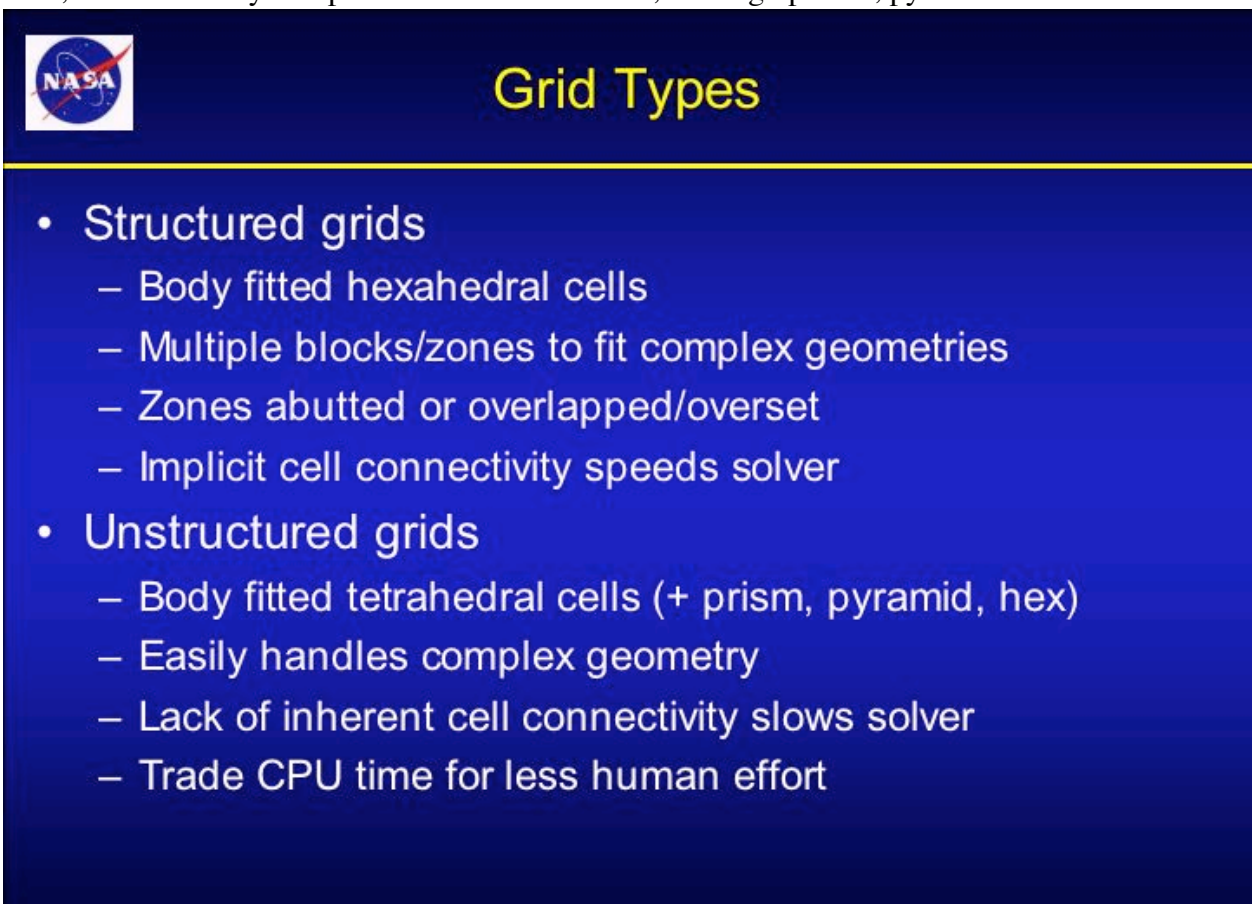

are often used. Prisms are particularly well suited for use within the boundary layer. Unstructured grids can readily handle complex geometry, easing the grid generation problem. The lack of any inherent connectivity between cells means that nearby neighbors of each cell must be explicitly spelled out for the solver, resulting in a slower speed than a comparable structured-grid solver. However, the ease of grid generation, as well as the relative ease of adaptation, discussed in subsequent slides, makes unstructured solvers very attractive – it's a tradeoff of more CPU time for less human effort.

To set the stage for later discussion it is worthwhile to give a very broad overview of the CFD solution process in today's environment. First, it is important to emphasize that the quality of a CFD result hinges on the quality of the grid.

When generating the grid (typically a substantial undertaking), the CFD expert or grid expert (ideally one person expert in both areas) decides where to place points. Typically points are clustered to resolve geometrical features and flow features using established "best practices". After the grid is generated, the flow solver is run and the process stops with a solution on this grid.

Adaptation carries the process one or more steps further. Given a solution on the baseline grid, obtained as described above, points are added to better resolve "important" features, typically

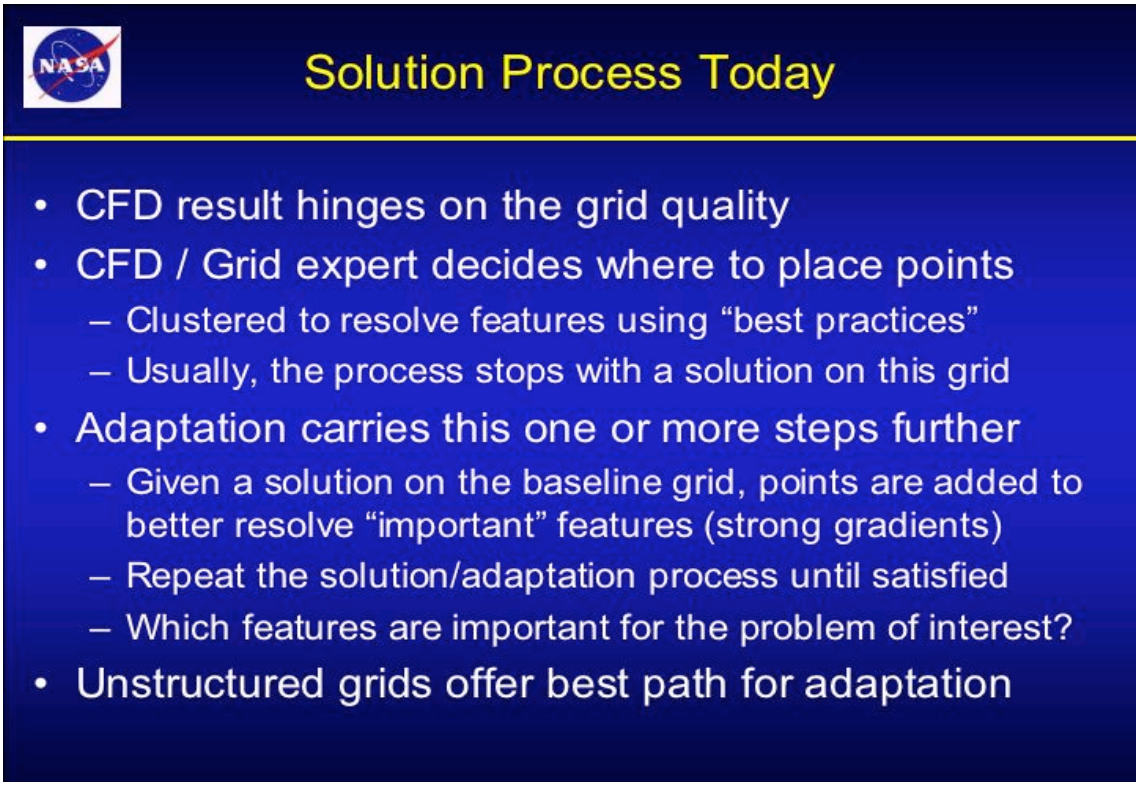

indicated by regions of strong gradients. Then the solution/adaptation process is repeated until the user is satisfied (or gives up…). A fundamental question, sometimes not easily answered even by the expert, is what features and regions are important to resolve in order to get the desired results for the problem of interest?

Adaptation may be done in with formal process as described above. More often however, there is a less formal approach to "adaptation" that occurs when the original grid fails to produce the desired result. In these cases the grid is changed in a manner deemed to result in a better CFD prediction from subsequent simulations. Unstructured grids offer the best path for adaptation as points can more easily be inserted into the grid as needed.

This slide and the next are intended to give an overview of some of the current capabilities of CFD.

Here we show a solution-adapted grid for a Modular Transonic Vortex Interaction (MTVI) configuration. This modular wind-tunnel model was used to provide detailed experimental data for CFD validation with a wide variety of vortical flows. The image in the lower right shows the grid in the vortical flow regions after refinement. The refinement was based on the locations of strong off-body vorticity as computed by the flow solver. The chine and wing leading edge vortices are indicated. The image in the upper left shows the corresponding flow visualization from the experiment. In this case the adaptation process was critical to the prediction of vortex bursting, which in turn was critical to predicting pitch up.

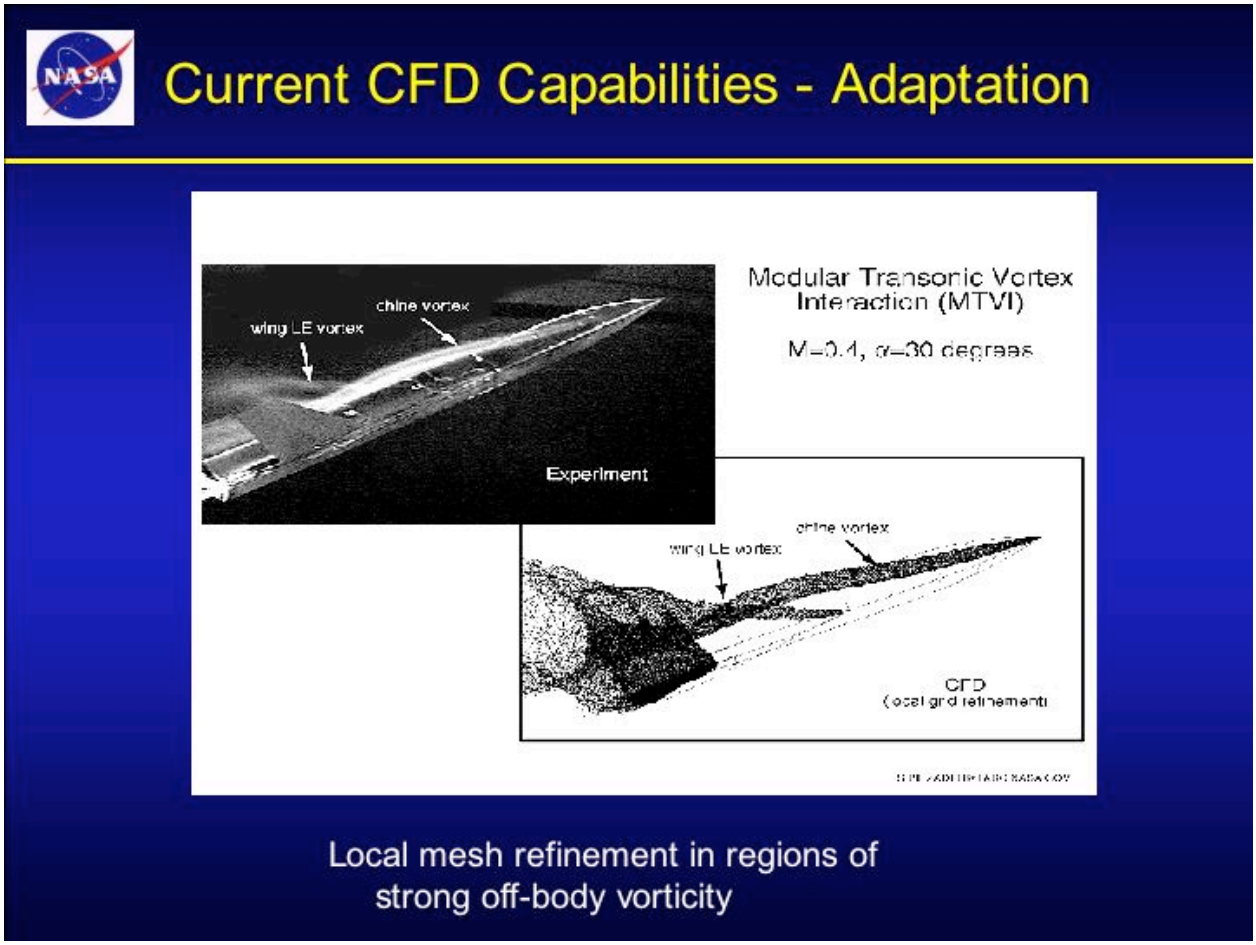

This slide shows a sampling of the wide range of problems that have been tackled using CFD.

Today's CFD codes are capable of simulating a wide variety of steady and unsteady flows over complex configurations. An example of a steady flow simulation is that shown for the S3 Viking in the lower left.

Complex flow interactions may arise even for simple geometries, but are the norm for complex configurations. An example is the full-stack Space Shuttle simulation in the upper left. Here we see the multiple shock waves formed during ascent.

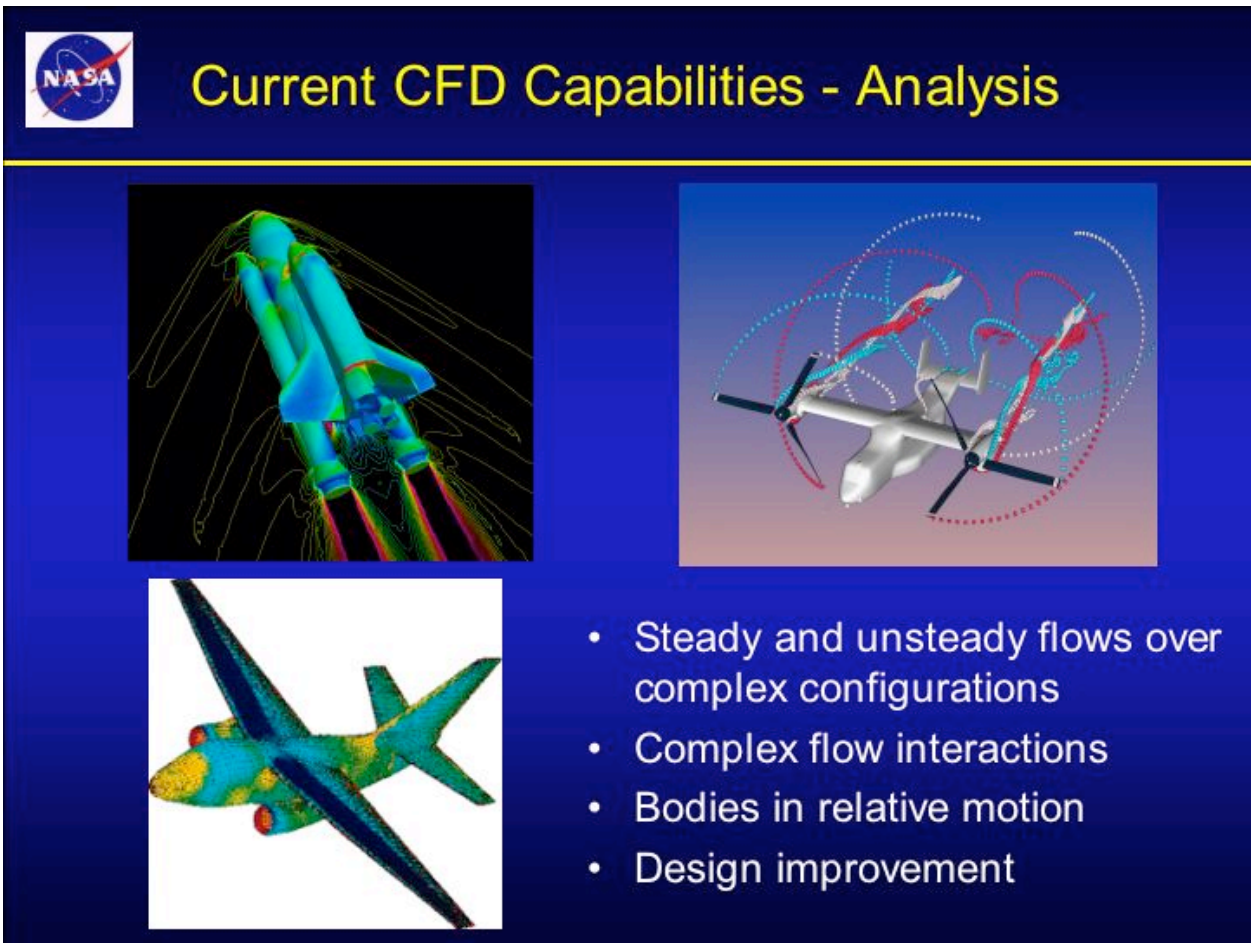

Unsteady simulations may include bodies in relative motion. In the upper right is shown a computation of the V22 Tilt Rotor with the spinning blades in the cruise mode. The complex wake structure is made visible by particle traces emanating from the rotors and wing tips.

Although not shown, analysis capabilities have been extended into the design environment in limited applications. More precisely, the use of CFD in design has been largely limited to design improvement, rather than for conceptual design. However, efforts are underway to bring CFD earlier into the design process.

Now let's look at two applications of CFD to stability and control of the F/A 18 aircraft, in the areas of Forebody Controls and Abrupt Wing Stall.

In the High Alpha Technology Program, a number of novel control effectors for maneuvering at high angles of attack were studied. Among them was an actuated nose strake, designed to be deployed on an as-needed basis. Early on in the wind tunnel program, using a generic strake design, it was observed that the direction of the yawing moment produced by the strake changed with deployment angle. CFD was used to confirm the experimental observation for the final strake design intended for use on the aircraft. The image on the top right shows (left to right) the vortical structure with the strake retracted, extended 10 degrees, and 90 degrees (full extension). The image on the lower right shows the computed yawing moments (symbols) for those three strake positions, as well as the wind tunnel measurements for a range of deflections (line). It is seen that CFD simulation has correctly predicted the yawing moment reversal, and has done a reasonable job at predicting the magnitude of the yawing moment.

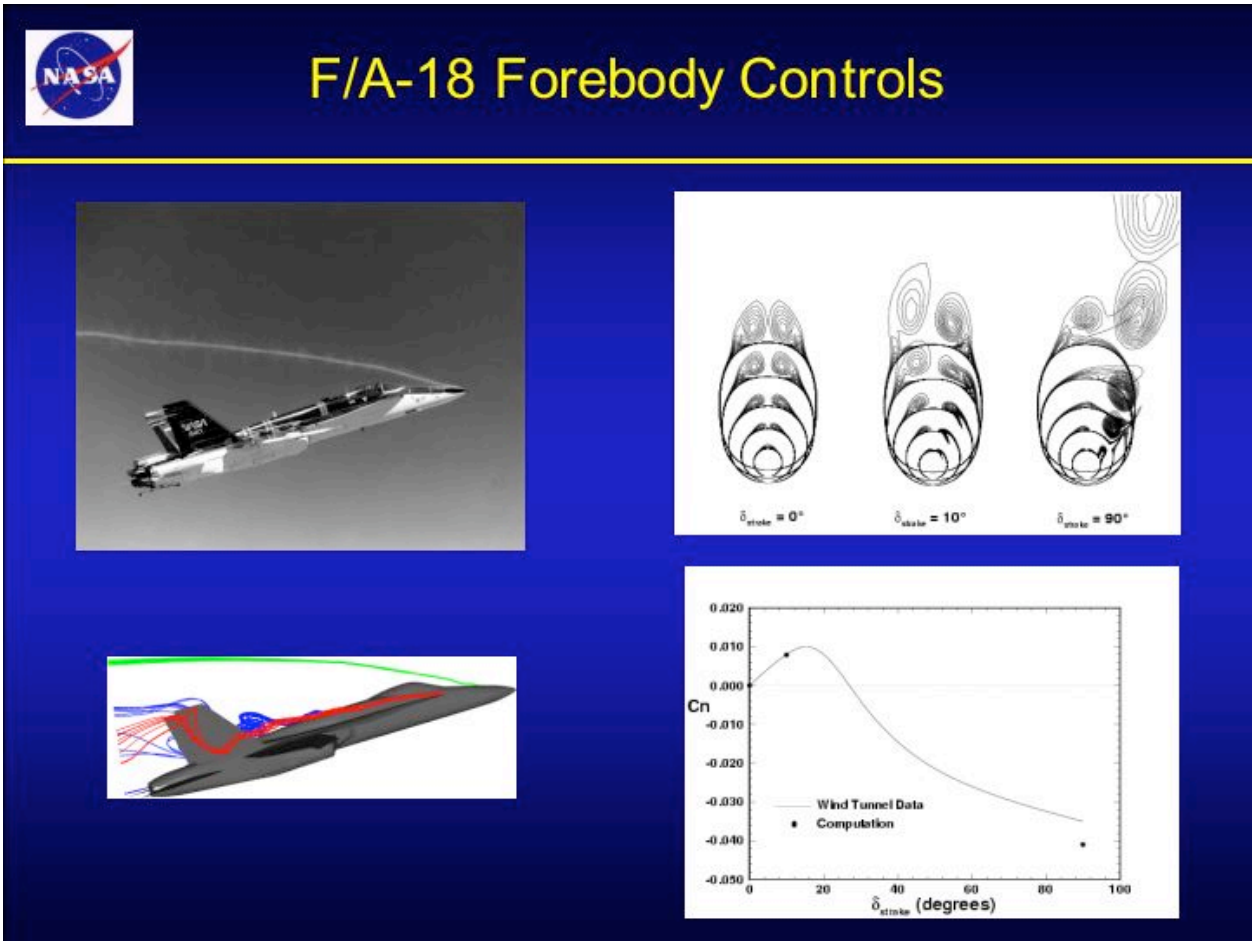

The image in the top left shows in-flight visualization of the vortex generated by the nose strake (fully deployed); the image in the lower left shows the corresponding visualization from the computation (green trace). Also shown are the traces of the LEX vortices (red and blue traces); the asymmetrical LEX vortex breakdown induced by the nose strake deployment is evident.

This slide shows a more recent application, by Jim Forsythe using the Cobalt code, for the prediction of Abrupt Wing Stall on the F/A-18E. The "E" variant exhibited an abrupt wing drop for a certain range of Mach number and angle of attack. Initial attempts using the RANS approach with either the Spalart-Allmaras (SA) model or the Shear Stress Transport (SST) model missed the shock location and thus gave incorrect predictions for the lift coefficient. In the case of the SA model the lift coefficients were in general too large, and though the break in the lift curve slope was correctly predicted near nine degrees angle of attack, the reduction in slope was under predicted. Conversely, using the SST model, the lift coefficients in the attached flow region were well predicted, but the break was predicted to occur too soon.

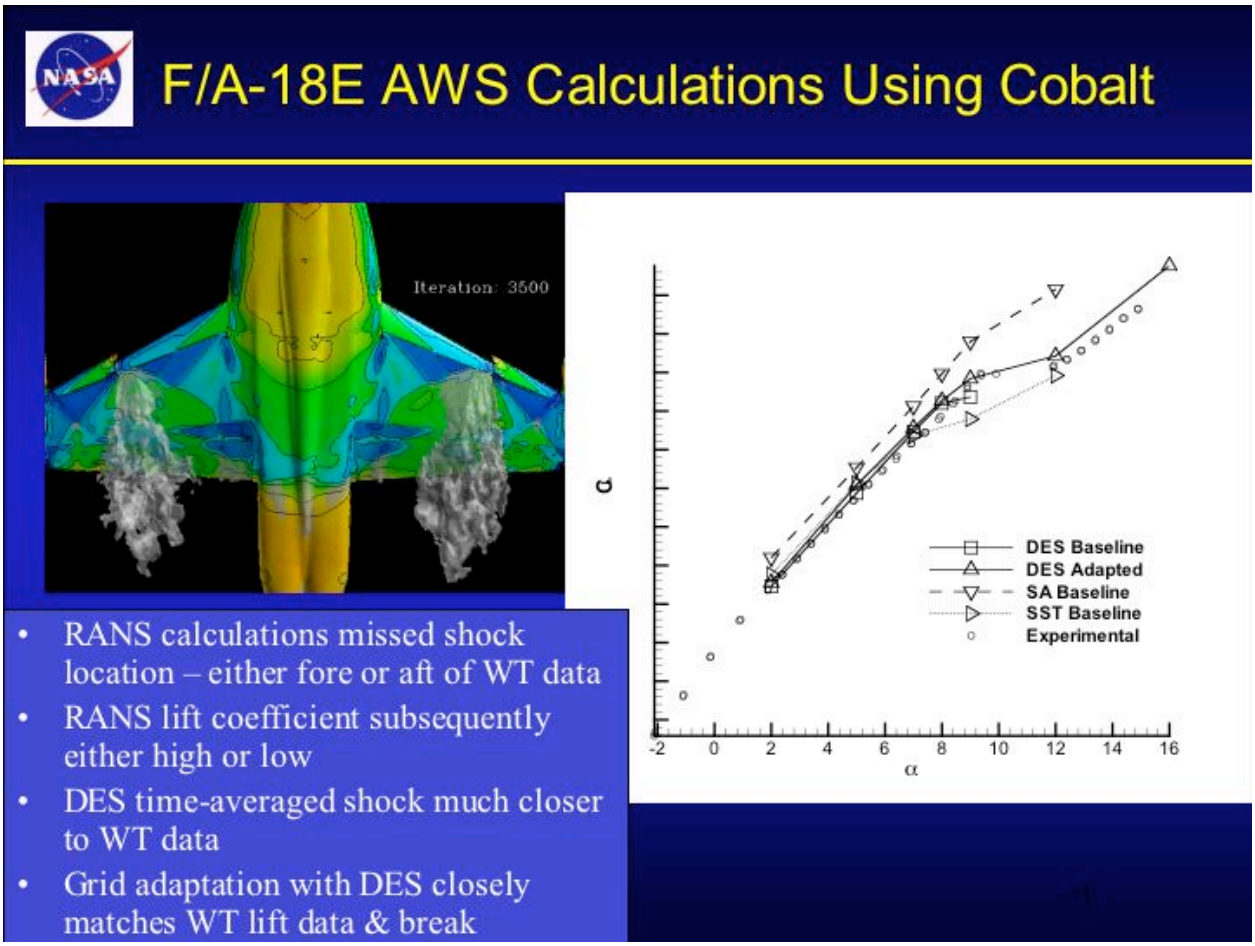

Subsequently, the DES approach was employed, using the SA model near the body. Time averaged output showed a more accurate prediction of the shock location compared to the standard RANS model with SA. Improvement was also seen in the lift coefficient near the break, but the break occurred too early. Finally, the grid was solution-adapted to the initial DES result, and the DES solutions were re-run. Time averaged lift coefficients from the DES simulations showed excellent agreement with the data, predicting the break near nine degrees correctly, as well as the correct reduction in slope, and the subsequent increase in lift curve slope beyond 12 degrees.

Now that we have shown a sampling of CFD capabilities, including some directly addressing S&C, lets consider some important CFD issues that can directly impact S&C calculations.

Massively separated flows will be quite common for S&C applications. Such flows require a grid of high resolution to adequately capture the slow-moving wake regions, which tends to slow the convergence of the CFD solution even for nominally steady flows. But in reality such flows are usually unsteady and we may need to do time dependant simulations in order to properly capture the relevant effects.

Then the question arises, is URANS sufficient or do we need DES? The preceding slide gave evidence that the DES is required in at least some situations. This leads to greater computational cost.

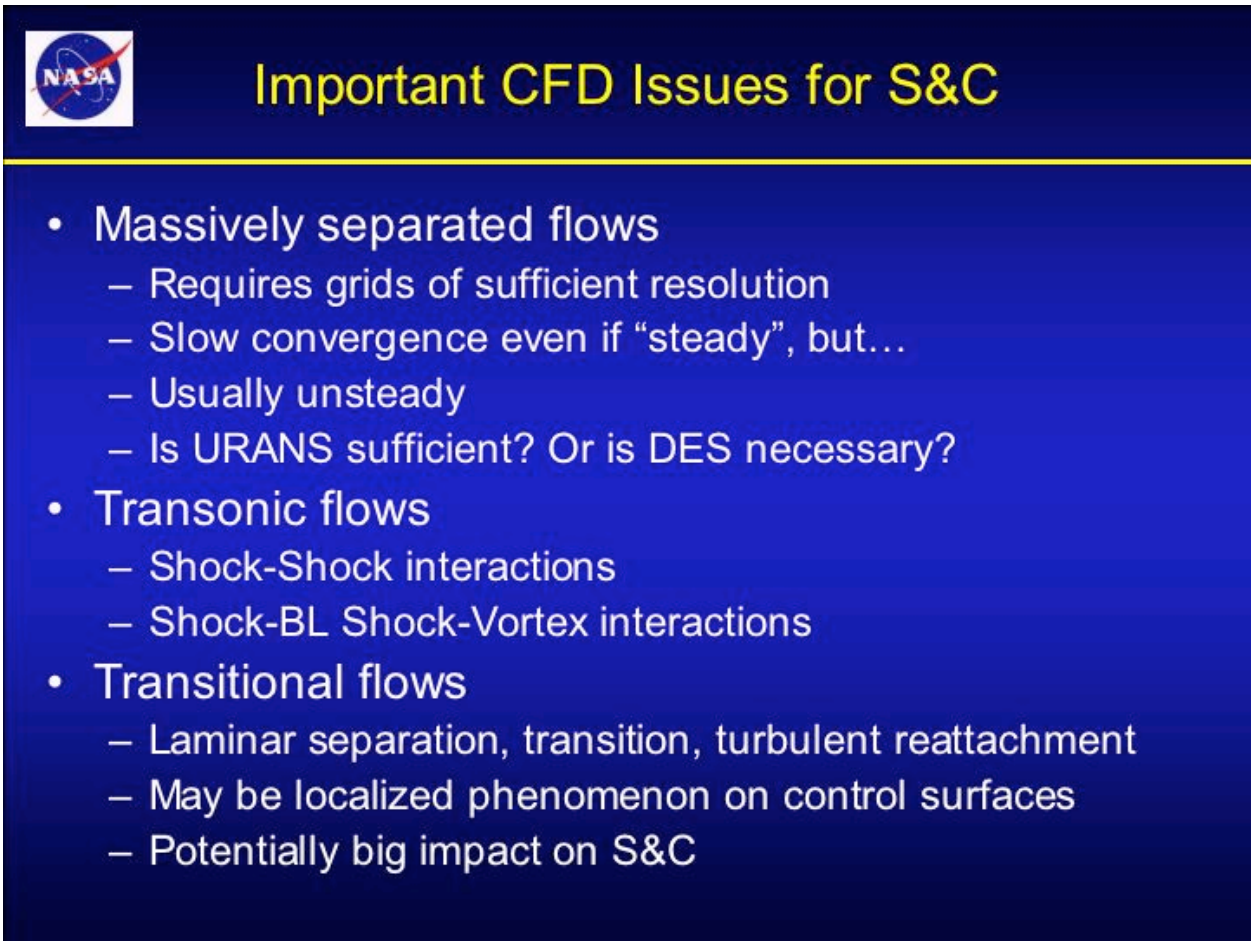

In many cases we will be dealing with transonic flows, and so we may need to resolve complex flows involving shock-shock interactions, shock-boundary layer interactions, and shock-vortex interactions.

Another potentially large issue involves transitional flows, especially ones involving laminar separation, transition to turbulent flow, and subsequent reattachment. Physics-based prediction of transition is not generally available in Navier-Stokes solvers. Even for high Reynolds Number flows based on vehicle length scales, there may be localized transitional phenomenon on control surfaces or near wing leading edges that can have a huge impact on S&C.

We will often have to model the entire configuration. This of course requires double the number of grid points needed compared to situations where symmetry can be assumed.

Such things as differential control surface deflections, sideslip or roll, as well as lateral flow asymmetries arising in a nominally symmetric configuration will dictate a full grid.

Derivatives are often of interest. Thus we need to be able to calculate these reliably for both static and dynamic situations. It should be noted that often derivatives change sign over very small ranges of flow conditions, so simple finite differencing over (say) angle of attack ranges of a degree or so may not yield sufficient accuracy. Other methods for evaluating derivatives, such as complex arithmetic or differentiated source code can provide more reliable derivatives, but are not available in all solvers.

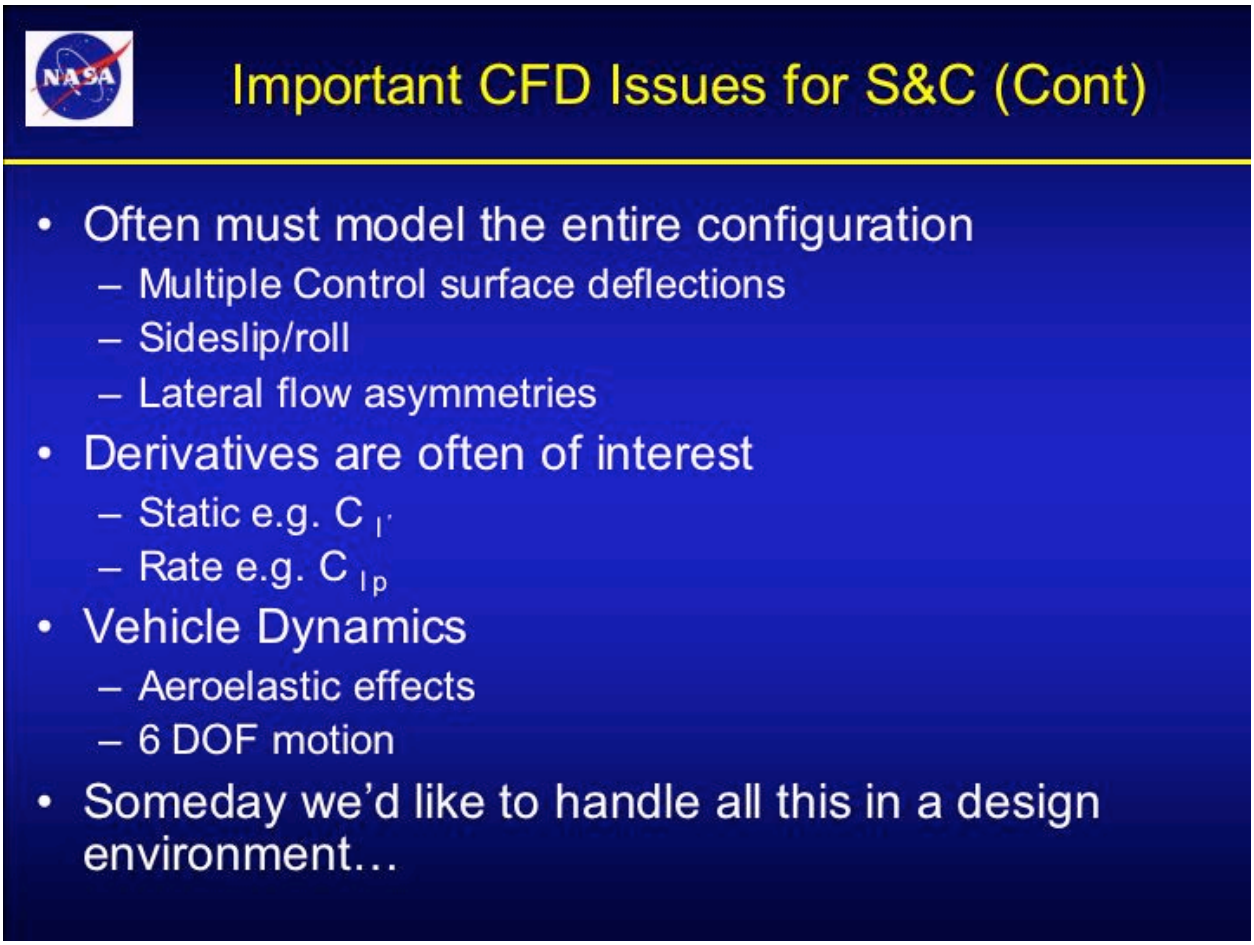

There are many cases where vehicle dynamics are important. These can range from situations in which aeroelastic effects are important to 6 DOF motion.

Finally, at some point many years from now we would like to handle all these situations in a design environment, so that S&C considerations can be designed in from the start.

This slide lists some additional technology barriers that inhibit engineering applications of CFD to S&C.

As things stand today, the whole process of obtaining a CFD solution, but particularly the grid generation aspect, requires a high degree of expertise. This applies to the use of the grid generation software and to the experience required to judiciously place grid points for an accurate CFD result.

The time to obtain a solution is also currently too long for day-to-day engineering calculations. Flow codes are not nearly as efficient at solving the equations as they could be, and the increasing need to simulate unsteady flows just compounds the problem.

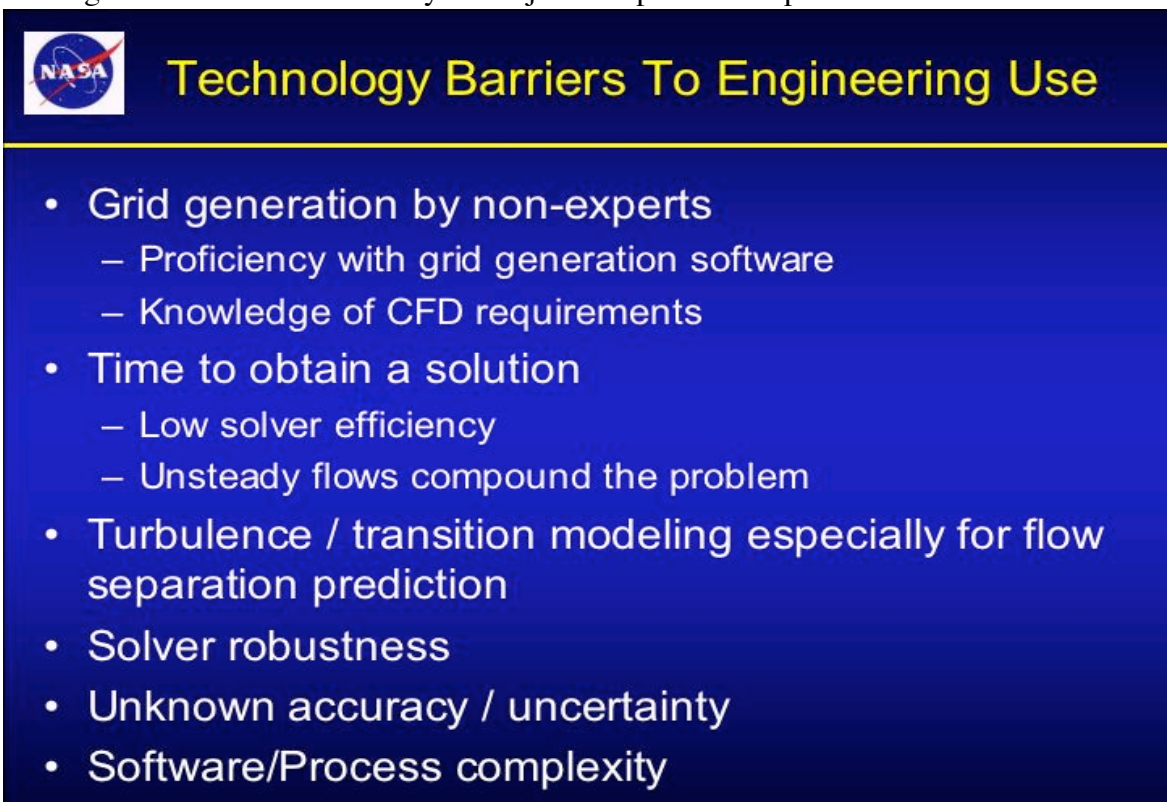

The issues of transition and turbulence modeling, especially for separated flows, is one that causes considerable debate even among CFD experts

Solvers are not robust enough. Not every attempt to get a CFD solution is successful, especially at extreme conditions. Even on a good day, if the solver runs 90% of the cases, the remaining 10% seem to require 90% of the user's time.

When a solution is obtained, there may be questions that arise as to it's accuracy, as well as the uncertainty associated with the result. Mike Hemsch will cover these issues in more detail in a separate presentation.

Finally, the whole process is fairly complex, even for the CFD expert, so much work needs to be done to simplify the process for "routine" engineering applications.

Here we list a few capabilities that are emerging and should have a beneficial impact on S&C applications.

The first is error-based adaptation, which should go a long way to automating the CFD process. We will cover this topic in more detail in the next few slides.

There has been some very promising work to couple high-fidelity, Navier-Stokes solvers with lower order, potential or Euler solvers in a design environment. In the long term, this will allow inclusion of CFD earlier in the design process than is currently practical.

A significant effort is underway to increase the basic flow solver efficiencies toward their full potential. This work is particularly targeted towards the multigrid methods. Impressive results

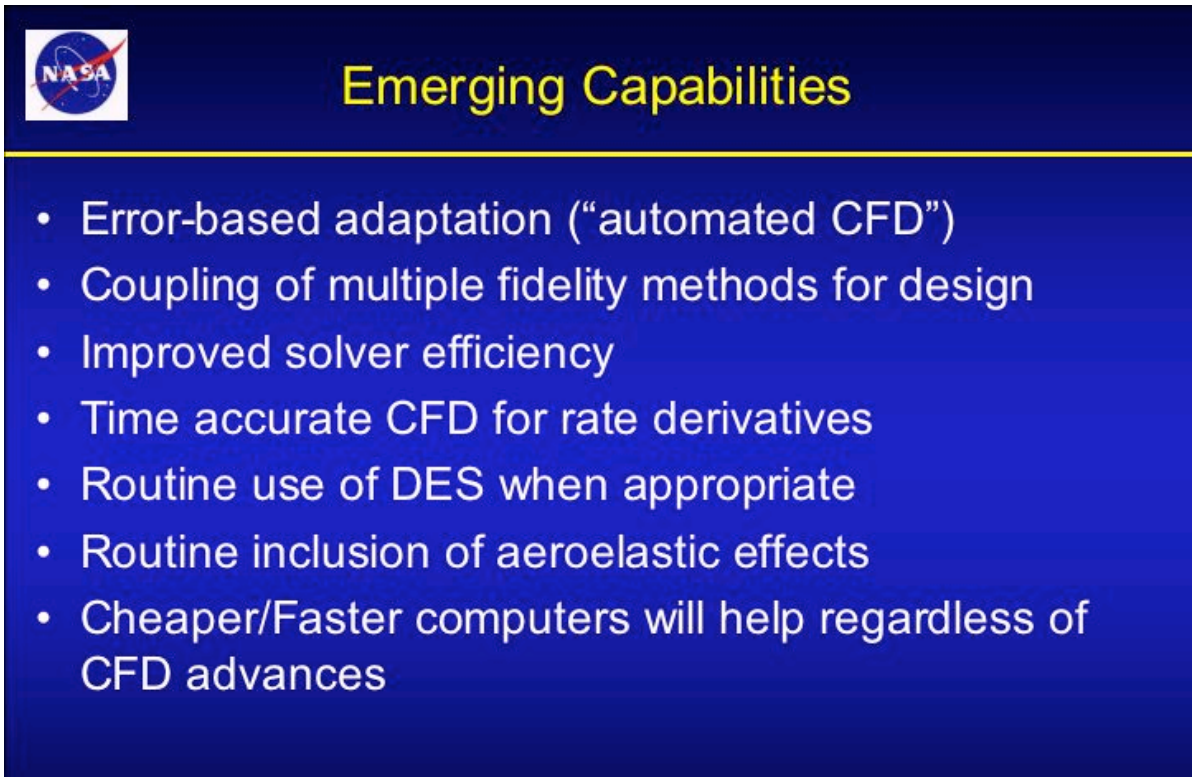

have been obtained for simple flows, but application to complete configurations with complex flows is some years away.

Progress has been made in the computation of unsteady flows, using dual time step methods, as well as the efficient evaluation of rate derivatives for constant rotation rate cases.

The DES method described earlier, has been applied to a wide range of massively separated flows, and may eventually become a routinely used tool for such cases. Likewise, the inclusion of aeroelastic effects, particularly static deformations, is becoming more widely used.

Finally, cheaper faster computers will help bring CFD into the S&C world, regardless of advances on the algorithmic front.

This slide illustrates some trends in computing resources at NASA Langley Research Center – no doubt similar trends can be observed elsewhere.

In the early 1990's we were doing all of out CFD simulations on Cray vector supercomputers. At that time we had approximately 20,000 hours available for the center, and they cost around $100 per hour. By 1995 the SGI Origin class machine was beginning to be widely used, perhaps doubling the available hours and halving the computing costs. Parenthetically it should be noted that a great deal of human effort was required to make the change from the single processor vector machines to the parallel processors of the SGI. Nonetheless, the Origin class machines took over as the primary computing platform by the late 1990's. Both the Cray and the SGI machines were developed with high-end computing as their primary market niche.

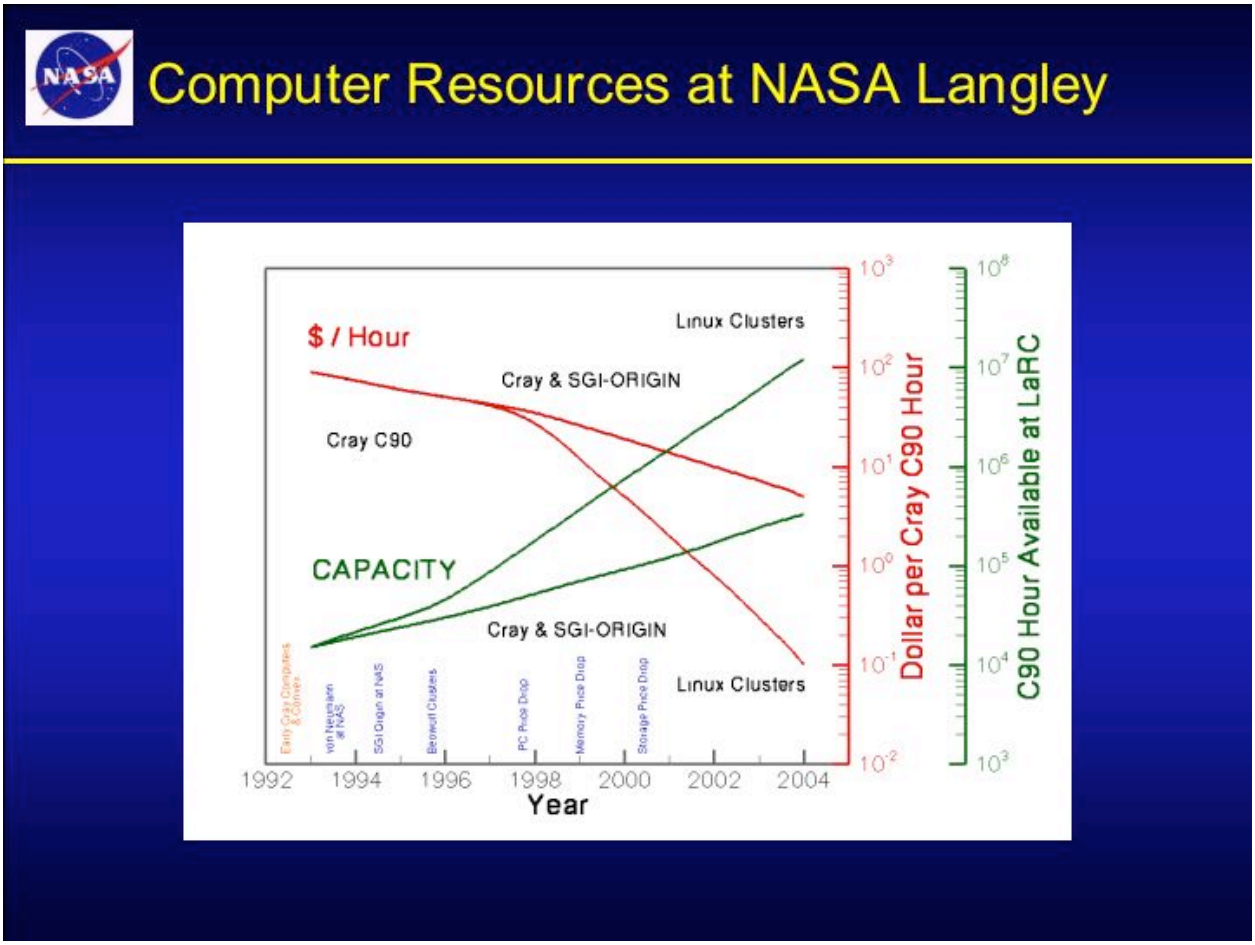

Meanwhile, Linux "Beowulf" clusters of commodity machines appeared, and have been steadily gaining ground. Price drops in CPUs, memory chips and storage have made the Linux clusters very compelling. Having made the effort in converting to parallel processing for the SGI Origin, the effort necessary to adopt the Linux clusters was comparatively minimal. Today there are roughly the equivalent of 10 million Cray C90 hours available, at a cost of about $0.10 per hour.

Now we want to turn attention to a newly emerging technology that shows a tremendous potential as a way to move forward to what might be termed “automated CFD”.

As discussed earlier, current adaptive methods can be somewhat ad-hoc. Recall that the CFD expert must identify the key feature or features to adapt to. Usually, human intervention is required to control the process and decide when to stop. Furthermore, we are still left with the question of what features are important for the problem at hand.

Ultimately, what we want is to provide engineers with a tool that is useful for timely engineering trade studies. So, we must ask, is there a less ad-hoc means of adaptation, and one that can be automated?

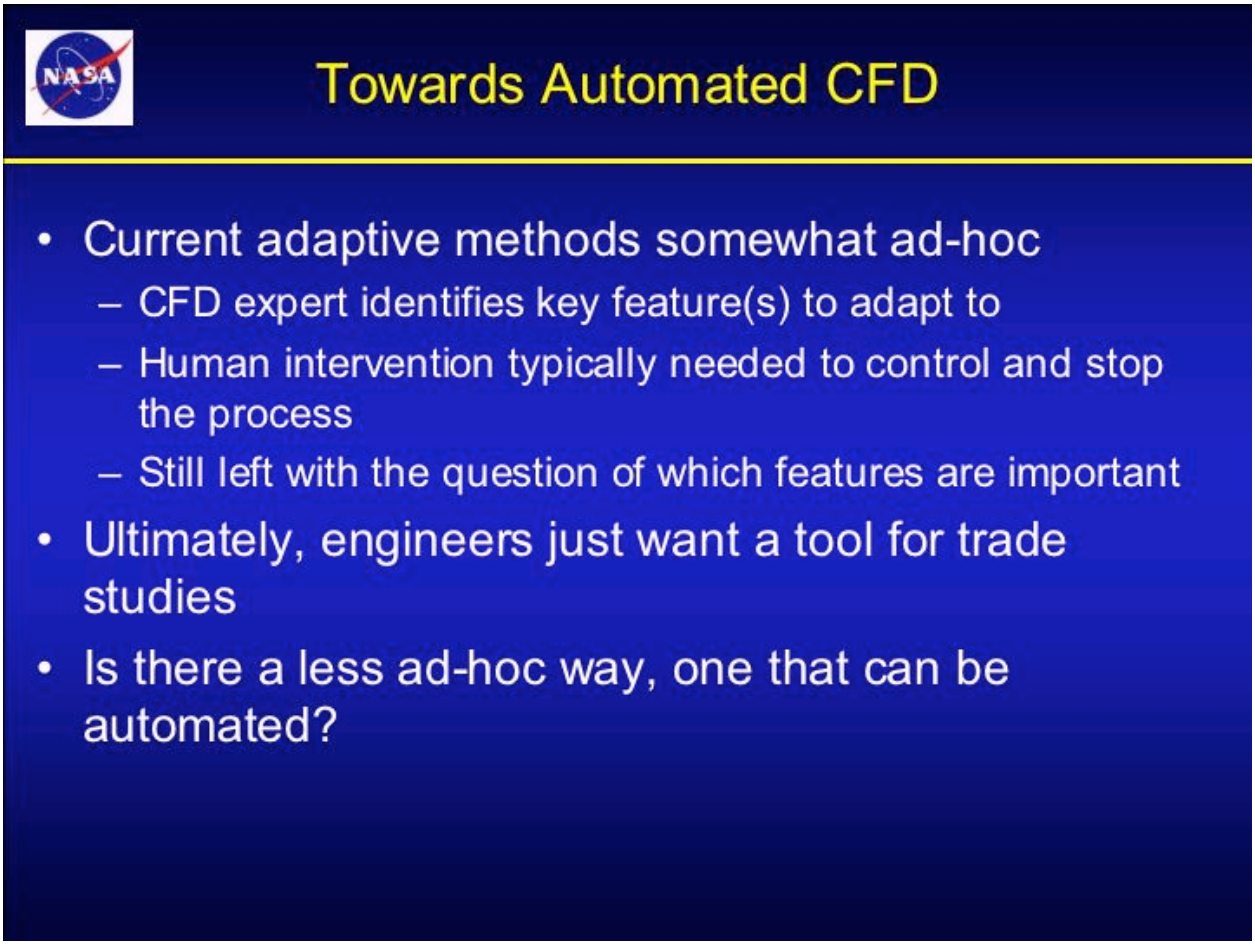

Not surprisingly, we believe the answer is “yes” – error-based adaptation.

Error-based adaptation method described here is the result of some pioneering work in two dimensions by Darmofal and Venditti a MIT. It is currently being extended to three dimensions by Mike Park at NASA Langley. As applied to CFD, the method is quite new; though a similar methodology seems to have been used for some time in computational structures.

This methodology is based upon a solution of the adjoint (dual) equations for the Navier-Stokes or Euler (primal) equations, which is used to determine a computable error estimate. In this approach, the engineer defines an error tolerance on an integral quantity of interest. For example,

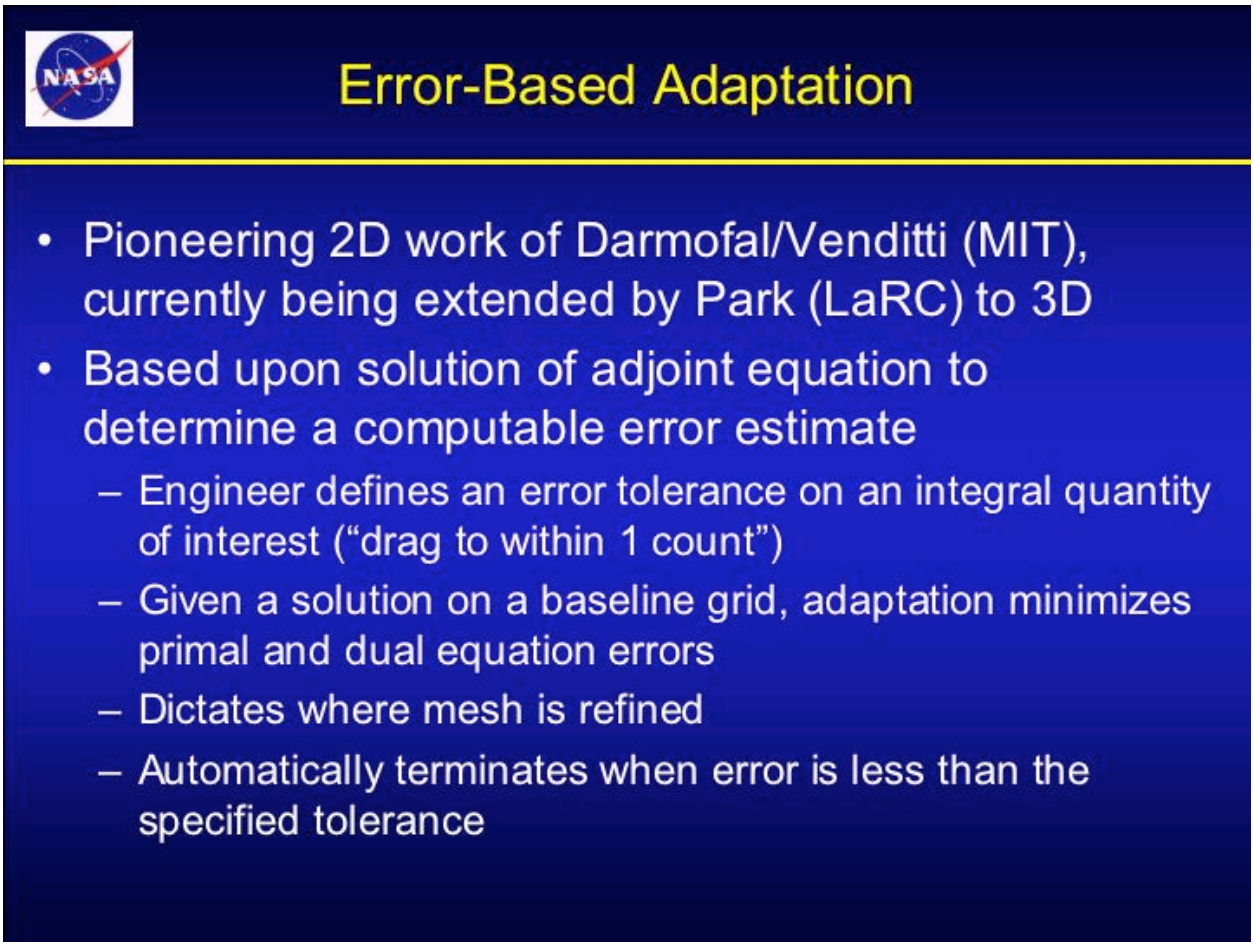

he may want to know the drag to within one count. (It should be noted that “within one count” refers to the best solution that can be obtained given the choice of numerical scheme, turbulence model, geometrical fidelity, etc. – not necessarily to within 1 count of the “true” answer.) Given a solution on a baseline grid, the method adapts the grid to minimize the primal and dual equation errors, and thus dictates where the mesh is refined. Furthermore, it provides a means to automatically terminate the process when the error is less than the specified tolerance.

This slides shows an application of error-based adaptation to supersonic inviscid flow past a staggered pair of airfoils – a Mach 3 biplane if you will.

On the left is the flow pattern that is obtained on a very-well resolved grid. The shock waves and their interactions are all captured quite well, with very sharp resolution of the shocks. On the top right is a solution adapted grid where the pressure gradient was the feature chosen as the adaptation criterion. The final grid contains nearly 38,000 points, and the computed drag coefficient on the lower airfoil drag is 767 counts.

On the lower right is the grid that results when the error-based grid adaptation method is used.

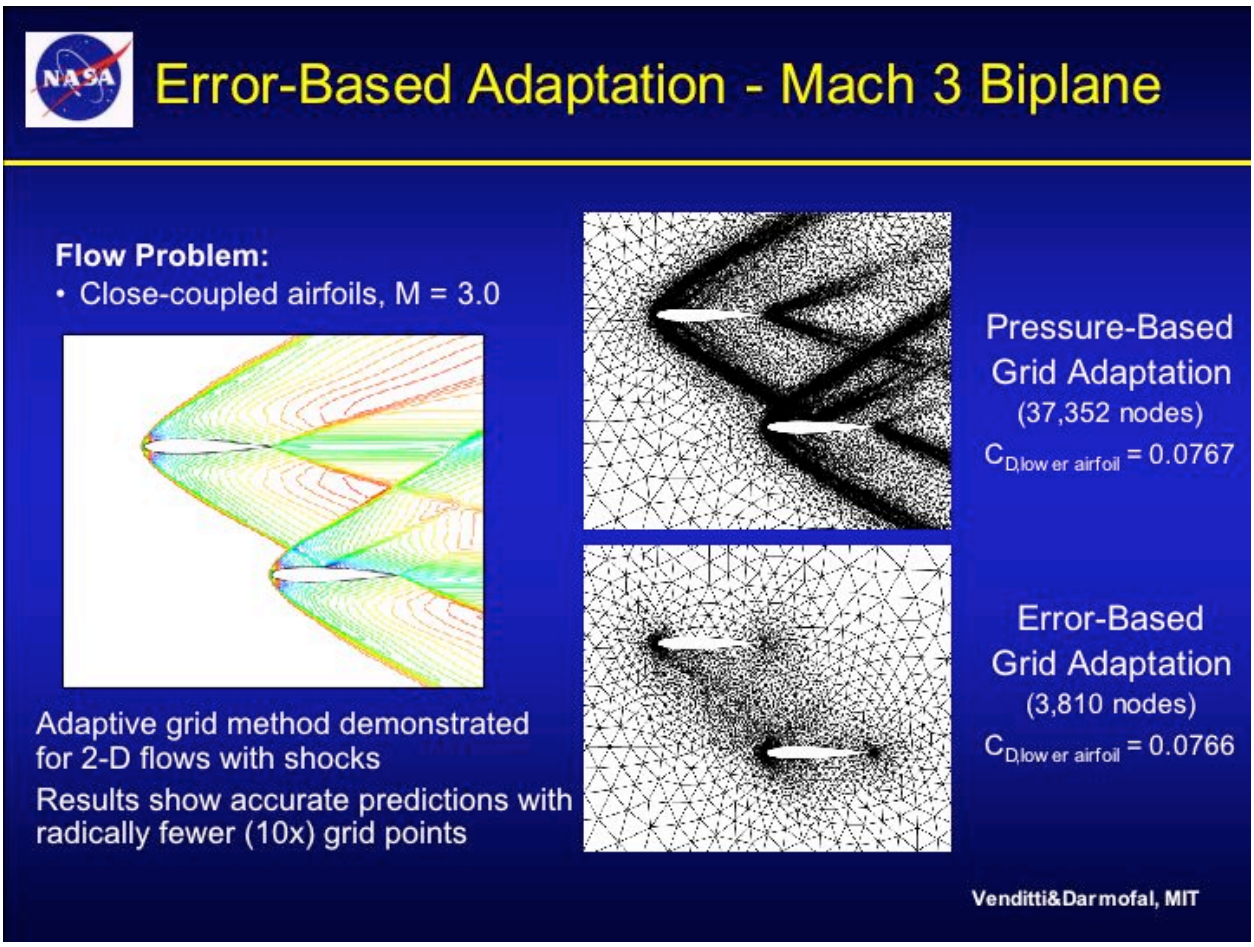

With only 3800 nodes – ten times fewer than the pressure based adaptation – the computed drag coefficient on the lower airfoil is 766 counts.

Notice that the error-based method has not resolved many of the flow interactions that one might think are absolutely necessary to resolve for an accurate prediction of the drag in this supersonic flow. In fact, apart from the leading and trailing edges, the only readily discernable feature that has been adapted to is the part of the bow shock from the upper airfoil that impinges on the lower airfoil.

The relatively simple example on the previous slide suggests that in cases of complex flow interactions, intuitive, ad-hoc adaptation schemes may be quite wasteful of grid points, leading to needlessly long computation times.

The remarkable savings in grid points has been seen in many other cases, including viscous flows. However, we should be careful not to oversell this methodology at this point. It is still evolving, and some significant difficulties lie ahead. For viscous flows, much development work still needs to occur in the adaptation mechanics for highly anisotropic cells. Unsteady flows, which may be quite important in S&C applications, still need theoretical development. The method as developed to date relies on a convergent solution to a steady state. Finally, it should be noted that the development of the adjoint solver is very labor intensive.

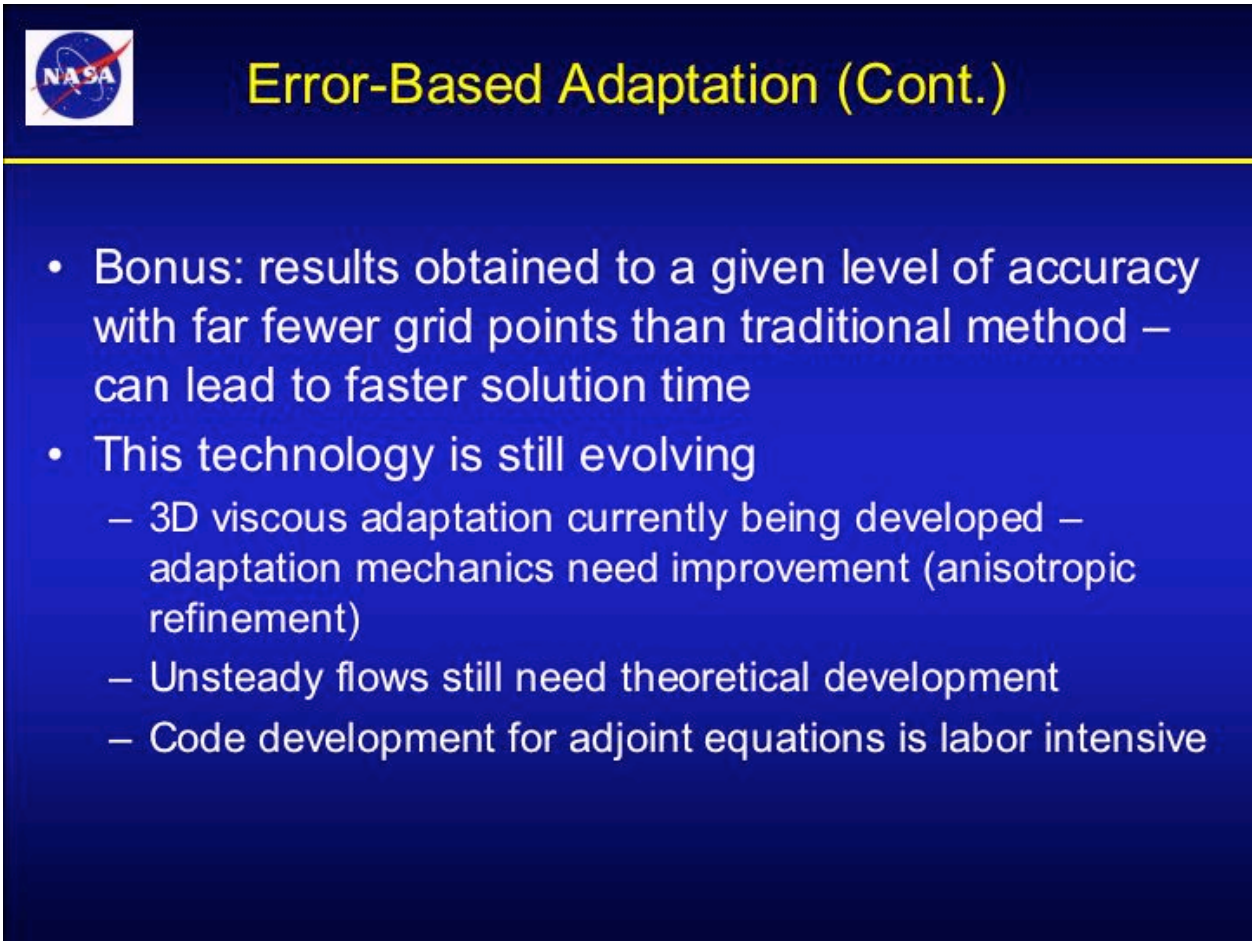

We believe that now is the time to promote a more aggressive use of CFD for S&C. There will be failures – they are to be expected – but that is the only way to make progress.

We feel that a coordinated effort between experiment and CFD is required. In addition to the usual force and moment data that the comes out of experiments geared toward S&C, more detailed information, including flow visualization, pressure data and velocity distributions are needed for CFD calibration. Such coordinated studies should include fundamental studies on simple configurations – to allow for careful grid and time step convergence studies – as well as studies on complete configurations of interest to industry in order to maintain relevance.

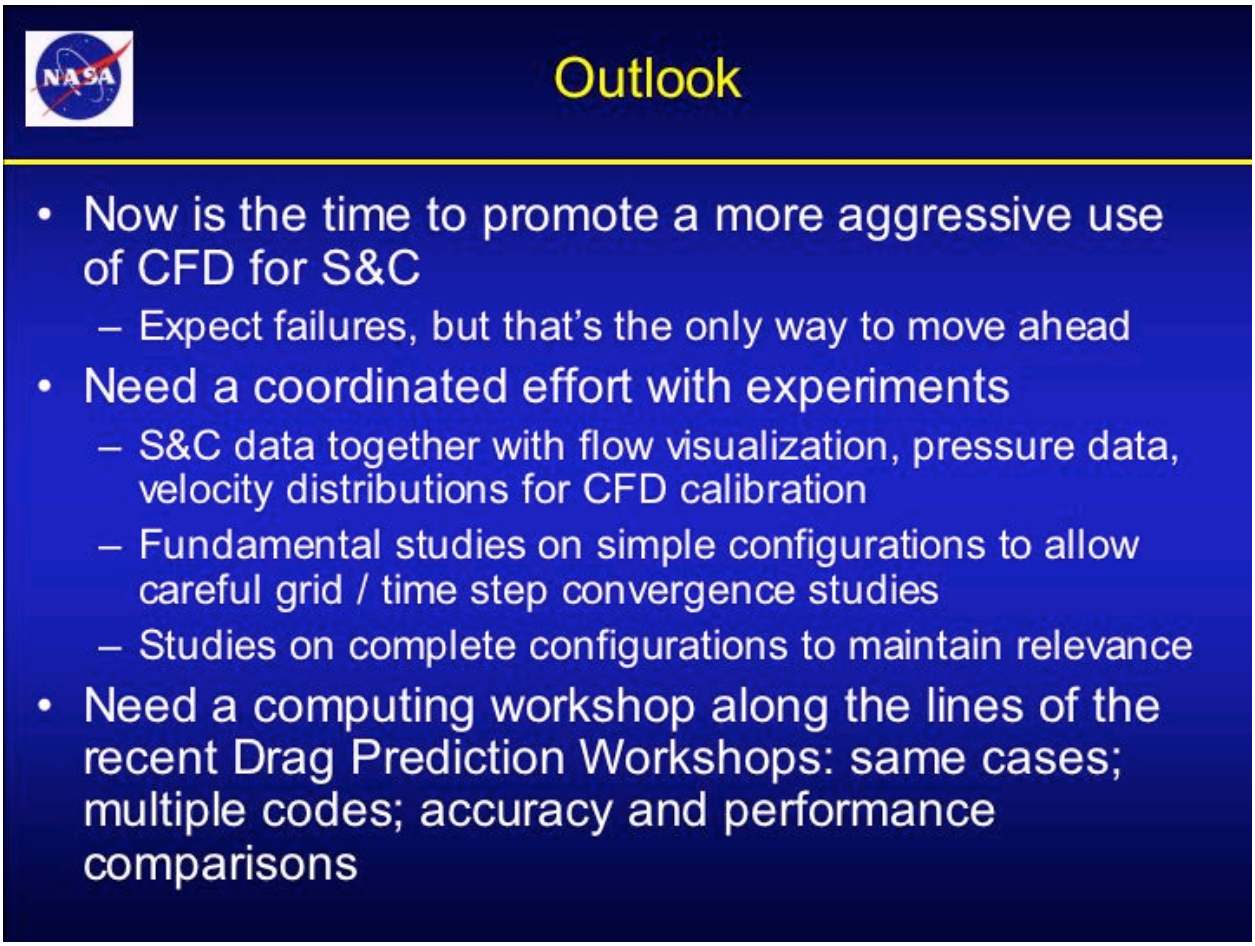

Finally, we believe that a computing workshop along the lines of the recent Drag Prediction Workshops should be held for a problem of interest to the S&C community. In these workshops, multiple codes are applied to the specified configuration (using supplied grids and/or grids of the participant's making), with comparisons made to experiment for assessing accuracy.